



HAL
open science

Modèles mathématiques et méthodes de traitement du signal pour explorer les mécanismes biologiques sur plusieurs échelles : de la dynamique intracellulaire aux séries temporelles neuronales

Matteo Dora

► **To cite this version:**

Matteo Dora. Modèles mathématiques et méthodes de traitement du signal pour explorer les mécanismes biologiques sur plusieurs échelles : de la dynamique intracellulaire aux séries temporelles neuronales. Neurosciences. Université Paris sciences et lettres, 2022. English. NNT : 2022UPSLE033 . tel-04573907

HAL Id: tel-04573907

<https://theses.hal.science/tel-04573907v1>

Submitted on 13 May 2024

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

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

THÈSE DE DOCTORAT
DE L'UNIVERSITÉ PSL
Préparée à l'École Normale Supérieure

**Mathematical models and signal processing methods
to explore biological mechanisms across multiple scales:
from intracellular dynamics to neural time series**

Soutenue par

Matteo Dora

Le 23 mai 2022

École doctorale n°158

ED3C

Spécialité

Neurosciences



Composition du jury :

Thierry Galli

Institute of Psychiatry and Neuroscience of Paris *Président*

Dimitri Van De Ville

École Polytechnique Fédérale de Lausanne *Rapporteur*

Christophe Pouzat

Université de Strasbourg *Rapporteur*

Stéphane Dieudonné

École Normale Supérieure, Paris *Examineur*

Stéphane Jaffard

Université de Paris-Est Créteil *Examineur*

Thibault Lagache

Institut Pasteur *Examineur*

David Holcman

École Normale Supérieure, Paris *Directeur de thèse*

ÉCOLE NORMALE SUPÉRIEURE

THESIS
in Neuroscience

**Mathematical models and signal processing methods
to explore biological mechanisms across multiple scales:
from intracellular dynamics to neural time series**

March 2022

Advisor

David Holcman
École Normale Supérieure, Paris

Candidate

Matteo Dora

Contents

Acronyms	iii
Preface	v
Overview of the contents	vii
I Modelling and analysis of intracellular dynamics	1
1 The endoplasmic reticulum	3
2 The active-flow network model	7
3 Refined properties of the active-flow network	25
4 Quantitative analysis of ER protein dynamics	41
II Neural time series	69
5 Denoising neural time series	71
6 Artefact removal for clinical electroencephalography	97
7 Properties of the wQN algorithm	117
Concluding remarks	127
Bibliography	131

Acronyms

AEEG ambulatory electroencephalography.

AFN active-flow network.

BIS bispectral index.

BM3D block-matching and 3D filtering.

BSS blind source separation.

CCA canonical correlation analysis.

CNN convolutional neural network.

DCT discrete cosine transform.

DWT discrete wavelet transform.

EEG electroencephalography.

EMD empirical mode decomposition.

EMG electromyogram.

EOG electrooculogram.

ER endoplasmic reticulum.

ERAD endoplasmic reticulum associated degradation.

FLIP fluorescence lifetime in photobleaching.

FRAP fluorescence recovery after photobleaching.

FWHM full width at half maximum.

GECI genetically encoded calcium indicator.

GEVI genetically encoded voltage indicator.

ACRONYMS

- GFP** green fluorescent protein.
- ICA** independent component analysis.
- IMF** intrinsic mode function.
- MAD** median absolute deviation.
- MFPT** mean first-passage time.
- MLI** molecular layer interneuron.
- MSD** mean-squared displacement.
- MSE** mean-squared error.
- NMSE** normalized mean-squared error.
- PA** photoactivation.
- PCA** principal component analysis.
- PMD** penalized matrix decomposition.
- PSF** point spread function.
- ROI** region of interest.
- SIM** structured illumination microscopy.
- SNR** signal to noise ratio.
- SPT** single-particle tracking.
- SVD** singular value decomposition.
- ULOVE** ultrafast local volume of excitation.
- UPR** unfolded protein response.
- VST** variance stabilizing transformation.
- WQN** wavelet quantile normalization.
- WT** wavelet transform.

Preface

In his 1960 article *The Unreasonable Effectiveness of Mathematics in the Natural Sciences*, physicist Eugene Wigner wrote that “the miracle of the appropriateness of the language of mathematics for the formulation of the laws of physics is a wonderful gift which we neither understand nor deserve”. He cautiously limited his observations to the laws of physics, but since then, the language of mathematics seems to have irreversibly pervaded all branches of knowledge, sprinkling applications in essentially every field. While coherent mathematical descriptions comparable to those established in physics remain confined to the inanimate world, mathematical methods have become an essential ingredient also in the study of living systems. We do not know if mathematical theories will ever work with biology as well as they did with physics. Perhaps they will—even if I do not believe it. However, to face the new complexity of challenges encountered in the study of biological systems, every tool in our toolkit is required, and no progress can be made without an interdisciplinary approach that includes, among others, mathematical and computational methods. In his essay, Wigner compares the route to modern physics to that of a man with a bunch of keys who has to open a series of consecutive doors, and always guesses the right key on the first or second trial. I like to think that in biology, instead, each door is secured by multiple locks and latches, which have to be all simultaneously unlocked if we want to go through. During my doctoral studies I tried to forge—or more simply to find in the bunch—the keys to open some of these locks. Coming from a background in computer science and theoretical physics, mathematical and computational methods have been my preferential tools. I must add that these keys are quite worthless by themselves, given that unfastening a single lock is not sufficient to unlock the door. Hopefully though, coupled with others, they will contribute to find a way through.

Overview of the contents

In this dissertation, I investigate biological phenomena related to the brain through mathematical models and signal processing methods. The leitmotiv, constantly reappearing through the chapters of this work, is the analysis of spatiotemporal series which naturally arise when analysing out of equilibrium biological systems at different scales. The thesis is organized in two parts, going from the finest to the coarsest scale.

In the first part of the thesis, I present my work on the analysis of intracellular protein dynamics in the endoplasmic reticulum (ER), an organelle found in all living eukaryotic cells formed by a network of tubular membrane structures. The ER plays a key role in protein transport, and its dysfunction has been associated with numerous diseases, including, in particular, neurodegenerative disorders such as Alzheimer's and Parkinson's. The ER also interacts with essentially every other cell organelle, regulating their biological functions. Understanding how molecules are redistributed over time within the network is thus crucial to determine how their functions are achieved. Just before the beginning of my doctoral studies, new experimental observations based on the analysis of single-particle tracking (SPT) suggested that transport of proteins in the ER lumen could be characterized by regions of active flow (along the network edges) and regions of confined motion (in the network junctions). This constituted a significant discrepancy with respect to previous studies based on fluorescence recovery after photobleaching (FRAP), in which the ER dynamics was successfully modelled as a passive, diffusive process.

Thus, in the first part of my doctoral studies, I set out to resolve this inconsistency. Based on the elementary principles of directed flow and confined motion that had been revealed by SPT experiments, I introduced a new mathematical model to describe the ER protein dynamics, which I present in detail in chapters 2 and 3. To emphasize its difference from passive diffusion, we named this model the active-flow network (AFN). In its essence, the AFN consists in a random walk on a directed graph, where the edges change their direction at random time intervals. This constitutes a major difference from previously studied graph models, since the trajectories of the particles moving on AFN are affected by the continuous changes in the edges directions. While random walks and random graphs have been studied separately, the AFN model is, to my knowledge, a unique case of random motion happening on a stochastic topology, where the two processes evolve simultaneously. An in-depth study of the model using both mathematical analysis and numerical

simulations allowed me to reveal a few interesting mechanisms which are characteristic of the AFN. The most important of these is the emergence, for certain parameter values, of synchronization between the motion of independent particles on the network, which results in transportation by aggregates of particles (*packets*). This behaviour, which is also preserved at steady-state, differentiates the AFN model from a passive, diffusion-like network. However, I was also able to prove that, for a wide range of parameters values, the AFN can redistribute material with a timescale very similar to that of a passive network. Moreover, I showed that at a coarse-grained resolution, the AFN can be well described by an effective diffusion model. In summary, with the AFN, directed flow and transport by packets can coexist with a coarse-grained diffusion-like behaviour, making the model a good candidate to reconcile the conflicting experimental findings.

After this theoretical chapters, in chapter 4 we turn to the analysis of experimental data to further investigate how molecules are spatially distributed within the ER network. Protein dynamics within the ER has been traditionally characterized by experimental techniques such as FRAP and fluorescence lifetime in photobleaching (FLIP). However, these methods are limited in terms of temporal or spatial resolution, and do not allow to simultaneously track high speed dynamics across the entire cell. Single-particle tracking, on the other hand, can provide a high resolution measure for isolated molecules in the periphery of the ER, but cannot deal with denser ER regions or with ensemble behaviour. The morphological characteristics of the ER thus pose a challenge to the quantitative analysis of global protein redistribution. To overcome these issues, I have developed a novel analysis pipeline based on photoactivation fluorescent imaging data provided by my collaborators¹. The pipeline couples statistical analysis with machine-learning aided deconvolution and segmentation, to perform a joint analysis of the ER morphology and dynamics. This new tool enabled, for the first time, extraction of quantitative information on the dynamics of ER-localized molecules at the whole-cell scale and with high temporal resolution, allowing tracking of ensemble behaviour both in ER lumen and membrane. By partially deconvolving the effect of the ER morphology, I could show that diffusion is the primary mode of motion in the bulk of the network, and highlight clear differences in the redistribution timescales depending on size of proteins, their topology, and their folding state. Surprisingly, the analysis revealed heterogeneous regions in the material redistribution, suggesting that, despite its continuity, the ER network is not uniformly connected. Overall, these results provide a new insight in the biophysical nature of the ER.

In the second part of my thesis, we turn to biological problems at a higher scale. Capitalizing on the experience with image and signal processing techniques that I gained in my study of the ER, in chapter 5 I present my work on a new technique to reduce photon-counting noise in neural time series acquired by optical recording. These recordings were collected by my collaborators² using a recently developed experimental method which leverages on special proteins, called genetically encoded voltage indicators (GEVIS), capable of changing their fluorescence level

¹I need to thank Christopher Obara, from Jennifer Lippincott Schwartz' group (HHMI, Janelia Research Campus, Virginia), for his awesome experimental work on the ER.

²My gratitude goes to Vincent Villette, from Stéphane Dieudonné's group at École Normale Supérieure, for collecting these impressive optical recordings, and to Stéphane Dieudonné himself and Jonathan Bradley for their support.

depending on the neuron membrane potential. When coupled with advanced optical techniques, this gives the ability to record single-cell voltage signal in awake, behaving animals. The method could provide unique insights in the properties of neuronal microcircuits, but it is currently limited to recording of a maximum of two to three cells simultaneously, since the signal to noise ratio (SNR) tends to degrade quickly as more regions are scanned. To address this issue, I developed a new denoising technique which takes advantage of recurrent patterns in the signal to improve the signal to noise ratio. The method builds on top of the well-known block-matching and 3D filtering (BM3D) algorithm from image processing, readapting it, by appropriate refinement of the filtering and block matching strategies, to work with highly irregular time series characterized by spike patterns. By coupling this technique with robust noise model estimation, I could achieve a 10-fold reduction of the noise variance, opening new possibilities for the simultaneous recording of the activity of an ensemble of cells.

In the final chapters 6 and 7, we move up one last level in the biological scale, as I tackle the problem of artefact removal from electroencephalography (EEG) recordings that capture the activity of the entire brain. Motivated by applications in the clinical context, especially for the monitoring of general anaesthesia, I sought a method to attenuate artefacts in EEG recordings with just one channel (or very few), which constitute the norm in intensive care units and operating rooms. To this aim, I proposed a novel wavelet-based technique which works by remapping wavelet coefficients according to a reference distribution extracted from portions of nearby uncontaminated EEG signal. Despite its relative simplicity, I could show that this technique can achieve better results than state of the art single-channel artefact removal techniques for various types of artefacts, and be at the same time computationally efficient. In chapter 7, I present some preliminary results showing how the method can better preserve the signal spectral features compared to traditional wavelet thresholding techniques widely employed to remove EEG artefacts.

In conclusion, this thesis illustrates how an interdisciplinary approach combining experimental data with mathematical modelling and signal processing can provide new tools for the understanding of a wide variety of biological mechanisms, ranging from protein transport in the ER to clinical EEG monitoring. Especially, denoising and deconvolution can push our capacities a little forward, allowing to see and measure what was before invisible to the naked eye, extracting meaningful information from apparently erratic and confusing measurements.

Schematic summary of the results

Part I: Intracellular protein dynamics

- Introduction of the active-flow network (AFN), a novel graph-based stochastic model to describe motion on the ER network. The AFN is a unique example of a random walker moving on a constantly reshaping graph topology. *(Chapter 2)*
- Emergence of *packet* transport in the AFN, i.e. a mechanism allowing the redistribution of particles in synchronized groups, which results from the topological characteristics of the AFN. *(Chapters 2 and 3)*
- Effective diffusion dynamics of the AFN at a coarse-grained scale which makes it a good candidate to reconcile the discrepancy in experimental findings (active flow detected via SPT, passive diffusion studied with FRAP). *(Chapter 3)*
- Introduction of a new analysis pipeline for photoactivation data of proteins in the ER, which couples deconvolution of the morphology with analysis of protein dynamics at the whole-cell level. *(Chapter 4)*
- Brownian motion signature and quantitative estimation of the timescales associated to redistribution of ER-localized proteins, as a function of their size, topology, and folding state. *(Chapter 4)*
- First quantitative observation of spatial anisotropy in protein redistribution on the ER, suggesting a non-uniform connectivity of the ER network. *(Chapter 4)*

Part I: Neural time series

- Introduction of a new method for the reduction of photon-counting noise in optical recording of single-neuron membrane potential, allowing to achieve a 10-fold reduction of the noise variance. *(Chapter 5)*
- Introduction of wavelet quantile normalization (WQN), a novel technique for the suppression of artefacts in EEG recordings, suitable for single-channel signals acquired in the clinical setting. The technique shows state of the art performances at the removal of artefacts of different nature. *(Chapter 6)*
- Comparison of the WQN technique with traditional wavelet thresholding, showing how the proposed technique can better preserve spectral features, and an example of its application to the monitoring of general anaesthesia. *(Chapter 7)*

PART I

Modelling and analysis of intracellular dynamics

CHAPTER 1

The endoplasmic reticulum

This journey starts with the endoplasmic reticulum. The endoplasmic reticulum—or ER for short—is an organelle of the eukaryotic cell with a fundamental role in managing the cell logistics. It consists in a network of tubular membranes which extend from the nucleus to the cell periphery, and it is associated to multiple functions, including transport and storage of proteins, lipids, and calcium ions. In particular, newly synthesized proteins produced by ribosomes are inserted in the endoplasmic reticulum (ER) network where they are then sorted, filtered and transported to their final destination. In the ER lumen, proteins also find a suitable environment for maturation: lumen-resident enzymes catalyse the chemical transformations needed to stabilize the structure and chaperones facilitate the folding process which brings the proteins into their final, fully functional form [Alberts et al. 2015]. Thanks to a system of biochemical tags, proteins in the ER can be recognized and filtered based on their level of maturity and final destination. Some of them, necessary to the functioning of the ER itself, are retained in the lumen. Others, addressed to different cell organelles, are packaged and delivered to their final destination via transport vesicles. This powerful logistics system is not limited to simple delivery, but also provides a quality control on the exported proteins, as the exit from the ER network is highly regulated. Misfolded or not fully assembled proteins, potentially harmful for the cell, are retained in the ER, where chaperones guide their folding. If this process fails, proteins are marked for degradation and redirected to a proper disposal pathway such as endoplasmic reticulum associated degradation (ERAD) [Alberts et al. 2015; Meusser et al. 2005]. Decades of research have revealed the central role of the ER in these essential functions to the cell. Yet, little we know about the exact mechanisms by which the ER governs the transport, filtering, and export of proteins.

Our interest in the ER is driven not only by sheer curiosity, hoping to unravel some of the secrets of this fascinating biological system, but also by the practical implications the ER has on disease pathologies, in particular for neurodegenerative disorders. Given the variety of essential functions that this organelle serves, it is easy to imagine the kind of dire consequences that can result from dysfunction of the ER. Indeed, malfunctioning of the ER is often a question of cell life or death. This




Figure 1.1 The endoplasmic reticulum of a human cell (COS-7), as seen with fluorescence microscopy. The dashed line indicates the approximate boundaries of the cell. **A** This dark region corresponds to the cell nucleus. It is not an empty region though, but as the ER wraps around the nucleus it falls outside the focal plane of the microscope. **B** A region characterized by dense matrices, which appear as continuous sheets of membrane linked by tubules. *Courtesy of Christopher Obara.*

is what happens, for example, when the ER is overloaded with unfolded — or misfolded¹—proteins altering the organelle homeostasis (i.e. its steady state balance), a condition commonly referred to as *ER stress*. When the level of stress becomes unsustainable, indicating that the ER quality control and degradation system cannot cope with the accumulation of unfolded proteins, the cell triggers a mechanism called unfolded protein response (UPR). UPR makes the cell produce more ER to clear out the proteins accumulated in the ER lumen and restore its normal functioning. If the ER fails to recover, its prolonged stress leads to the cell suicide in the form of apoptosis [Alberts et al. 2015; Schröder and Kaufman 2005; Xu 2005].

The scenario described above is relevant for many neurological disorders which involve misfolding and aggregation of proteins into neurotoxic species [Sweeney et al. 2017]. Alzheimer’s, Parkinson’s, amyloid lateral sclerosis (ALS), and prion disorders, among others, are all associated to specific misfolded proteins [Ghemrawi and Khair 2020; Sweeney et al. 2017]. Neuronal cells seem to be particularly sensible to the accumulation of misfolded proteins as they are generally unable to perform mitosis, and thus cannot dilute the unfolded material by cell division [Remondelli and Renna 2017]. These diseases are often linked to increased ER stress and UPR [Ghemrawi and Khair 2020; Remondelli and Renna 2017; Schröder and Kaufman 2005],

¹Notice the small distinction we make between *unfolded* and *misfolded* proteins: with the first we mean polypeptide chains which have not yet folded in their final structure, while with the latter we indicate proteins which are indeed folded, but in the wrong way.

although the exact role in development of the disease played by these conditions remains debated [Hashimoto and Saido 2018; Sweeney et al. 2017].

In this context, shedding some light on the mechanisms at the base of the transport and filtering system in the ER can help to better understand some of the processes associated to these neurological disorders and, in the long term, to devise better strategies for pharmacological treatment. It is, of course, a very long way—and a foggy one. One of the possible starting points in this direction is the characterization of protein dynamics in the ER lumen. This is what we will try to do in the next three chapters. In chapters 2 and 3 we will explore a hypothetical mathematical model of motion which was derived from previous observations on luminal protein dynamics [Holcman et al. 2018]. Then, in chapter 4, we will move to experimental data, presenting a set of new quantitative methods to analyse protein dynamics on the complex and constantly reshaping ER structure. At this stage, however, our discussion has to be understood as pure, fundamental research, with no tangible objective in sight. Despite the awareness of the potential practical consequences, our exploration is chiefly driven by pure curiosity, knowing that, as said by Randy Schekman, “the pursuit of science for its own curiosity-driven thirst for understanding, inevitably has practical applications”.

CHAPTER 2

The active-flow network model

This chapter—and the next one—resumes the modelling work that I have realized at the very beginning of my doctoral studies. Previous experimental observations suggested the presence of alternating flow along the tubules of the endoplasmic reticulum (ER), and of confined diffusion in the network junctions. We thought it would have been interesting to build a model accounting for these two elementary mechanisms, and see how they played out in the grand picture of ER dynamics. I built this simple mathematical model considering a random walk on a directed graph, where the direction of the edges changes randomly over time. Despite its simplicity, the model revealed some interesting and unexpected properties. In particular, a peculiar mechanism of synchronization that causes particles to travel together in groups. We will see, at the end of this first part of the thesis, that the model had some unfortunate fate when faced with the reality of experimental measurement. It remains, nevertheless, a curious and rare study of motion on a randomly reshaping network topology.

This chapter was adapted from “Active flow network generates molecular transport by packets: case of the endoplasmic reticulum”, by Matteo Dora and David Holcman, Proceedings of the Royal Society B (2020).

The ER is a fundamental organelle located inside the cytoplasm, involved in synthesis and export of proteins and membrane lipids across the cell [Alberts et al. 2013]. The ER is formed by a network of membrane sheets, called cisternae, interconnected by membrane tubules. This structure deforms constantly [Lee and Chen 1988], but at steady-state it can be described by a network presenting in most cases three-way junctions (figure 2.1 A) [Nixon-Abell et al. 2016; Terasaki, Chen, and Fujiwara 1986], where each vertex (junction or cisterna) is connected in average to three neighbouring vertices by small tubules (edges).

One of the fundamental roles of the ER is to redistribute proteins, but the exact mechanism and the associated time scales remain controversial. Initial experimental evidences [Nehls et al. 2000] considered that molecules could simply be mov-

ing by diffusion. At the molecular population level, this was studied via fluorescence recovery after photobleaching (FRAP), a technique which consist in exposing a small region of the cell to a high intensity illumination to locally bleach fluorescent molecules, and then monitoring the recovery of fluorescence caused by new material moving in [Nehls et al. 2000]. These experiments were analysed under the assumption that transport is diffusive in the entire network, estimating effective diffusion coefficients in the order of $1 \mu\text{m}^2 \text{s}^{-1}$ [Nehls et al. 2000]. As we shall see below, this diffusion coefficient is sufficient to describe the redistribution in a passive network and the associated timescales.

However, a recent work based on the analysis of single particle trajectories in the ER lumen revealed a different picture, reporting directional flow in tubules and a slower, confined motion in the network junctions [Holcman et al. 2018]. Hozé and Holcman used a stochastic processes approach [Hozé and Holcman 2017] to reveal that a purely diffusive model was not necessarily sufficient to describe the dynamics in the ER lumen: while the motion in nodes was characterized by confined diffusion, with an effective diffusion coefficient $D_{\text{node}} \approx 0.4 \mu\text{m}^2 \text{s}^{-1}$, the motion along tubules was characterized as an alternating flow with mean velocity $v_{\text{eff}} \approx 30 \mu\text{m} \text{s}^{-1}$. This result was further supported by photoactivation experiments, consisting in photoconverting special fluorescent molecules in a small region and following their redistribution, which showed qualitative patterns seemingly incompatible with pure diffusion [Holcman et al. 2018].

To reconcile these experimental results, in this chapter we introduce a network model aiming at describing how independent molecules could be redistributed inside the ER. Our aim will be to define and quantify the timescale of mixing inside the ER lumen. Based on the results from single particle trajectory data, we consider a new transport mechanism on a graph where the direction of the flow along edges is inverted at random times. This model accounts for single molecular trajectories showing randomly alternating directionality in crossing tubules connecting two neighbouring junctions [Holcman et al. 2018]. We call this model the active-flow network (AFN). Using the AFN, we study the consequences of the alternating flow on the redistribution of particles such as ions or proteins inside the ER.

The characterization of redistribution on a network falls into the branch of research studying random walk processes on graphs. However, here we will need to develop a theory to account for the AFN. The timescale of redistribution can be studied by computing the mean time it takes for molecules released from a source node S to reach a target node T located at a distance $d(S, T)$. Evaluation of this mean first-passage time (MFPT) is relevant to quantify how long could take for proteins to reach an exit site of the ER [Kurokawa and Nakano 2019; Kurokawa, Okamoto, and Nakano 2014]. The MFPT has been studied for passive, undirected networks, both for single particles [Bénichou and Voituriez 2014; Sood, Redner, and Ben-Avraham 2004] or multiple particles [Weng et al. 2017], quantifying its dependence on the distance $d(S, T)$ and on the network topology. However, the mean time properties are not known in the case of a network whose topology changes in time (i.e. what we

call an *active network*). Similar active networks have been described in the context of transportation (e.g. railways, vehicular traffic) or communication (e.g. internet, telephony) [Ford Jr and Fulkerson 2015]. In these situations, unidirectional flow emerges as a consequence of the limited capacity of the channel, which does not allow simultaneous bidirectional flow, thus practically affecting the topology of the network by constraining the connectivity. Instead, in the ER, we consider the switching of edge direction to be an intrinsic property of the network, maybe the result of currently undetermined biological machinery. For this reason, we focus here on unidirectional flows which are independent of capacity limitations.

The chapter is organized as follows. First, we present the ER network reconstructed from live cell microscopy. Using this reconstructed network and the AFN properties, we then study the redistribution of molecules across the network using the classical notion of first-passage time [Schuss 1980]. We show that the time of redistribution increases drastically for sparsely connected regions, located at the boundaries of the network. In the second part, we study the AFN on a hexagonal lattice modelled after the ER, and we estimate the timescale of material redistribution for the mean particle and also the fastest particle. In the third part, we show that the AFN can exhibit a novel type of transportation, by redundant *packets* of particles. In fact, under randomly alternating flow, all edges incident to a node could temporarily be inward directed, trapping particles in the node. We call this situation a *capture state*. When the flow of at least one of the edges changes its direction, particles are released. We observe that this mechanism can cause the synchronization of group of particles which start to travel together. In the final section, we study the steady-state and transient properties of AFN, confirming redundant transport in packets, that we suggest is a possible mechanism of transport inside the ER. Thus transport in the ER modelled as an AFN is very different from flows or diffusion in classical networks [Lovász et al. 1993]. The unidirectional property of the network is reminiscent of diode networks, introduced in percolation theory [Broadbent and Hammersley 1957; Redner 1981].

2.1 Results

2.1.1 First-passage simulations in a reconstructed ER network

To study the AFN, we use a graph (figure 2.1 A) which was built based on the ER topology reconstructed from SIM [Young, Ströhl, and Kaminski 2016]. The AFN consists of a graph with N nodes and a Poissonian rate $\lambda_{\text{switch}} \equiv \tau_{\text{switch}}^{-1}$ that controls the switching of the flow direction in each tubule. This random switching process is independent for each edge. This property makes the ER network quite different from classical networks (figure 2.1 B). When the flow in tubules connecting a node to the rest of the network are all inward directed (blue arrows, figure 2.1 B), the particles located inside such a node are trapped.

To evaluate the redistribution efficiency of an AFN, we shall focus on estimating the mean time for a particle starting from a node S to visit the network. We start by running stochastic simulations as follows. Each particle attempts to jump through an outward directed edge with a Poissonian rate $\lambda_{\text{wait}} \equiv \tau_{\text{wait}}^{-1}$. When more than one outward edges are available, a particle can randomly select one of them with equal probability, and moves to the connected node in a time τ_{tubule} . In the following, we will consider that the time spent crossing the tubule is negligible with respect to τ_{wait} , i.e. we set it to zero in our simulations. If a particle is in a node whose incident edges are all inward directed, it is blocked. The particle will thus wait the next Poissonian event to attempt a new jump.


We estimated numerically the mean time for a particle starting in node S to arrive for the first time at node T on the reconstructed ER network (figure 2.1 c). Interestingly, the simulations reveal that the ER network can be divided into three regions (figure 2.1 c). A first subregion consists of nodes close to the source S . For those nodes, the MFPT increases quickly with the graph distance $d(S, T)$. At intermediate distances, the mean time increases in an approximately linear way. Finally, for nodes located far away, with reduced connectivity with the rest of the network, the mean time increases drastically, because it is a rare event to enter these regions. We called these inaccessible regions exponential regions. These regions at the periphery of the ER network are quite inaccessible (figure 2.1 c).

We then studied the effect of the switching timescale τ_{switch} on the MFPT. In principle, due to the presence of capture states, the MFPT in the active network should increase compared to an undirected network. Instead, we found that the MFPT has a minimum at $\tau_{\text{switch}} = 0.25\text{s}$ (figure 2.1 d), making the MFPT timescale very similar to the undirected case for a wide range of switching timescales τ_{switch} . The mechanisms underlying this phenomenon are described in more detail in section 2.1.5.

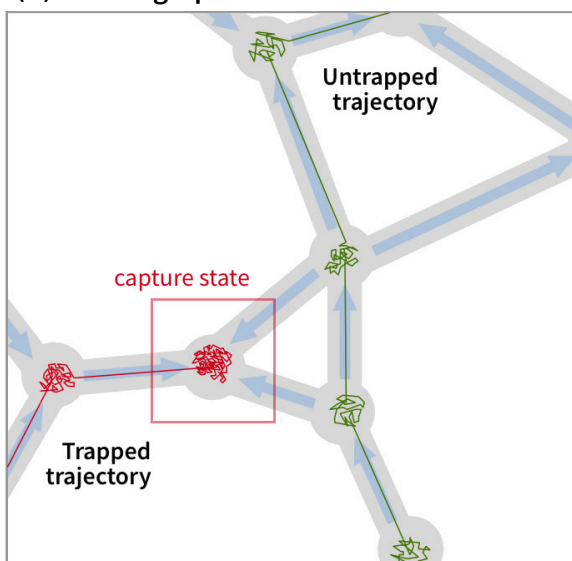
To conclude, we found that an active network (with parameters $\tau_{\text{switch}} = 30\text{ ms}$ and $\tau_{\text{wait}} = 100\text{ ms}$ derived from the results in [Holcman et al. 2018]) can redistribute material with a timescale similar to a classical undirected network. Based on these parameters, the timescale amounts to about 25 min for a particle to arrive to any node, except for the peripheral nodes located in the exponential region (figure 2.1 d). The MFPT depends on the total number of nodes in the network, because on average each trajectory will visit a large portion of them before arriving for the

Figure 2.1 First passage time simulations in an active network. **A** ER network reconstructed from structured illumination microscopy (SIM) data [Holcman et al. 2018]. The average length of an edge is $0.5\ \mu\text{m}$. **B** Schematic representation of an active network model: the flow in a tubule is unidirectional and alternates at random time. A capture state (red box) is defined as a node with all converging arrows (in blue). **C** Heatmap of the mean first-passage time (left) and scatter plot of the MFPT for each node versus the distance from the source S (right), for a switching timescale $\tau_{\text{switch}} = 30\text{ ms}$. A boundary separates the linear and exponential regions (dashed lines). **D** Variation of the MFPT for a node at distance 25 from the source, depending on the switching timescale τ_{switch} . The minimum is obtained for $\tau_{\text{switch}} \approx 180\text{ ms}$. The MFPT for diffusion (dashed line) is $\approx 19\text{ min}$. The network contains 1920 nodes.

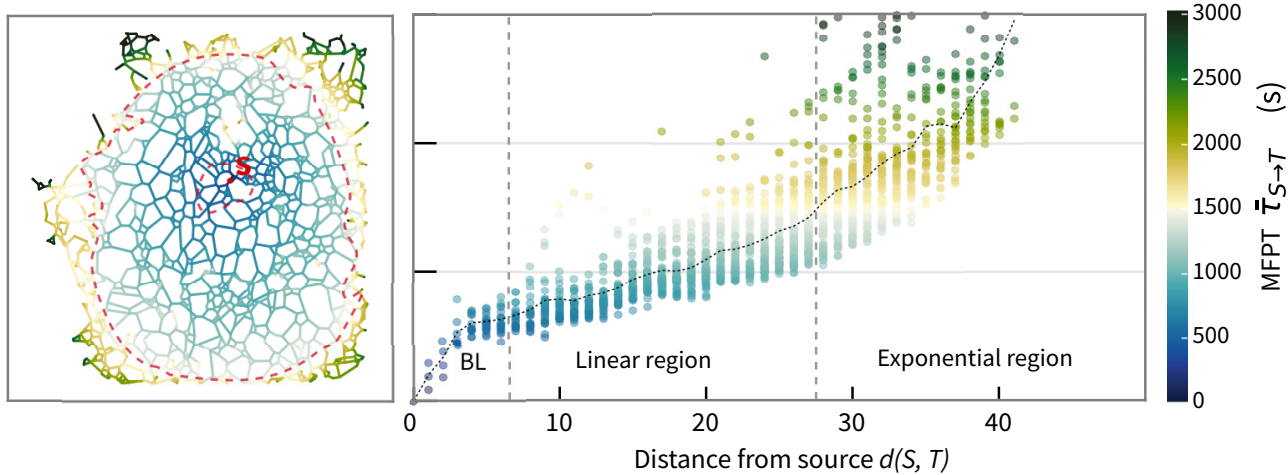
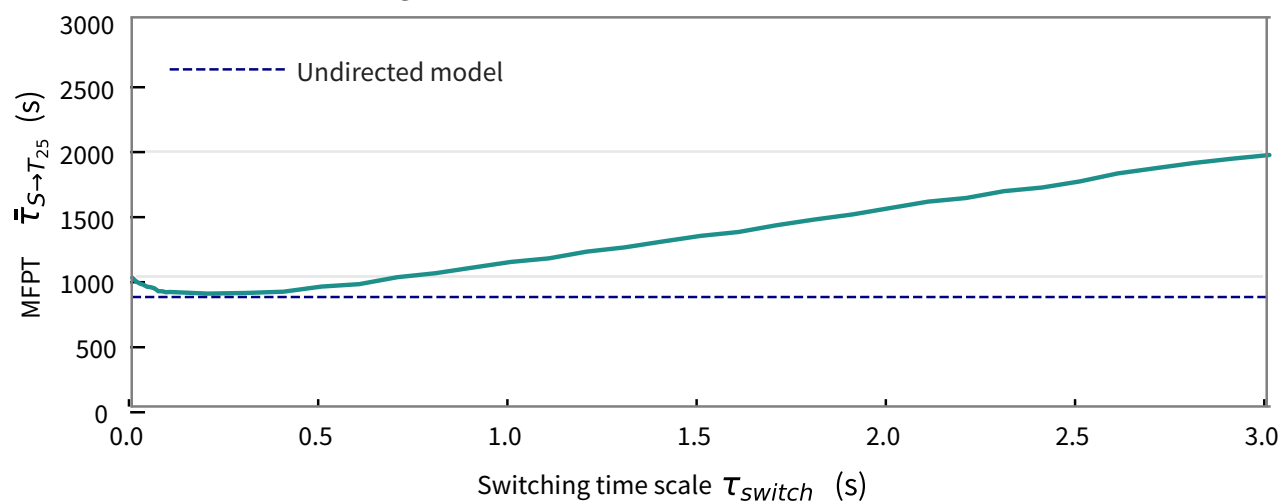
(A) ER network



(B) Active graph



(C) Mean First Passage Time

(D) Effect of switching rate λ 

first time to the target node, a situation similar to the classical escape through a narrow window [Holcman and Schuss 2015]. We next investigate how these properties would apply to a general hexagonal lattice which approximately describes the structure of the ER.

2.1.2 Transport on an equivalent hexagonal lattice

To assess the role of the ER topology on transportation, we tested whether the MFPT could be similar between the SIM imaging reconstructed ER network (figure 2.1) and a hexagonal lattice. We chose a hexagonal, honeycomb network because in average a node of the ER network has three neighbours [Nixon-Abell et al. 2016; Terasaki, Chen, and Fujiwara 1986]. We evaluate the arrival time in two networks: one with a periodic boundary condition and the other with reflecting nodes at the boundary (figure 2.2). Using AFN parameters compatible with the ones used in figure 2.1, we found that the MFPT in the network with reflecting boundary can be similarly decomposed into three subregions, with the exponential regions corresponding to the network corners which are difficult to access. Predictably, in the periodic boundary network, these exponential regions are absent. We conclude that a finite honeycomb network could recapitulate the properties of the AFN ER network.

2.1.3 Transportation of the fastest particles

Various fast molecular mechanisms take place in the endoplasmic reticulum, such as transport of protein, unfolded protein response (UPR), or calcium signalling. In this context where many particles are involved, the statistic of the mean as given by the MFPT may become irrelevant, because extreme events [Schuss, Basnayake, and Holcman 2019] such as the arrival time of the fastest particle to their target is sufficient to activate a cellular process. For example, the fastest folded proteins reaching exit sites with COPII-mediated trafficking [Budnik and Stephens 2009] will be transported to the next organelles. We thus focus here on the arrival time of the fastest particles among many, starting at a single source node S and arriving at a target node T . We recall that, in the absence of edge directionality, an ensemble of particles starting at S will simply distribute by diffusion with a timescale determined by the first eigenvalue of the Laplace operator on the graph [Maier and Brockmann 2017]. After a transient regime, the density of particles converges to the uniform distribution: at steady state, the number of particles per node is simply given by the ratio of the total number of particles to the total number of nodes $N_{\text{particle}}/N_{\text{node}}$. However, for the AFN, the flow in tubules (edges) can create capture states (Fig. figure 2.1 B). By numerical simulations in the honeycomb network (figure 2.2 A), we found that these capture states can create a synchronization between independent particles and also alter the steady state distribution.

We found that the trajectories of the fastest particles concentrate toward the shortest paths between S and T as the total number of particles increases (figure 2.3 A

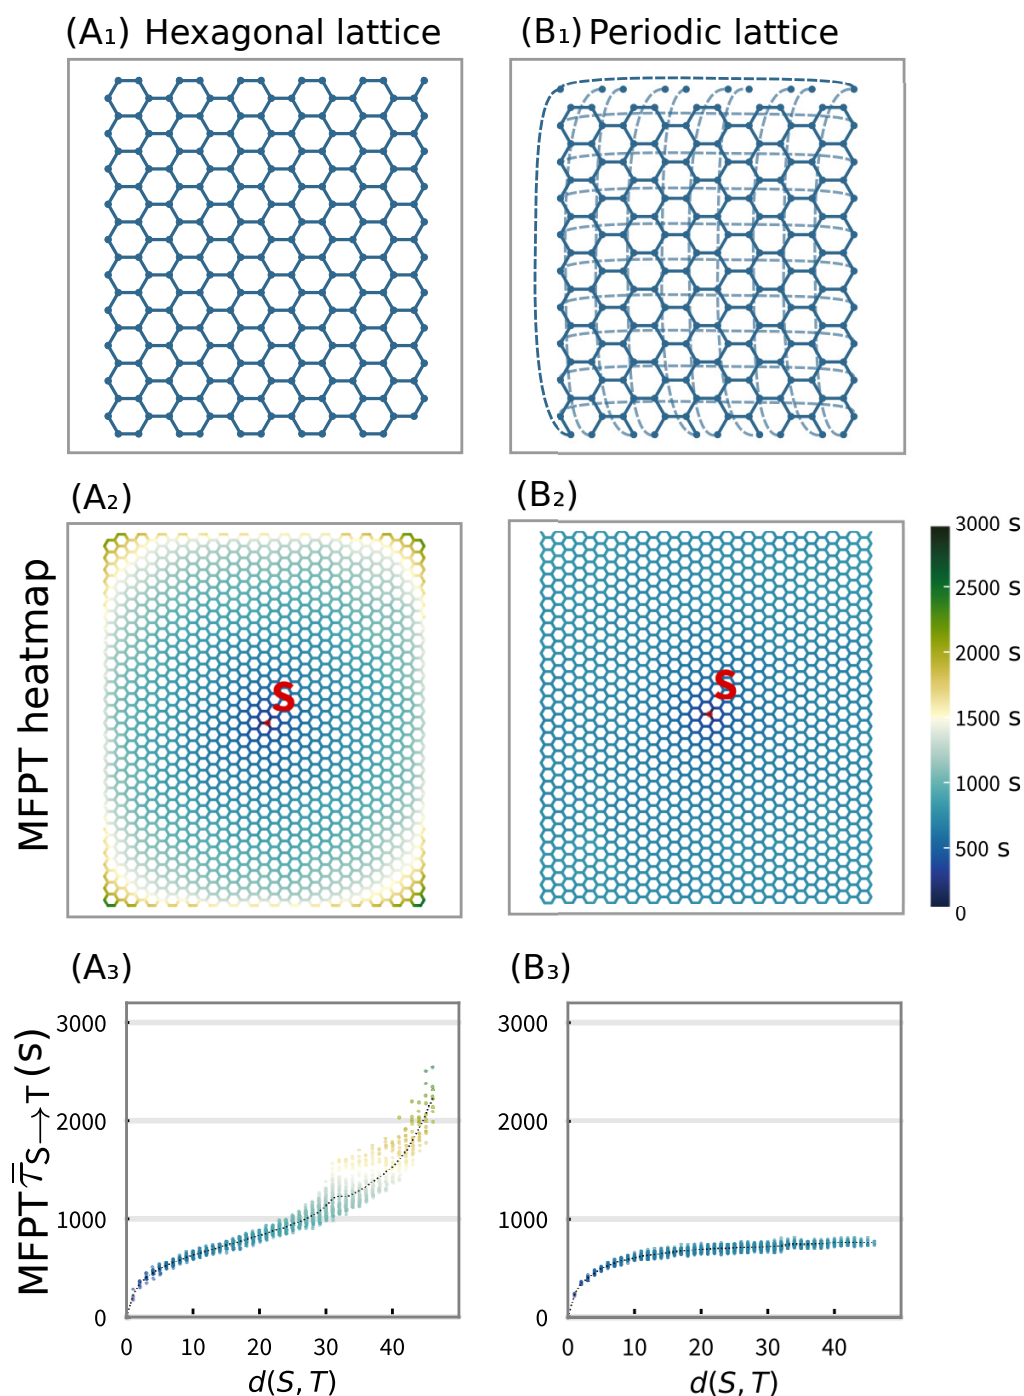


Figure 2.2 Mean first-passage time (MFPT) in a honeycomb network. Hexagonal lattice for a finite (A) and a periodic (B) network. The MFPT is plotted as a function of the distance from the source S to a target T located at distance $d(s, T)$. Parameters are similar to figure 2.1, with $\tau_{\text{switch}} = 30$ ms, $\tau_{\text{wait}} = 100$ ms, and network $N = 2000$ nodes.

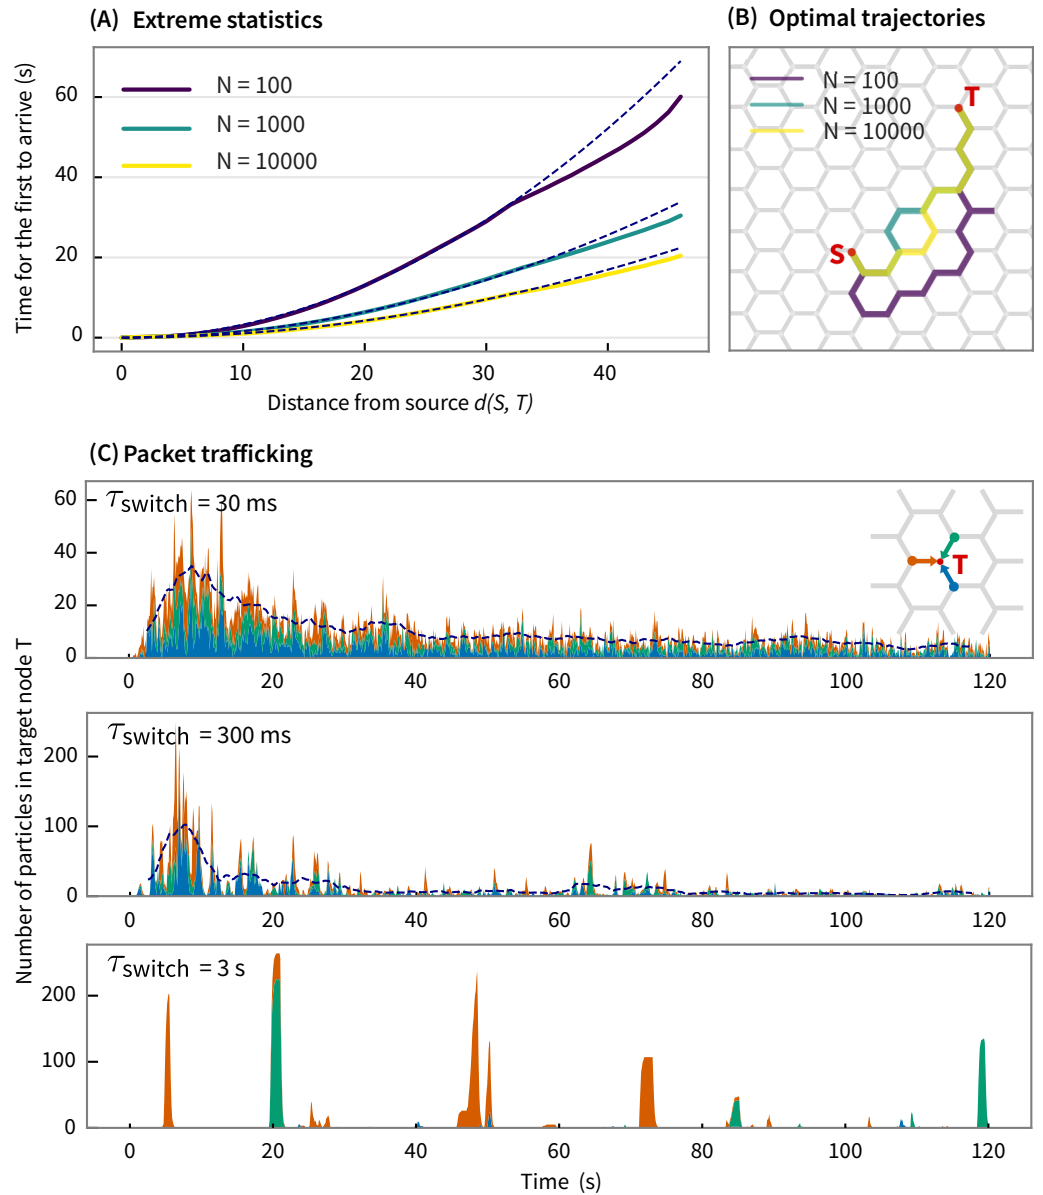


Figure 2.3 Extreme statistics and packet transport mechanism in the AFN. **A** Trajectories of the fastest particles released at S for different number N . **B** Mean arrival time for the first particles at a target node T , for N particles. The stochastic simulations (dashed line) is compared to the fit (continuous line) $\frac{c_1 \delta^2}{c_2 + \log N}$ ($c_1 = 0.07$ and $c_2 = -2.38$), where $\delta = d(S, T)$ is the network distance. **C** Number of particles in the target node coloured by the edge where they are coming from. Top: diffusion-like behaviour (for small switching timescale $\tau_{\text{switch}} = 30 \text{ ms} \ll \tau_{\text{wait}} = 100 \text{ ms}$) leads to a uniform steady-state distribution. Middle: intermediate $\tau_{\text{switch}} = 300 \text{ ms}$. Bottom: large switching timescale $\tau_{\text{switch}} = 3 \text{ s}$ causes groups of particles to arrive synchronously to the node.

and figure 2.3 B, see also figure 3.1 in the next chapter). The MFPT for the fastest among $N \gg 1$ particles can be derived from first principles [Basnayake, Schuss, and Holcman 2018]. We leave the detailed calculations for chapter 3, but provide here a brief summary of the main steps. Considering an effective diffusion process in the network, we determine the distribution of the arrival time τ^1 of the first among N particles, i.e. $\tau^1 = \min(t_1, t_2, \dots, t_N)$ where t_n indicates the arrival time of the n -th particle. The distribution of τ^1 can be computed from the arrival distribution of an arbitrary particle:

$$P[\tau^1 > t] = P^N[t_1 > t] = \int_{\Omega} p(\mathbf{x}, t) d\mathbf{x}, \quad (2.1)$$

where the density p is solution of

$$\frac{\partial}{\partial t} p = \frac{\lambda_{\text{eff}} a^2}{2} \left(\frac{3}{2} \frac{\partial^2}{\partial x^2} p + \frac{3}{2} \frac{\partial^2}{\partial y^2} p + \sqrt{3} \frac{\partial^2}{\partial x \partial y} p \right) = \mathcal{L}(p), \quad (2.2)$$

and the effective per-tubule transition rate in the limit of fast switching is $\lambda_{\text{eff}} \sim \frac{1}{3 \tau_{\text{wait}}}$ (see the chapter 3 for the details). The mean time for the fastest among N is

$$\bar{\tau}^N = \int_0^{\infty} P[\tau^1 > t] dt = \int_0^{\infty} (P[t_1 > t])^N dt, \quad (2.3)$$

leading to the asymptotic expression (see chapter 3)

$$\bar{\tau}^N \approx c_N \frac{\delta_{\min}^2}{D \log N}, \quad (2.4)$$

where δ_{\min} is the graph distance between the source S and the target node T , $D = \lambda_{\text{eff}} \frac{a^2}{3}$ is the effective diffusion coefficient [Holcman and Schuss 2015] (a is the tubule length) and c_N is a constant that only depends on the network topology. We used equation 2.4 to fit the results of numerical simulations, obtaining good agreement (figure 2.3 B). This approximation, which is only valid when $\tau_{\text{switch}} \ll \tau_{\text{wait}}$, captures well the statistics of the fastest in the AFN. The arrival time for the fastest particle is quite different from the MFPT, which is given by

$$\langle \tau \rangle \approx \frac{3\sqrt{15} N_{\text{nodes}}}{2\pi \lambda_{\text{eff}}} \ln N_{\text{nodes}} + \mathcal{O}(1), \quad (2.5)$$

where N_{nodes} is the number of nodes in the network. At this stage we conclude that when the number of particles is sufficiently high ($N \geq 1000$ in a network with about 2000 nodes), the mean arrival time for the fastest particle can be two orders of magnitude faster (≈ 20 s, see figure 2.3 A) than the MFPT (≈ 25 min). We note moreover that this timescale decreases only moderately even when increasing the number of particles by a factor 10 (figure 2.3 A). The timescale of redistribution of the fastest particle is associated to extreme statistics [Bray, Majumdar, and Schehr 2013; Majumdar, Sabhapandit, and Schehr 2016; Schehr and Majumdar 2014] leading to the extreme behaviour [Basnayake, Schuss, and Holcman 2018; Basnayake et al. 2018; Coombs 2019; Schuss, Basnayake, and Holcman 2019; Sokolov 2019].

2.1.4 Active-flow network can generate synchronized packets

Another interesting effect of the AFN is the correlation induced by the edge flow on independent particles when the timescale of switching is large enough compared to the time spent in the nodes (i.e. $\tau_{\text{switch}} \gg \tau_{\text{wait}}$). In fact, due to the occurrence of capture states (figure 2.1 B), particles can accumulate in nodes. This accumulation creates a synchronization of the particle motion. As soon as one of the edges reverses its direction, allowing exit of the particles from the node, the particles start flowing together towards the only accessible neighbouring node. This behaviour can be repeated in consecutive capture states, resulting in a large fraction of the particles travelling together in synchronized packets (figure 2.3 C, with $\tau_{\text{switch}} = 3$ s). When $\tau_{\text{switch}} \ll \tau_{\text{wait}}$, the synchronization caused by capture states disappears, and we recover an almost classical diffusion regime, where the motion of particles is uncorrelated. Interestingly, for the AFN, the distribution of material at steady-state does not converge to a uniform density, where the material is evenly partitioned across the nodes, but aggregates keep appearing and disappearing (figure 2.4) so that at each time instant the particles are concentrated in just a few nodes (we will further examine this mechanism in chapter 3).

2.1.5 Redistribution timescale of the active-flow network

In this section we investigate the redistribution timescale and the MFPT of the AFN and compare it to an equivalent undirected network. In particular, we study two competing mechanisms governing the AFN dynamics. First, in the AFN particles can be slowed down by being trapped in capture states, thus we expect the exploration of the network to be slower with respect to the undirected network, where such situation does not occur. Second, the presence of directionality in the flow tend to drive particles to new regions, increasing the exploration speed, because the particles cannot linger between two nodes by going back and forth along the same edge. We formalize this effect by evaluating the probability that a particle moves to a new node and successively jumps back to the node it came from (we call this situation backtracking). We will study these mechanisms with more detail in chapter 3, but we provide here a summary of the main points.

To compute the consequences of capture states on the time it takes to travel along a trajectory, we estimated the steady-state probability p_{trap} for a particle to be trapped in a capture state. For the AFN, the edges switch at a Poissonian rate $\lambda_{\text{switch}} \equiv \tau_{\text{switch}}^{-1}$, and the waiting time in a node follows an exponential distribution $p_{\text{wait}}(t) = \frac{1}{\tau_{\text{wait}}} e^{-t/\tau_{\text{wait}}}$. For a general node degree d (i.e. the number of incident edges), the steady-state probability that a particle is in a trapping state can be computed by considering a Markov chain approximation of the random walk on the AFN (with a special construction detailed in chapter 3). We can condition the probability on the two possible previous states: either the particle was already trapped, or it jumped in the capture state from a neighbouring node. Note that it is necessary to consider previous trajectory steps because they define the directionality of the

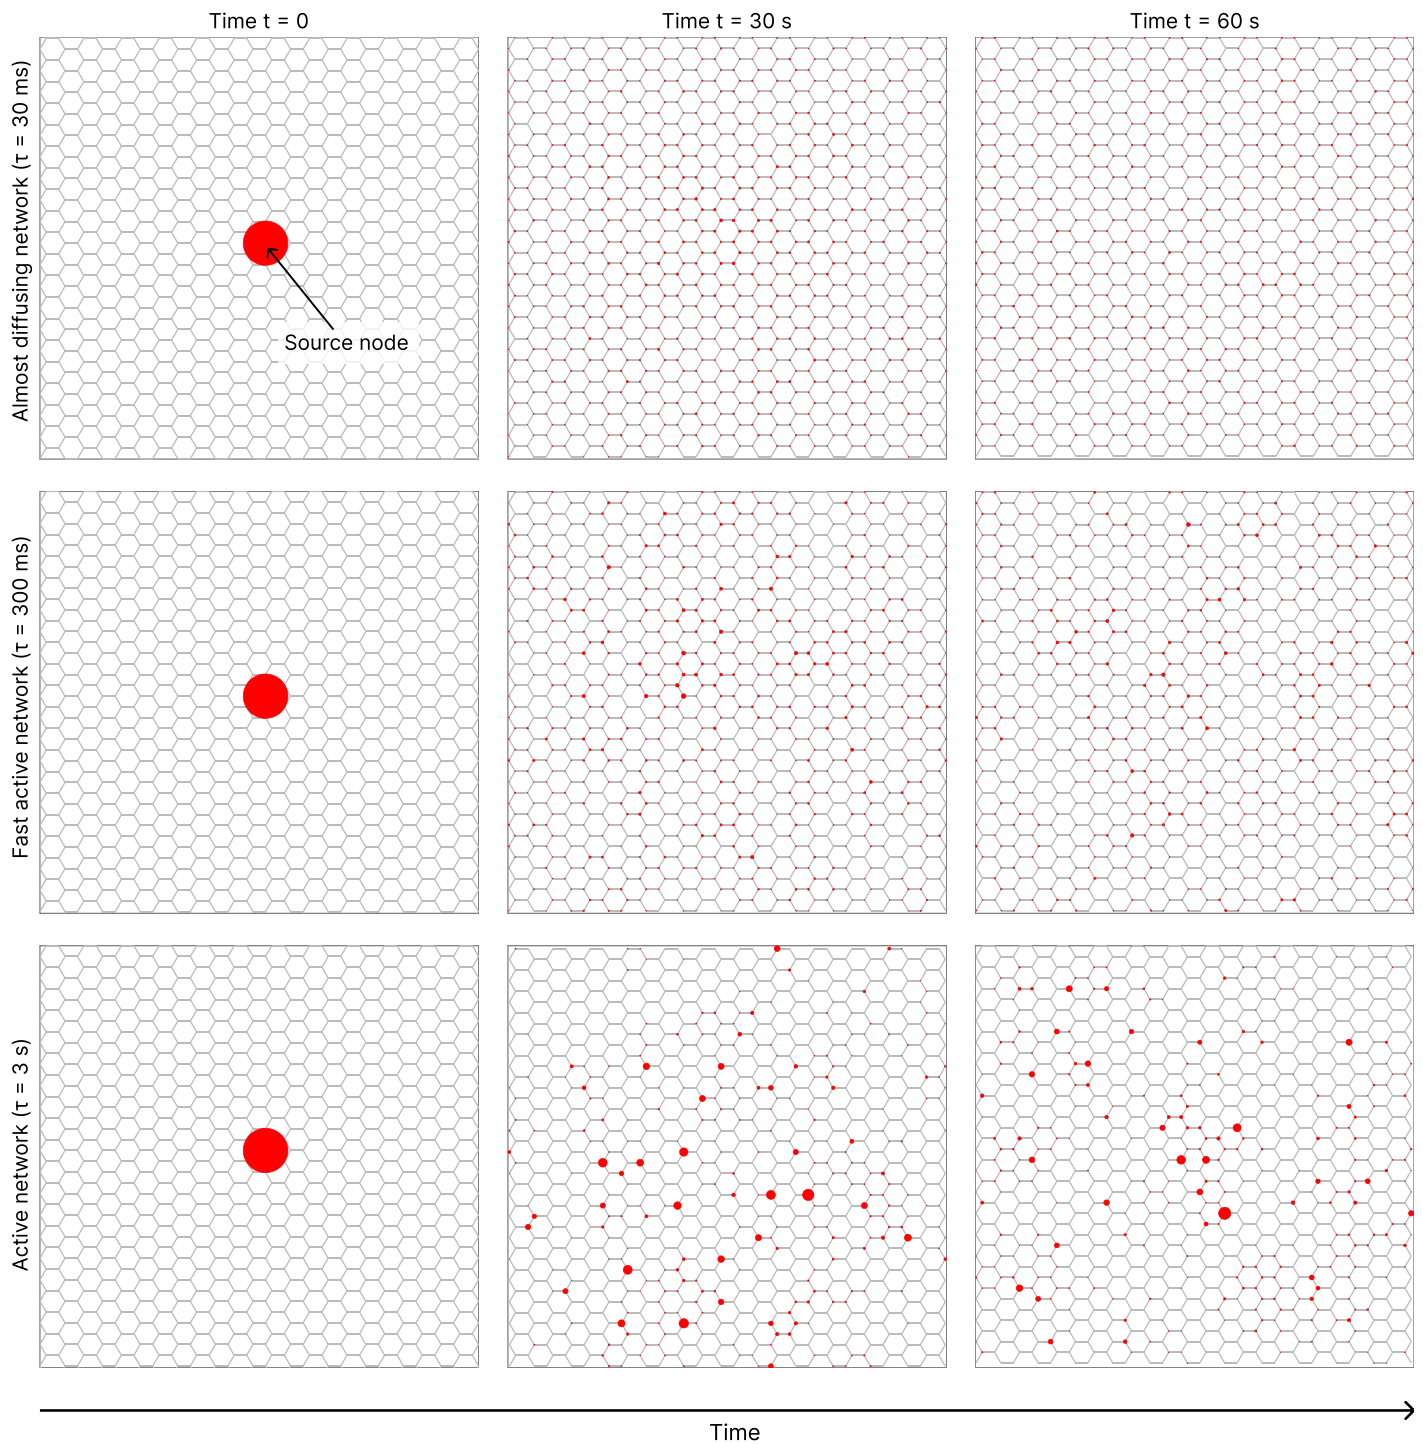
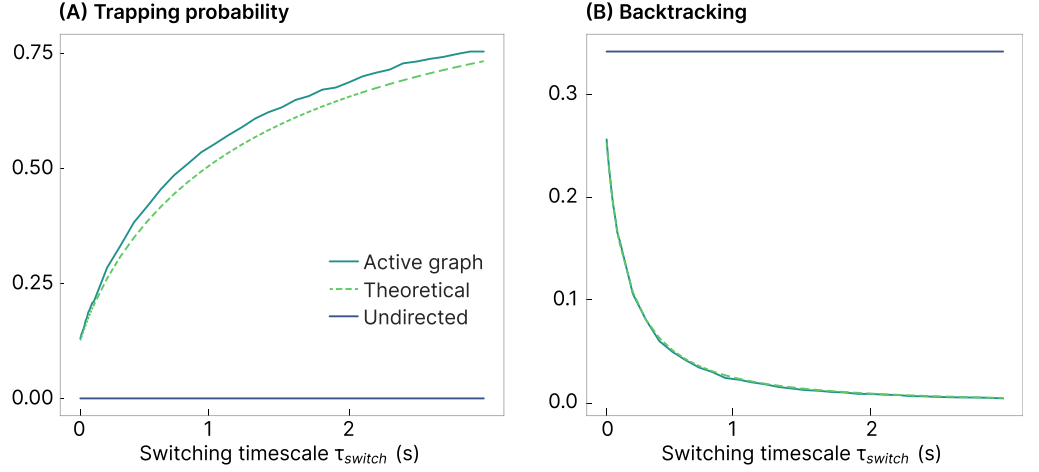


Figure 2.4 Example of transportation regimes of the AFN. Snapshots of the distribution of particles on the network following their release from the source node (in the middle). Each node is indicated by a red circle proportional to the number of particles contained. Columns indicate snapshots at time 0, 30 s, and 60 s. First row: almost diffusing active network ($\tau_{\text{switch}} = 30$ ms), the particles distribute uniformly on the network. Second row: fast active network ($\tau_{\text{switch}} = 300$ ms), some deviation from the uniform distribution of particles starts to be seen. Last row: slow active network ($\tau_{\text{switch}} = 3$ s), aggregates of particles in nodes (packets) are clearly visible.

Figure 2.5 Effect of the switching time τ_{switch} on the trapping and the backtracking probability. In the undirected network model (blue), active flow network. The simulation (solid green) is compared to the analytical approximation (dashed green). **A** Probability of trapping p_{trap} as a function of τ_{switch} . **B** Probability of backtracking p_{back} as a function of τ_{switch} .



edges that the particle crossed, and thus provide information regarding the state of the topology (this is remarkably different from a standard undirected network where the topology is not stochastic). This leads at steady-state to the Markov equation (see chapter 3)

$$p_{\text{trap}} = (1 - p_{\text{trap}}) \int_0^{\infty} \left(\frac{1 + e^{-2\lambda_{\text{switch}} t}}{2} \right) \left(\frac{1}{2} \right)^{d-1} p_{\text{wait}}(t) dt + p_{\text{trap}} \int_0^{\infty} \left(\frac{1 + e^{-2\lambda_{\text{switch}} t}}{2} \right)^d p_{\text{wait}}(t) dt, \quad (2.6)$$

where the first term denotes the probability for the particle of not being trapped in its previous step and becoming it in the current, and the second term represents consecutive trapping. A direct integration for $d = 3$ gives

$$p_{\text{trap}} = \frac{24q^3 + 34q^2 + 11q + 1}{192q^3 + 130q^2 + 23q + 1}, \quad (2.7)$$

where $q = \tau_{\text{wait}}/\tau_{\text{switch}}$. In figure 2.5, we compare the approximated solution of equation 2.7 (green dashed line) with the result of numerical simulations (green solid line). Interestingly, in the limit of fast switching $\tau_{\text{switch}} \ll \tau_{\text{wait}}$, the trapping probability converges to $(1/2)^d$, as there is no correlation between the state of the edges. It however never reach zero, as in the undirected case, since the edge directionality is still present. For $\tau_{\text{switch}} \gg \tau_{\text{wait}}$, the Markovian approximation tends to worsen because the correlation between the previous states becomes stronger.

In the limit of fast switching, motion on the AFN can be considered similar to a random walk on an undirected network, but with a total waiting time in nodes which includes the delay due to capture states. This can be computed by summing over all possibilities to be trapped n times:

$$\tau_{\text{tot}} \sim \sum_{n=1}^{\infty} (p_{\text{trap}})^{n-1} (1 - p_{\text{trap}}) \Gamma(n, \lambda_{\text{wait}}), \quad (2.8)$$

where $\Gamma(n, \lambda_{\text{wait}})$ is the gamma distribution for the sum of n independent exponential waiting times with rate $\lambda_{\text{wait}} \equiv 1/\tau_{\text{wait}}$. Then, the mean time is given by

$$\langle \tau_{\text{tot}} \rangle = \sum_{n=1}^{\infty} (p_{\text{trap}})^{n-1} (1 - p_{\text{trap}}) n \tau_{\text{wait}} = \frac{1}{1 - p_{\text{trap}}} \tau_{\text{wait}}. \quad (2.9)$$

To conclude, due to capture states, an AFN can retain particles in nodes much longer than the waiting timescale τ_{wait} . Moreover, at a coarse-grained level (i.e. with an observation timescale much larger than τ_{wait} and τ_{switch}), transport in AFN can be described as an effective diffusion-like model where $\langle \tau_{\text{tot}} \rangle$ is the effective waiting time in a node.

We now quantify the phenomenon of backtracking, which consists in a particle jumping back to the node it came from, thus wasting time by visiting again the previous node. For undirected networks, the probability for a particle to jump back to the node it came from depends only on the degree of connectivity: if the current node has d incident edges, the probability of going back is simply $p_{\text{back}}^{\text{undirected}} = \frac{1}{d}$. However, in the AFN this probability is affected by the direction of the edge, which flips at a rate $\lambda_{\text{switch}} \equiv \tau_{\text{switch}}^{-1}$. To compute the probability of a backtracking event in the AFN, we study the trajectory of a particle starting in a node N and immediately coming back to it after visiting an adjacent node A . To be able to go back node N it is necessary that, after the waiting time t spent in A , the flow from N to A has reverted its direction to be now oriented from A to N . Moreover, the particle has to choose the edge that goes back to N among the other possible outward edges. To compute the probability of selecting this edge, we fix the \overrightarrow{NA} edge and consider the configurations with other k edges in an outward directed state: the probability of choosing the backward one is $\frac{1}{k+1}$. Since the state of the backward edge is fixed (outward flow), the total number of edge configurations is 2^{d-1} . The probability of a configuration with k other edges being outward directed is $(\frac{1}{2})^{d-1} \binom{d-1}{k}$. Thus, the probability of choosing the specific path from A to N is (see chapter 3 for the detailed derivation)

$$P \left[\text{choosing } \overrightarrow{NA} \mid \overrightarrow{NA} \text{ outward} \right] = \frac{1}{2^{d-1}} \sum_{k=0}^{d-1} \frac{1}{k+1} \binom{d-1}{k} = \frac{2^d - 1}{d 2^{d-1}}. \quad (2.10)$$

The conditional probability of going back to N through node A , when the previous transition of the particle was made from N to A is

$$P[(N, A) \mid (A, N)] = \int_0^{\infty} \frac{1 - e^{-2\lambda_{\text{switch}}t}}{2} p_{\text{wait}}(t) dt \times P \left[\text{choose edge } \overrightarrow{NA} \mid \overrightarrow{NA} \text{ outward} \right],$$

where we have integrated over the switching time t with respect to the waiting time distribution $p_{\text{wait}}(t)$ in a node and $(1 - e^{-2\lambda_{\text{switch}}t})/2$ which represents a switching event for the edge direction during time t . Finally, at steady state, the backtracking probability is given by the probability to jump back in the absence of trapping:

$$p_{\text{back}} = (1 - p_{\text{trap}}) \frac{2^d - 1}{d 2^{d-1}} \int_0^{\infty} \frac{1 - e^{-2\lambda_{\text{switch}}t}}{2} p_{\text{wait}}(t) dt. \quad (2.11)$$

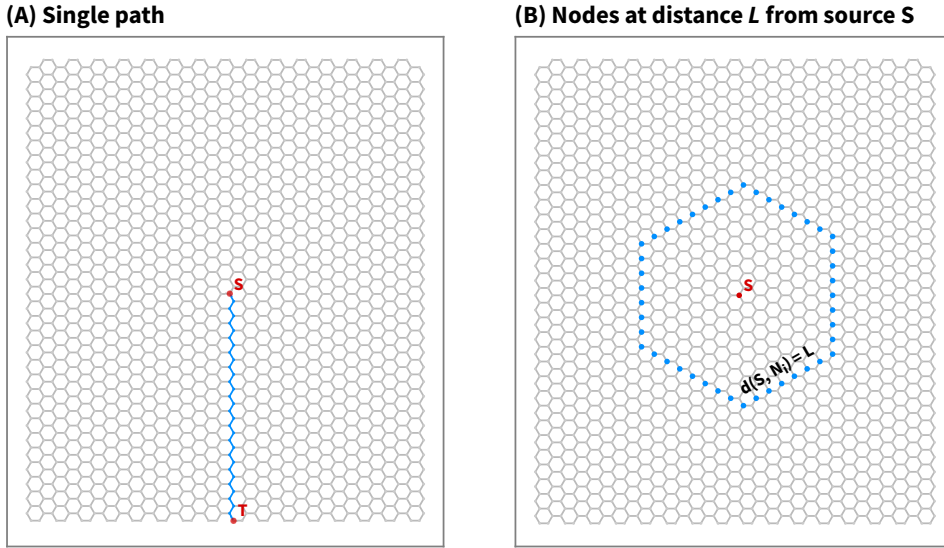


Figure 2.6 Structures used for the evaluation of redistribution kymographs. **A** The number of particle is evaluated along nodes on a single path, from S to T , at increasing distance. **B** The average number of particle is calculated over all nodes located at distance $d(S, T) = L$ from the source node.

For degree $d = 3$, direct integration leads to expression

$$p_{\text{back}} = (1 - p_{\text{trap}}) \frac{7}{12} \frac{q}{2q + 1}, \quad (2.12)$$

where the trapping probability p_{trap} is defined in equation 2.7 and $q = \tau_{\text{wait}}/\tau_{\text{switch}}$. We found a good agreement between this analytical result and numerical simulations (figure 2.5 B).

2.1.6 Transient material redistribution of the AFN

To study how particles are redistributed following local release at the source location S , we generated numerical simulations in the honeycomb network. Following the release, we monitor the number of particles in nodes along a single path (figure 2.6 A) and also over all nodes at a given distance $d(S, T) = L$ (figure 2.6 B). We analyse kymographs created by evaluating the average number of particles as a function of distance and time. Interestingly, the kymograph of the diffusion network is characterized by a propagation front $r^2 = ct$, where c is a constant that depends on the network topology and the diffusion coefficient (figure 2.7A). However, for the AFN, numerical simulations reveal that the propagation front can deviate significantly (figure 2.7 B–C) from what is observed for diffusion. In particular, as the switching time τ_{switch} increases, we observe that the material gets more fragmented, highlighting the transport mechanism by redundant, synchronized packets. Interestingly, packets can appear at a long distance quite rapidly due to the motion of the fastest packet [Basnayake and Holcman 2019; Schuss, Basnayake, and Holcman 2019]. By looking at kymographs representing the time evolution of the average density of particles over all nodes at a given distance from the source (figure 2.7, right column), we found that the AFN (figure 2.7 B–C) behaves similarly to the undirected network (figure 2.7 A) for values of τ_{switch} in the range 30–300 ms. Instead, looking

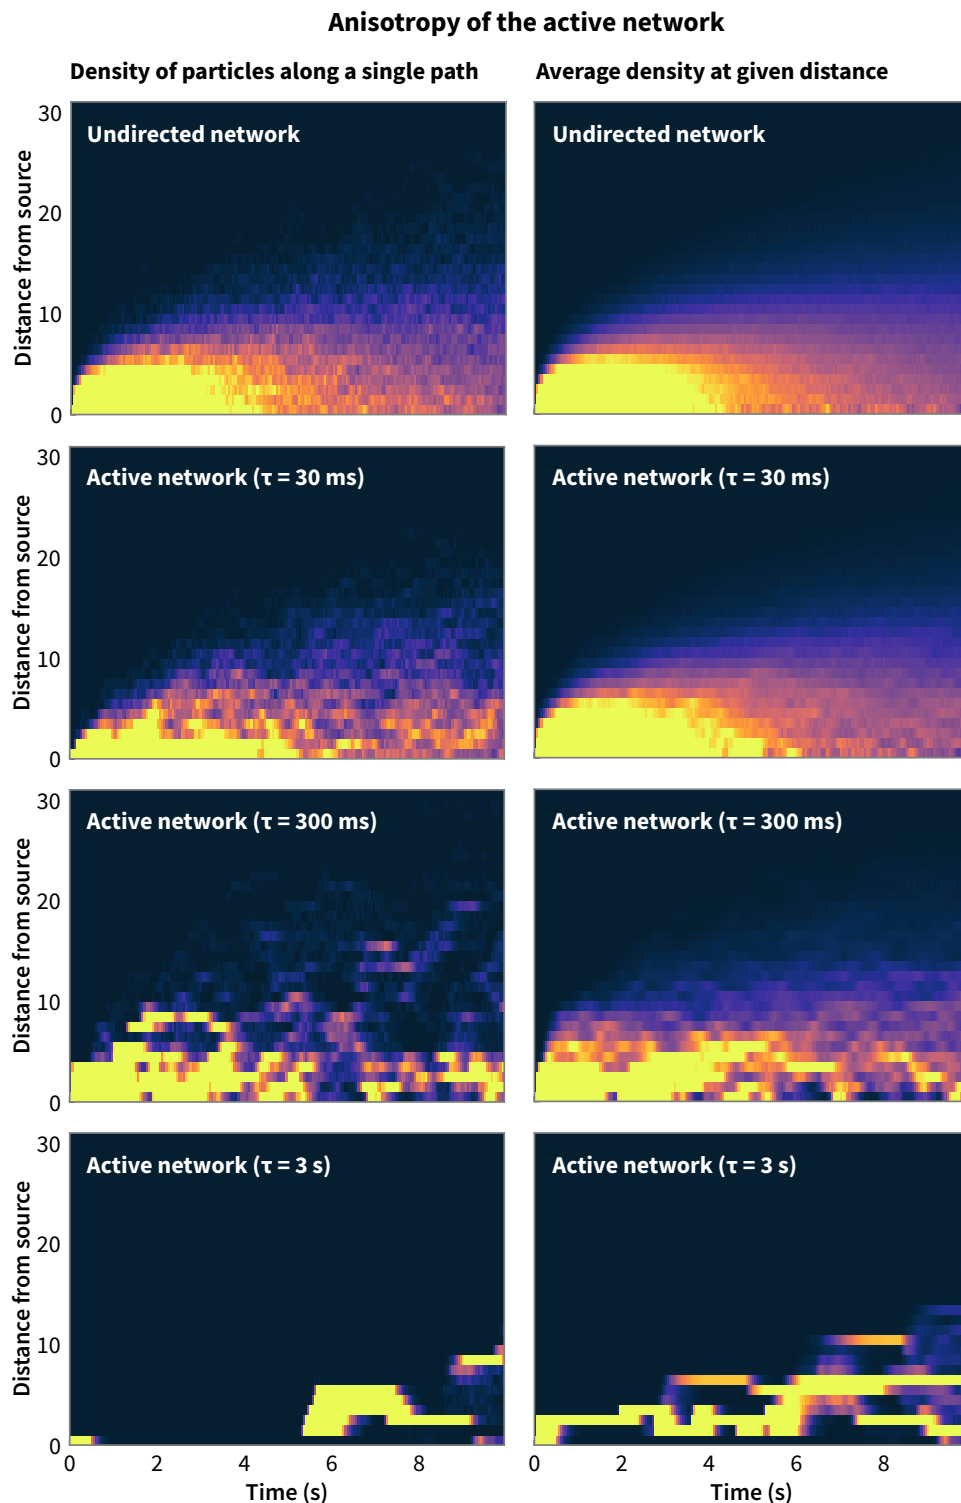


Figure 2.7 Kymographs of particle distribution following release at a single source node. The density of particles is computed along a single path (left column) or by averaging over all equidistant nodes (right column). Brighter colour means higher number of particles. Results were obtained simulating 10 000 particles on a honeycomb network for the various models. **A** Undirected network. **B** AFN with $\tau_{\text{switch}} = 30$ ms **C** AFN with $\tau_{\text{switch}} = 300$ ms **D** AFN with $\tau_{\text{switch}} = 3$ s

at the density along a single path, we observe the anisotropic behaviour of the AFN, as shown in figure 2.7 D ($\tau_{\text{switch}} = 300$ ms). For an undirected network (figure 2.7 A), representing a diffusive situation, there is almost no difference between the single path density and the average density, due to the isotropy of the dynamics. To conclude, we predict that following a release at one location, stable aggregates of material could appear quickly far away from the source, as observed along single paths of the network. These aggregates are the result of coordinated travel of many particles caused by capture states. This prediction could be tested using photoactivation experiments, where materials are transiently photoconverted and followed using live cell fluorescent microscopy [Holcman et al. 2018; Lippincott-Schwartz, Altan-Bonnet, and Patterson 2003].

2.2 Discussion and conclusions

We studied here the properties of the active flow network (AFN), a model of material transport inside the ER. We showed how the AFN can generate atypical protein redistribution, where molecules travel in packets that form and separate transiently while visiting all accessible nodes of the network. This property can allow proteins to be delivered in groups, which is a more reliable mode of transportation compared to individually uncorrelated motion occurring for pure diffusion models. Moreover, we showed that the arrival of the fastest particles or packets occurs along the shortest paths (figure 2.2 B). The present approach predicts that the timescale of redistribution (in the bulk of the network, for τ_{switch} in the range 30–300 ms and $\tau_{\text{wait}} = 100$ ms) is in the order of 20 min. However, if we consider only the first to arrive among more than 1000 simultaneously released particles, the timescale of arrival is in the order of 15–30 s (table 2.1). These average values are comparable to what is observed for random motion on an undirected network (i.e. diffusion-like).

2.2.1 Packet transportation guarantees a reliable delivery

The AFN allows the redistribution of many particles by redundant packets, a mechanism that guarantees a reliable delivery of groups of proteins to their final destination. Indeed, when enough molecules are transported at the same time, the biological function that has required this protein synthesis can be performed with a higher efficiency, compared to a transport mechanism where molecules arrive one at a time. Moreover, proteins undergo maturation while travelling on the endoplasmic reticulum, a process which requires a variable time. The redundancy resulting from the delivery of proteins in groups ensures the presence of at least some correctly folded and matured copies. Delivery in group can also play an important role at ER exit sites, where ensembles of particles are segregated and exported by COPII-coated vesicles [Budnik and Stephens 2009].

Particle synchronization is achieved by transiently blocking their motion in capture states. This mechanism should in principle impact the transport performance

because it effectively slows down the particle redistribution. Yet numerical simulations (figure 2.5) show that with the help of the unidirectional properties of the network nodes, particles explore the network faster due to a reduced backtracking probability. This faster exploration of the network can thus compensate the slowdown caused by capture states. For a wide range of parameter values, the active network can exhibit packet transportation and at the same time have an average redistribution timescale comparable to that of an undirected diffusive network. Even if packet transportation does not imply a performance hit, this active transport may require energy in the form of ATP consumption [Holcman et al. 2018; Nehls et al. 2000]. The exact mode of transportation in the ER remains unclear, although the range of switching timescale of tubule flow has been estimated to be within 30 ms and to 3 s [Holcman et al. 2018]. Such range should be further explored.

Moreover, the present analysis predicts that packets can persist at steady-state (more details on this are given in chapter 3). This implies that, even after a long time from the release of material in a source region, it should still be possible to see a non-uniform distribution in nodes. In particular, most of the nodes of the ER network should be almost empty, with only a small fraction of nodes containing most of the particles.

2.2.2 A theory to reconcile different types of experiments

Under the classical FRAP experiments, material redistribution in the ER was mostly viewed as driven by diffusion [Lippincott-Schwartz, Altan-Bonnet, and Patterson 2003; Nehls et al. 2000] and the dynamics is characterized by the overall diffusion coefficient D_{ER} . This model is relevant at a coarse spatiotemporal scale, where the dynamics in tubules can be neglected. At a finer scale, single-particle tracking experiments have suggested the presence of an active flow in tubules. Moreover, photoactivation experiments have shown a fast and unexplained propagation at large distances [Holcman et al. 2018]. These two statements seem in contradiction with the diffusion framework. However, based on the single-particle trajectories observations, the present model exhibits at small timescale the behaviour of an active flow, and at large timescale it can reproduce the properties of diffusion, thus effectively reconciling the FRAP diffusion-based approach with the presence of an active flow. Moreover, we have shown here that switching directionality generates a novel type of dynamics, characterized by motion by synchronous packets. This type of motion cannot be observed at the coarse spatiotemporal scale of FRAP experiments, nor it can be measured within a small set of single-particle trajectories. We suggest that high-resolution photoactivation experiments should be able to resolve the appearance of points where material is aggregated, as we predicted here by the AFN model.

2.2.3 Transport in other active networks

The network redistribution we explored here share some similarities with other networks. For example the one of the slime mould *Physarum polycephalum* [Alim et al. 2013], which grows as a random network of tubes, where peristalsis is used to drive internal cytoplasmic flows. However, the scale here is very different: for the slime mould, the scale is of few millimetres, and the fluid in tubules is modelled by continuous Navier-Stokes equations to link velocity and pressure. The fluid oscillates back and forth across the slime mould network with a timescale of 100 s. Monitoring the phase of contraction in each tubule revealed that contractions are cyclic in time. Interestingly, simulation results [Alim et al. 2013] revealed that the phase spatial gradient of this contraction is linear across the organism. Thus, a peristaltic wave can redistribute materials within the entire organism. However, in the ER, peristalsis could only be used to generate the fast motion in tubules. The physical mechanism causing the retention of particles in the ER nodes remains however to be understood. Another subcellular network is the one generated by mitochondria [Hoitzing, Johnston, and Jones 2015; Zamponi et al. 2018]. By fusing and dividing, mitochondria in eukaryotic cells can form tubular networks. How the material could be redistributed in such networks remains unclear, but it is possible that the difference in electrochemical membrane potential between each individual could contribute to create a local electric field, influencing the redistribution of ions and charged proteins.

To conclude, the present study suggests that, in the ER, the switching rate of tubule orientation could be a key modulator for material redistribution. The fact that the ER network covers most of the cell cytoplasm allows it to redistribute packets of proteins at a timescale of few seconds at any location. It would be interesting, in future works, to analyse this redistribution with respect to the distribution of the ER exit sites [Budnik and Stephens 2009], a key structure for the connection of the ER to other organelles.

Table 2.1 Parameters and results of the simulations on the reconstructed ER network. τ_{wait} is the exponential waiting time in nodes, τ_{switch} the time for edge to switch, τ_{tubule} the time needed to cross an edge. The results include the mean first-passage time (MFPT) and the mean extreme first-passage time (MEFPT) (i.e. the time for the first particles among N to arrive) computed for $N = 1000$. In both cases the target node is located at a distance of 25 edges from the source node.

Model type	τ_{wait} (ms)	τ_{switch} (ms)	τ_{tubule}	MFPT (min)	MEFPT (s)
Undirected	100	NA	0	19	30
AFN	100	30	0	20	33
AFN	100	300	0	19	15
AFN	100	3000	0	45	47

CHAPTER 3

Refined properties of the active-flow network

This chapter complements the previous one by providing the detailed derivations of the results that have been discussed in chapter 2, and more insights in the active-flow network (AFN) properties. First, I present the analytical derivations for the trapping and backtracking mechanisms, that we used to interpret the change in the mean first-passage time (MFPT) with respect to the switching rate and to define the coarse-grained effective diffusion equivalence. Second, I provide complementary details about particle distribution in packets at steady-state. In the last section, I report the details of the asymptotic computation of the MFPT and extreme statistics in the honeycomb network, approximated by using an embedding on a continuous two-dimensional space.

This chapter was adapted from “Active flow network generates molecular transport by packets: case of the endoplasmic reticulum”, by Matteo Dora and David Holcman, Proceedings of the Royal Society B (2020).

3.1 Trapping and backtracking

To explain how the edge switching timescale can influence the speed of transport, we consider two mechanisms that cause the trajectory of a particle to be slowed down. The first is the case of a particle that jumps back to the node it came from, thus wasting time by touching a node it had already visited. We call this kind of move *backtracking*. Note that this effect is present also in the pure random walk on an undirected network (diffusion-like model). The second effect is instead a peculiarity of the active network, where edge directionality can create capture states from which the particle cannot escape until one of the edges changes its direction. We recall that the dynamics of the active network model is governed by two Poisson rates: first, particles try to jump to neighbouring nodes at a rate $\mu = \tau_{\text{wait}}^{-1}$; second, the direction of each edge switches at a rate $\lambda = \tau_{\text{switch}}^{-1}$. A particle attempting to jump while in a capture state (i.e. all adjacent edges are inward directed) will be prevented from moving, and will have to wait for the next jump event before continuing

the network exploration. Considering an arbitrary trajectory, we want to calculate the steady-state probability that next move will result in backtracking (p_{back}) or in trapping (p_{trap}).

In the undirected network model, the probability for a particle to backtrack depends only on the node connectivity. If the current node has d incident edges, the probability of going back will be $p_{\text{back}}^{\text{undirected}} = \frac{1}{d}$. As edges have no directionality, no capture states can exist, and we define the probability of trapping $p_{\text{trap}}^{\text{undirected}} = 0$.

We consider now a process X moving on an active network where edges switch at a Poissonian rate $\lambda = \tau_{\text{switch}}^{-1}$ and the waiting time τ_{wait} follows the distribution $p_{\text{wait}}(t) = \mu e^{-\mu t}$. We also assume a regular network structure with constant node connectivity d . In principle, the probability of a particle to be trapped or to backtrack depends on its whole previous trajectory. In fact, the probability of hitting an edge in a given direction at a certain time $t + \Delta t$ will depend on its previously observed state at time t . However, this correlation is exponentially decreasing with Δt . For example, consider the probability that an edge is inward directed at $t + \Delta t$, given it was inward directed at time t . This corresponds to the probability of 0 or an even number of switching events in the interval Δt :

$$P[\text{edge inward, } t + \Delta t \mid \text{edge inward, } t] = \sum_{k=0}^{\infty} e^{-\lambda \Delta t} \frac{(\lambda \Delta t)^{2k}}{(2k)!} = \frac{1 + e^{-2\lambda \Delta t}}{2} \quad (3.1)$$

When $\lambda \Delta t \gg 1$, the probability goes to $\frac{1}{2}$. This means that after a sufficiently long time, we can ignore the information given from the trajectory history and just consider the probability of finding an edge in either state to be uniformly equal to $\frac{1}{2}$. In order to calculate the probability of trapping and backtracking, we approximate X with a process \tilde{X} where we neglect all the history but the two last steps, thus we write $P[\tilde{X}_i = x \mid \tilde{X}_{i-1}, \tilde{X}_{i-2}]$. While \tilde{X} is not itself a Markov process, we can build an equivalent Markov process T by grouping its states in tuples:

$$T_i = (\tilde{X}_i, \tilde{X}_{i-1}) \quad (3.2)$$

The T process can be understood to represent a transition between two nodes, and it clearly fulfils the Markovian property as the probability of a transition only depend on the previous one.

3.1.1 Trapping probability

To compute the trapping probability p_{trap} , we consider a node $A = X_i$ with neighbours N_1, N_2, \dots, N_d , the probability of trapping is:

$$P[T_i = (A, A)] = \sum_{k=1}^d P[T_i = (A, A) \mid T_{i-1} = (N_k, A)] P[T_{i-1} = (N_k, A)] + P[T_i = (A, A) \mid T_{i-1} = (A, A)] P[T_{i-1} = (A, A)] \quad (3.3)$$

where we conditioned the probability on all possible previous states T_{i-1} . The probability of remaining in a trap given the previous transition was already trapped is

equivalent to the probability that no edge has changed direction between the transition times t_{i-1} and t_i . Using the result from equation 3.1 and integrating over possible waiting times ($\Delta t = t_i - t_{i-1}$):

$$P[T_i = (A, A) | T_{i-1} = (A, A)] = \int_0^\infty d(\Delta t) \left(\frac{1 + e^{-2\lambda\Delta t}}{2} \right)^d p_{wait}(\Delta t) \quad (3.4)$$

Instead, when the previous transition started from a neighbour N , we only know that the corresponding edge was inward directed, and we assign probability of $\frac{1}{2}$ to the remaining ones:

$$P[T_i = (A, A) | T_{i-1} = (N, A)] = \int_0^\infty d(\Delta t) \left(\frac{1 + e^{-2\lambda\Delta t}}{2} \right) \left(\frac{1}{2} \right)^{d-1} p_{wait}(\Delta t) \quad (3.5)$$

We use the relation

$$\sum_{k=1}^d P[T_{i-1} = (N_k, A)] = 1 - P[T_{i-1} = (A, A)], \quad (3.6)$$

and consider that at steady-state we must have $P[T_i = (A, A)] = P[T_{i-1} = (A, A)] \doteq p_{trap}$. Then, equation 3.3 becomes:

$$\begin{aligned} p_{trap} = & (1 - p_{trap}) \int_0^\infty d(\Delta t) \left(\frac{1 + e^{-2\lambda\Delta t}}{2} \right) \left(\frac{1}{2} \right)^{d-1} p_{wait}(\Delta t) \\ & + p_{trap} \int_0^\infty d(\Delta t) \left(\frac{1 + e^{-2\lambda\Delta t}}{2} \right)^d p_{wait}(\Delta t) \end{aligned} \quad (3.7)$$

Solving for the case of the hexagonal lattice ($d = 3$) leads to:

$$p_{trap} = \frac{24q^3 + 34q^2 + 11q + 1}{192q^3 + 130q^2 + 23q + 1}. \quad (3.8)$$

3.1.2 Backtracking probability

We take a similar approach in order to calculate the backtracking probability. The probability of backtracking to node A at steady state is:

$$\begin{aligned} p_{back} &= \sum_{k=1}^d P[T_i = (N_k, A), T_{i-1} = (A, N_k)] \\ &= \sum_{k=1}^d P[T_i = (N_k, A) | T_{i-1} = (A, N_k)] P[T_{i-1} = (A, N_k)] \end{aligned} \quad (3.9)$$

To go back to node A , it is necessary that after the waiting time Δt the back-linking edge has changed its direction to be outward directed. Moreover, it is necessary that, even if other edges are available, the particle randomly chooses the one going back, thus

$$\begin{aligned} P[T_i = (N, A) | T_{i-1} = (A, N)] &= \int_0^\infty d(\Delta t) \left(1 - \frac{1 + e^{-2\lambda\Delta t}}{2} \right) p_{wait}(\Delta t) \\ &\times P \left[\text{choosing edge } \vec{NA} \mid \vec{NA} \text{ outward} \right]. \end{aligned} \quad (3.10)$$

To compute the probability of choosing the back-linking edge given that it is directed towards A, we consider all the possible configurations of other edges direction. If we assume that k edges are open, the probability of choosing the backward one is $\frac{1}{1+k}$. Since the backward edge is fixed, we have a total of 2^{d-1} configurations and the probability of a configuration with k edges open is $(\frac{1}{2})^{d-1} \binom{d-1}{k-1}$. The probability of choosing the backward edge is thus:

$$P[\text{choosing edge } \overrightarrow{NA} \mid \overrightarrow{NA} \text{ outward}] = \frac{1}{2^{d-1}} \sum_{k=1}^{d-1} \frac{1}{k} \binom{d-1}{k-1} = \frac{2^d - 1}{d 2^d}. \quad (3.11)$$

Substituting the results of equations 3.6, 3.10 and 3.11 into equation 3.9 we get:

$$p_{back} = (1 - p_{trap}) \frac{2^d - 1}{d 2^d} \int_0^\infty d(\Delta t) \left(1 - \frac{1 + e^{-2\lambda\Delta t}}{2}\right) p_{wait}(\Delta t) \quad (3.12)$$

In particular, for $d = 3$, we obtain the explicit expression

$$p_{back} = (1 - p_{back}) \frac{7}{12} \frac{\lambda}{2\lambda + \mu}. \quad (3.13)$$

3.1.3 Effective diffusion model for the AFN

We analyse here the behaviour of the active-network model in the limit of fast switching. We shall prove that when the switching rate $\lambda = 1/\tau_{switch}$ is very large, the active-network model can be considered equivalent to a random walk on the graph. As $\lambda \rightarrow \infty$, the correlation between the edge direction at different times vanishes. In fact, the probability of the edge to be in a certain state at time $t + \Delta t$ does not depend any more on its state at time t . As this probability becomes time independent, it is immediate to conclude that the inward and outward directed states must be equally probable.

$$P[\text{edge inward}] = P[\text{edge outward}] = \frac{1}{2}, \quad \lambda \rightarrow \infty \quad (3.14)$$

Note that the probability of a capture state does not vanish. In the limit $\lambda \rightarrow \infty$ there is no correlation, and we can compute this probability exactly:

$$p_{back} = (P[\text{edge inward}])^3 = \frac{1}{8} \quad (3.15)$$

where we have considered the case of a regular lattice of degree 3.

However, given p_{trap} (equation 3.8), we can calculate the total time that a particle waits in a node, including the trapping. The probability for a particle to stay trapped in a node n consecutive times is:

$$P[\text{trapped } n \text{ times}] = (p_{back})^n (1 - p_{back}). \quad (3.16)$$

Thus, the total time spent in the node is a random variable distributed as:

$$\tau_{tot} \sim \sum_{n=1}^{\infty} (p_{back})^{n-1} (1 - p_{back}) \Gamma(n, \mu), \quad (3.17)$$

where $\Gamma(n, \mu)$ is the distribution of the sum of n independent exponential waiting times with rate $\mu = 1/\tau_{\text{wait}}$. Taking the mean:

$$\langle \tau_{\text{tot}} \rangle = \sum_{n=1}^{\infty} (p_{\text{back}})^{n-1} (1 - p_{\text{back}}) n \tau_{\text{wait}} = \frac{1}{1 - p_{\text{back}}} \tau_{\text{wait}}. \quad (3.18)$$

As in the pure random walk, the process has no preferential directionality since particles will explore the network in all directions with equal probability. We can thus describe the active-network model as a random walk where the waiting time in nodes is given by the distribution of equation 3.17. This result will also hold for large but finite λ , if we consider a coarse-grained observation timescale ξ such that $\xi \gg \tau_{\text{wait}}$ and $\xi \gg \tau_{\text{switch}}$. Then, the model is equivalent to a random walk on an undirected network with effective waiting time:

$$\tau_{\text{wait}}^{\text{RW}} = \frac{1}{1 - p_{\text{back}}} \tau_{\text{wait}}. \quad (3.19)$$

3.2 Properties of the active network transport by packets

3.2.1 Statistics of the fastest particles

In this section, we study the properties of the fastest particles moving inside the AFN (figure 3.1 A–B). For an undirected network, the distribution of lengths of the fastest among N particles concentrates toward the shortest path as the number of particles increases (figure 3.1 A). However, for the AFN, the length of the shortest trajectories presents a heavier tail in the distribution, mainly independent of N , showing that the fastest trajectories are not necessarily along the shortest path, but take longer routes to avoid possible trapping nodes (figure 3.1 B). This effect is confirmed by plotting the distribution of lengths and arrival times for various switching times τ_{switch} (figure 3.1C–D).

By studying the average time spent in a node (figure 3.1 E), plotted by order of arrival, we find that the average time per node is almost constant for both undirected and directed networks, and that the lengths increase similarly (figure 3.1 F). However, trajectory lengths of the fastest particle are longer in the active network (figure 3.1 F), confirming the deviation of the trajectory of the fastest from the shortest path. To conclude, extreme statistics of particle redistribution in an AFN show some deviation with respect to undirected networks, as they are characterized by longer trajectories due to the avoidance of capture states.

3.2.2 Transient to steady-state regime

To explore the differences in the redistribution of materials between the undirected and active network, we study here the transient regime to steady-state of $N = 10\,000$ particles simultaneously released from a single source node, following the time evolution of the distribution of the number of particles per node. Initially, all particles

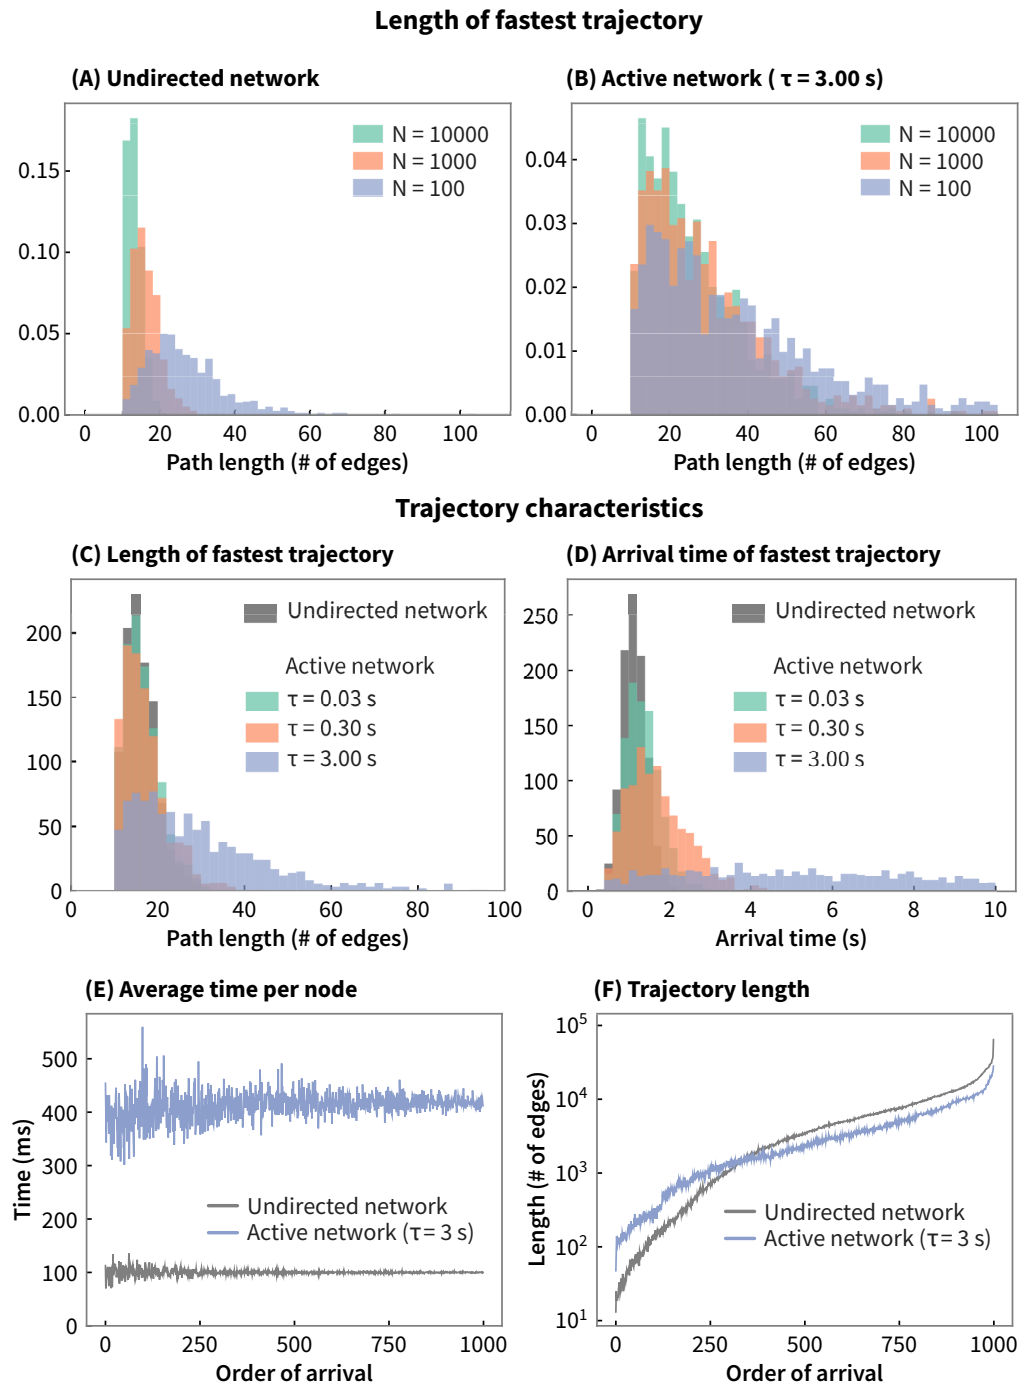


Figure 3.1 Characteristics of the fastest trajectories. Trajectories are considered for a target at distance 10 from the source node. **A–B** Distribution of the length of the first particles to arrive among N for undirected network and AFN. **C–D** Distribution of the length for the undirected and active network for various switching timescales τ . **E–F** Average time spent in nodes and trajectory length of the first 1000 particles among $N = 10\,000$, plotted by order of arrival.

are concentrated in the source node and thus the distribution is peaked at 0 (i.e. all nodes except the source are empty, figure 3.2 A–B). Intuitively, in the case of an undirected network, at steady-state the material tends to be uniformly distributed among the nodes (figure 3.2 A3). The distribution of the number of particle per node at a given time instant will thus be peaked around the value N/M , where N is the number of released particles and M is the total number of nodes in the network (figure 3.2 A1–A2). In the case of the active network, for sufficiently large switching time, this steady-state equilibration does not take place (figure 3.2 B1–2), and the distribution of particles per node remains peaked at zero, meaning that most nodes are empty. In fact, due to capture states, particles tend to aggregate in just a few nodes which behave as attractors (figure 3.2 B3). Thus, while in the undirected network particles appear to be approximately uniformly distributed, the active network is characterized by sparse high-density clusters. In the following sections we will better quantify this behaviour for the two cases.

3.2.3 Steady-state particle distribution in the undirected network

In this section, we study how particles are distributed at steady-state in an undirected network. A particle released from the source node, in the long time, will visit all nodes on the network. Thus, for sufficiently large time, the probability for a particle to be in a given node will be independent of the node position. We can then describe the probability of finding a particle in a node in a given node m by a Bernoulli random variable depending only on the size of the network:

$$\mathbb{P}[\text{particle is in node } m] \simeq \frac{1}{M}, \quad \text{for } t \text{ large}, \quad (3.20)$$

where M indicates the number of nodes in the network. If we now consider N particles, the probability of finding k of them in node m is given by the binomial distribution:

$$\mathbb{P}[k \text{ particles in node } m] = \binom{N}{k} p^k (1-p)^{N-k}, \quad (3.21)$$

where $p = 1/M$. We compare this analytical result with numerical simulations in figure 3.3 A, confirming its accuracy. Interestingly, simulations in the active network (figure 3.3 B) show a completely different result, with the distribution peaked at zero, a situation that we will discuss in the next section.

3.2.4 Steady-state particle distribution in the AFN

We now consider the active network in the limit of very low switching rate. In this case, edges keep their directionality for a very long time and particles will freely jump through the network until they fall in a trap. Initially, we neglect for simplicity the topology of the network and the size of attractors created by the trap nodes. Following this assumption, we expect all non-trap nodes to be empty and the particles to be randomly distributed in trap nodes at steady-state. We describe then the




Figure 3.2 Time evolution of the distribution of the number of particles per node. **A** Undirected network. **B** Active network with $\tau_{\text{switch}} = 3$ s. The two models are based on the same hexagonal lattice with $M = 2046$ nodes. $N = 10000$ particles are released from a source node at the middle of the network. First and second rows: Time evolution of the distribution of particles per node. Last row: Spatial distribution of particles on the network at steady state.

probability of finding k particles in a node based on whether it is a trap or not. If the node is not a trap, the answer is simply:

$$\mathbb{P}[k \text{ particles in node } x \mid x \text{ is not a trap}] \simeq 0. \quad (3.22)$$

If instead the node is a trap, we proceed as in the undirected network, but now we consider the possible distribution of particles only in the trap nodes.

$$\mathbb{P}[k \text{ particles in node } x \mid x \text{ is a trap, } M_{\text{trap}}] = \binom{N}{k} p^k (1-p)^{N-k}, \quad p = \frac{1}{M_{\text{trap}}}, \quad (3.23)$$

where M_{trap} is the total number of trap nodes. Note that the number of trap nodes is a random variable which depends on the network topology. For a node to be a trap all its edges have to be inward directed (each will be with probability $1/2$). On average, we can consider

$$\mathbb{P}[\text{node is a trap}] = \left(\frac{1}{2}\right)^d \equiv p_{\text{trap}}, \quad d = \text{average node degree} \quad (3.24)$$

The number of traps in a network will then be a binomially distributed with $n = M$ (total number of nodes) and $p = p_{\text{trap}}$.

$$\mathbb{P}[m \text{ traps in network}] = \binom{M}{m} (p_{\text{trap}})^m (1-p_{\text{trap}})^{M-m} \equiv \varphi_{\text{trap}}(m) \quad (3.25)$$

Note that once we fix the number of trap nodes in the network to m , the probability of selecting a trap node is m/M (recall that we neglect the topology). Finally, we

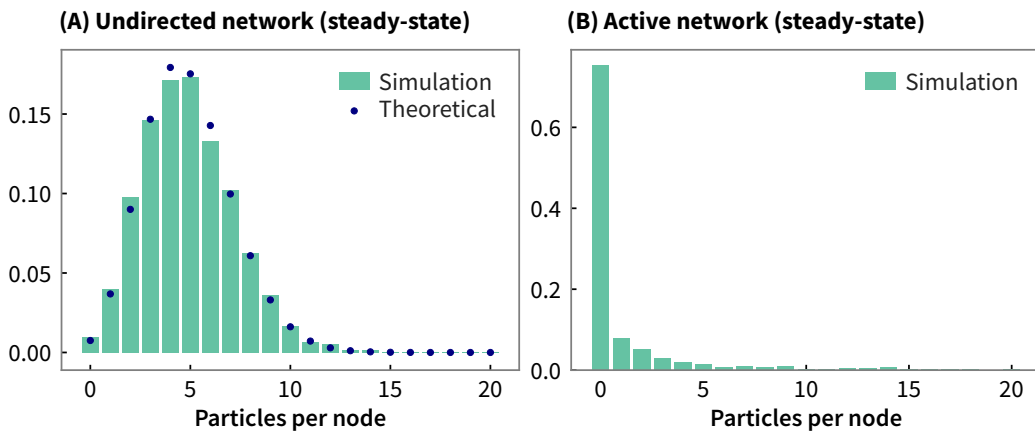


Figure 3.3 Distribution of particles per node at steady-state **A** Undirected network. The result of simulations is compared to the analytical result (equation 3.21). **B** Active network ($\tau_{\text{switch}} = 3$ s). The simulations use $N = 10000$ particles and a network of $M = 2046$ nodes.

remove the conditioning by summing over the possible number of traps, obtaining

$$\mathbb{P} \left[\begin{array}{c} k \text{ particles} \\ \text{in node } x \end{array} \right] = \sum_{m=0}^M \left[\left(1 - \frac{m}{M}\right) \delta_{k,0} + \frac{m}{M} \binom{N}{k} \left(\frac{1}{m}\right)^k \left(1 - \frac{1}{m}\right)^{N-k} \right] \varphi_{\text{trap}}(m) \quad (3.26)$$

While this derivation roughly predicts the peak in density of empty nodes we obtain from simulations (figure 3.3 B and figure 3.4 C), it does not match the tail of the distribution. In fact, our current derivation is based on the assumption that all trap nodes attract roughly the same number of particles. To find a better analytical form for the distribution of particles in nodes we need to take topology into account.

If we consider a fixed configuration of the network, with all edge directionalities fixed (i.e. a directed graph), trap nodes will act as attractors. We have thus paths directed to the attraction node from which particles cannot escape. We define the basin of attraction of a trap node to be composed of all nodes connected to it (taking directionality into account). In the framework of dynamical systems theory, trap nodes act as fixed points. For particular configurations of the network, we can also find loops surrounded by inward directed edges from which the particles cannot escape, corresponding to limit cycles in classical dynamical systems (figure 3.4 A, cycle A2). Since the occurrence of limit cycles is rare, we will neglect this kind of configuration in the following derivation.

We define the mass of an attractor as the number of nodes in its basin of attraction. Note that nodes on the boundary of the basin of attraction have edges linking to neighbouring attractors, meaning that particles starting in a boundary node will randomly fall in one of the possible attractors. We thus count only a fraction of their weight (proportional to the fraction of edges connected to other nodes belonging to the attractor) when calculating the mass of the attractor.

Similarly to the previous derivation based on trap nodes, we assume that particles starting in any node of the network will converge, at steady-state, to the corresponding attractor. We thus expect particles to be concentrated in attractors, distributed proportionally to their mass. To investigate this, we first need to derive the distribution of the mass of the attractors. From percolation theory [Bollobas, Bollobás, and Riordan 2006], we know that the mass distribution of similar clusters in undirected graphs can be described with an exponential law:

$$p_{\text{mass}}(m) \sim e^{-\beta m}, \quad \frac{1}{\beta} = m_{\xi}, \quad (3.27)$$

where m_{ξ} is the characteristic mass of the cluster. We derived the distribution of attractor mass from simulations on a hexagonal lattice with $M = 2046$ nodes and fitted it with an exponential law, finding $m_{\xi} \simeq 16$ (figure 3.4 B). The probability that a node x belongs to an attractor C_m of mass m is then:

$$\mathbb{P} [x \in C_m] \sim m \cdot p_{\text{mass}}(m) \quad (3.28)$$

up to a normalization factor. Now we consider the possible outcomes for the number of particles in a node, conditioned on the fact that it belongs to an attractor C_m . Since limit cycles are rare in the hexagonal lattice (the configuration requires fixing 12 edges, i.e. a probability in the order of 2^{-12}), we only consider the case of fixed points. In an attractor of mass m , at steady-state $m - 1$ nodes will be empty, and all particles will be concentrated in the centre of attraction. If we consider an initial condition of α particles per node, we have:

$$\mathbb{P}[x \text{ is empty} | x \in C_m] = \frac{m - 1}{m}, \quad (3.29)$$

$$\mathbb{P}[x \text{ contains } k \text{ particles} | x \in C_m] = \frac{1}{m} \delta_{k, \alpha \cdot m}, \quad (3.30)$$

Note that we approximated the number of particles in the centre of attraction to be exactly $\alpha \cdot m$. In reality, there will be fluctuations due to the presence of boundary nodes which connect different attractors, allowing particles to randomly distribute among them. For $\alpha \gg 1$, we assume these fluctuations to be negligible. Summing over possible attractor mass:

$$\mathbb{P}[k \text{ particles in node } x] = \begin{cases} A \cdot \sum_m p_{mass}(m) \cdot (m - 1), & k = 0; \\ A \cdot p_{mass}(k/\alpha), & k > 0; \end{cases} \quad (3.31)$$

where A is a normalization factor. We did not consider here limit cycles because they are rare in the hexagonal lattice (the configuration requires fixing 12 edges, i.e. a probability in the order of 2^{-12}).

In figure 3.4 C we compare the number of particles per nodes obtained from simulations, the analytical result of equation 3.31, and the approximation of equation 3.26 which was obtained by neglecting the topology. We found a good agreement when we accounted for the topology of the network. The effect of the fluctuations is evaluated figure 3.4 D for the smallest attractor that consists of 4 points (containing 5 particles per node). This result explains the deviation between equation 3.31 and the numerical simulations in figure 3.4 C.

To conclude, the distribution of particles per node is affected by the segregation of the AFN in attractors and by the fluctuation of the number of particles (located on edges between two attractors). This analysis explains that in long time limit, nodes of an AFN are mostly empty of particles and only a small fraction (attractor points) contains most of the particles.

3.3 Mean first-passage time and extreme statistics

3.3.1 Asymptotic computation of MFPT and flux

In this section, we compute the MFPT for a particle released from the source to reach a target site. We consider the case of a fast switching limit, where motion of the particle in the endoplasmic reticulum (ER) can be approximated as a random walk,

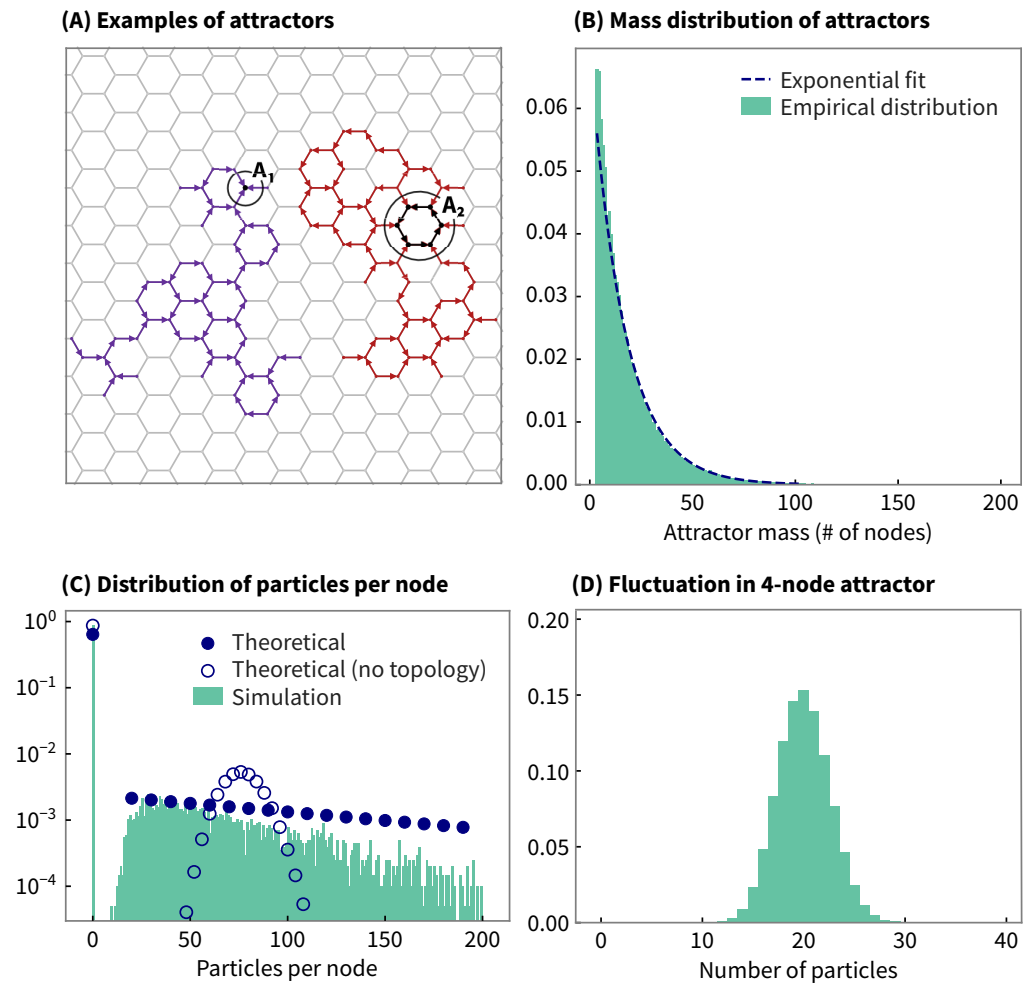


Figure 3.4 Steady-state particle distribution in the AFN. **A** Example of attractor configurations, showing an attractor point (A_1), and a limit cycle (A_2). **B** Distribution of attractor size in realization of an AFN, fitted by an exponential [Bollobas, Bollobás, and Riordan 2006]. **C** Distribution of particles per node obtained from simulations (green), analytical equation 3.26 (hollow circles), and equation 3.31 (filled circles). **D** Fluctuation of the number of particle per node for a 4-node attractor.

with an equal probability to jump to any neighbouring node. We define by λ the effective transition rate from a node to each neighbour. For the AFN, this effective rate is obtained from the effective diffusion model (see section 3.1.3) as $\lambda = 1/3 \tau_{\text{wait}}^{\text{RW}}$, where $\tau_{\text{wait}}^{\text{RW}}$ is the effective waiting time defined in equation 3.19. Coherently with the structure of the ER, we consider here a hexagonal lattice network where each node is connected to three neighbours by edges of length a .

To coarse-grain the random walk occurring on the two-dimensional hexagonal lattice, we use the Markov transition equation [Schuss 1980]. The transition probability function between time t and $t + \Delta t$ from any node \mathbf{x} to the three neighbouring ones is given by

$$\begin{aligned} P[\mathbf{x}, t + \Delta t] = & P[\mathbf{x}, t] (1 - 3\lambda\Delta t) + \lambda\Delta t P[\mathbf{x} + a, t] + \lambda\Delta t P[\mathbf{x} + ae^{2i\pi/3}, t] \\ & + \lambda\Delta t P[\mathbf{x} + ae^{4i\pi/3}, t] + o(\Delta t), \end{aligned}$$

where the first term in the l.h.s. represents the probability of no transition from \mathbf{x} to any nodes and the other term corresponds to a transition from each of the three nodes. In the small Δt limit, we obtain the master equation:

$$\begin{aligned} \frac{d}{dt}P[\mathbf{x}, t] = & -3\lambda P[\mathbf{x}, t] + \lambda P[\mathbf{x} + a, t] \\ & + \lambda P[\mathbf{x} + ae^{2i\pi/3}, t] + \lambda P[\mathbf{x} + ae^{4i\pi/3}, t]. \end{aligned} \quad (3.32)$$

In the limit of a large network compared to the length a , a Taylor expansion of equation 3.32 to second order with respect a leads to get the Master-Fokker-Planck equation for the transition probability $p(\mathbf{x}, t) = P[\mathbf{x}, t]$,

$$\begin{aligned} \frac{\partial}{\partial t}p(\mathbf{x}, t) = & \lambda a (1 + \cos(2\pi/3) + \cos(4\pi/3)) \frac{\partial}{\partial x}p(\mathbf{x}, t) \\ & + \lambda a (1 + \sin(2\pi/3) + \sin(4\pi/3)) \frac{\partial}{\partial y}p(\mathbf{x}, t) \\ & + \frac{\lambda a^2}{2} (1 + \cos^2(2\pi/3) + \cos^2(4\pi/3)) \frac{\partial^2}{\partial x^2}p(\mathbf{x}, t) \\ & + \frac{\lambda a^2}{2} (\sin^2(2\pi/3) + \sin^2(4\pi/3)) \frac{\partial^2}{\partial y^2}p(\mathbf{x}, t) \\ & + \lambda a^2 \sqrt{3} \frac{\partial^2}{\partial x \partial y}p(\mathbf{x}, t), \end{aligned}$$

which simplifies to the second order operator:

$$\frac{\partial}{\partial t}p = \mathcal{L}(p) \text{ where } \mathcal{L}(p) = \frac{\lambda a^2}{2} \left(\frac{3}{2} \frac{\partial^2}{\partial x^2}p + \frac{3}{2} \frac{\partial^2}{\partial y^2}p + \sqrt{3} \frac{\partial^2}{\partial x \partial y}p \right). \quad (3.33)$$

This is equivalent to the approximation of network by an ensemble of N hexagonal domains with size a and the MFPT reduces to the one of a diffusing particle, characterized by a diffusion tensor

$$\mathbf{D} = \frac{\sqrt{3}}{2D_a} \begin{pmatrix} \sqrt{3} & 1 \\ 1 & \sqrt{3} \end{pmatrix} \quad (3.34)$$

satisfying the stochastic equation

$$d\mathbf{x} = \sigma d\omega, \quad (3.35)$$

where $\mathbf{D} = \frac{1}{2}\sigma^T\sigma$, $D_a = \frac{\lambda a^2}{2}$ and ω is a two-dimensional Brownian motion. In that case, the computation of the MFPT reduces to the narrow escape to a small disk of radius a . The MFPT $\tau(\mathbf{x}) = u(\mathbf{x})$ is solution of the boundary value problem inside the network domain Ω :

$$\begin{cases} D_a \left(\frac{3}{2} \frac{\partial^2 u}{\partial x^2} + \sqrt{3} \frac{\partial^2 u}{\partial x \partial y} + \frac{3}{2} \frac{\partial^2 u}{\partial y^2} \right) = -1, & \text{for } \mathbf{x} \in \Omega \\ u = 0, & \text{on } \omega_a = \partial\Omega_a \\ \frac{\partial u}{\partial n} = 0, & \text{on } \omega_r = \partial\Omega_r, \end{cases} \quad (3.36)$$

where the absorbing condition is given on the small fraction ω_a of the boundary, while ω_r is the reflecting boundary of the network. For a two-dimensional planar domain, the solution of 3.36 can be obtained from the narrow escape theory (NET) [Holcman and Schuss 2015] and the solution is independent of the initial position \mathbf{x} . To determine the NET formula in dimension 2, we change coordinates to diagonalize the operator \mathcal{L} , and we obtain the two eigenvalues $\lambda_1 = 3 + \sqrt{2}$, $\lambda_2 = 3 - \sqrt{2}$. In the eigenvector coordinates (u, v) , the MFPT $U(\mathbf{x}) = u(\mathbf{x})$ satisfies

$$\begin{cases} D_a \left(\lambda_1 \frac{\partial^2 U}{\partial u^2} + \lambda_2 \frac{\partial^2 U}{\partial v^2} \right) = -1 \\ U = 0 & \text{on } \partial\Omega_a \\ \frac{\partial U}{\partial n} = 0 & \text{on } \partial\Omega_r. \end{cases} \quad (3.37)$$

This equation can be further transformed into the classical Laplacian operator. We use the final change of coordinates $u_1 = \sqrt{\lambda_1}u$ and $v_1 = \sqrt{\lambda_2}v$. Finally, since the MFPT is independent of the position, we get to leading order

$$\bar{\tau} = U = \frac{|S|}{\pi D_a} \ln \left(\frac{1}{\varepsilon} \right) + O(1), \quad (3.38)$$

where $|S|$ is the surface

$$|S| = \int du_1 du_2 = \sqrt{\lambda_1 \lambda_2} \int dv du = \sqrt{\lambda_1 \lambda_2} |S_g|, \quad (3.39)$$

where the surface of the network is $|S_g| = N3\sqrt{3}a^2$ when there N hexagonal fundamental domains of volume $3\sqrt{3}a^2$. For a small absorbing disk, $\varepsilon = \frac{a}{R}$, where R is the radius of the network, approximated as a disk of radius $R = a\sqrt{N}$. We conclude that the MFPT is

$$\tau \approx \frac{\sqrt{\lambda_1 \lambda_2} N 3\sqrt{3} a^2}{2\pi \frac{\lambda a^2}{2}} \ln \left(\frac{1}{\frac{a}{a\sqrt{N}}} \right) + O(1). \quad (3.40)$$

Thus,

$$\tau \approx \frac{3\sqrt{15}N}{2\pi\lambda} \ln N + O(1). \quad (3.41)$$

3.3.2 First arrival time in a honeycomb network

The first particles among $N_{\text{particles}}$ to arrive for the first time at a given node falls into the extreme statistics theory [Basnayake, Schuss, and Holcman 2018; Basnayake et al. 2018; Basnayake and Holcman 2019; Schuss, Basnayake, and Holcman 2019]. Using the operator 3.32, we can use the formulation of the Extreme Narrow escape [Basnayake, Schuss, and Holcman 2018] to derive the formula for the fastest among n to narrow opening which consists of few nodes ω_a forming a hexagonal unit in a two-dimensional hexagonal lattice.

Briefly, for $m = N_{\text{particles}}$ non-interacting i.i.d. Brownian trajectories moving on the hexagonal network Ω , the shortest arrival time τ^1 is by definition

$$\tau^1 = \min(t_1, \dots, t_m), \quad (3.42)$$

where t_i are the independent arrival times of the m particles in the network. The distribution of τ^1 is expressed in terms of a single particle,

$$P[\tau^1 > t] = P^m[t_1 > t]. \quad (3.43)$$

Here $P[t_1 > t]$ is the survival probability of a single particle prior to binding at the target site ω_a . In the approximation of a large network, this probability is computed by solving the diffusion equation [Schuss 2009]

$$\begin{cases} \frac{\partial p(\mathbf{x}, t)}{\partial t} = \mathcal{L}(p)(\mathbf{x}, t), & \text{for } \mathbf{x} \in \Omega, t > 0 \\ p(\mathbf{x}, 0) = p_0(\mathbf{x}) & \text{for } \mathbf{x} \in \Omega \\ \frac{\partial p(\mathbf{x}, t)}{\partial \mathbf{n}} = 0 & \text{for } \mathbf{x} \in \partial\Omega_r \\ p(\mathbf{x}, t) = 0 & \text{for } \mathbf{x} \in \omega_a. \end{cases} \quad (3.44)$$

The single particle survival probability is

$$P[t_1 > t] = \int_{\Omega} p(\mathbf{x}, t) d\mathbf{x}, \quad (3.45)$$

so that $P[\tau^1 = t] = \frac{d}{dt}P[\tau^1 < t] = m(P[t_1 > t])^{m-1}P[t_1 = t]$, where $P[t_1 = t] = \oint_{\partial\Omega_a} \frac{\partial p(\mathbf{x}, t)}{\partial \mathbf{n}} dS_{\mathbf{x}}$.

The probability density function (PDF) of the arrival time is

$$P[\tau^1 = t] = m \left[\int_{\Omega} p(\mathbf{x}, t) d\mathbf{x} \right]^{m-1} \oint_{\partial\Omega_1} \frac{\partial p(\mathbf{x}, t)}{\partial \mathbf{n}} dS_{\mathbf{x}}, \quad (3.46)$$

which gives the MFPT

$$\bar{\tau}^1 = \int_0^{\infty} P[\tau^1 > t] dt = \int_0^{\infty} [P[t_1 > t]]^m dt. \quad (3.47)$$

The shortest time estimation proceed as in [Basnayake et al. 2018]. They are expressed in terms of the shortest distance from the source S to the absorbing window ω_a , measured by the distance associated to the operator $\mathcal{L}(p)$, denoted by $\delta_{min} = d_{\mathcal{L}(p)}(S, A)$, where d is the associated distance here. Interestingly, the trajectories followed by the fastest among m are as close as possible from the optimal trajectories: the associated trajectories of the fastest among m , concentrate near the shortest path when the number m of particles increases. The mean extreme time for the hexagonal lattice,

$$\bar{\tau}_m^{hex} \approx \frac{\delta_{min}^2}{4D_a \ln \left(\frac{\pi\sqrt{2m}}{8 \ln(\frac{1}{a})} \right)}, \text{ for } \frac{m}{\ln(\frac{1}{a})} \gg 1 \quad (3.48)$$

where we recall that $D_a = \frac{\lambda a^2}{2}$. To conclude, we derived here an expression for the MFPT and the mean time for the fastest (MEFPT) associated to diffusion in two-dimensional honeycomb network.

3.4 Conclusions

In this chapter we have provided a detailed description and analysis of the properties of the AFN, comparing it to standard undirected networks. While exact calculations for the AFN are difficult to perform due to the correlation created by the stochastic directed flows, we have shown that it is possible to describe the main properties of the model by approximated and asymptotic expressions. These approximations show good agreement with numerical simulations. In particular, we have quantified the impact of the previously discussed trapping and backtracking mechanisms, and shown that the AFN model is characterized by a steady-state behaviour which is remarkably different from what is observed in undirected, diffusive networks.

CHAPTER 4

Quantitative analysis of ER protein dynamics

In this chapter we turn to the analysis of experimental data. We have seen in the previous chapters that protein transport in the endoplasmic reticulum (ER) could be characterized by local, transient phenomena. To observe these mechanisms experimentally, I needed a tool providing sufficiently high spatial and temporal resolution, capable of seizing both local heterogeneity and ensemble behaviour. To achieve this, I worked in collaboration with Christopher Obara, from Jennifer Lippincott-Schwartz' group at Janelia Research Campus, who curated the experimental part.¹ Together, we developed a system to track the ensemble dynamics of proteins in the ER at the whole-cell level, trying to overcome some of the limitations of FRAP and SPT. This new method showed that Brownian diffusion is still the primary mode of motion over large distances, but it can be affected by the ER structural properties. We reported the surprising observation that despite its continuity, the ER network is not always uniformly connected. We could also quantify differences in the redistribution timescales of molecules depending on their size, their topology, and their folding state, providing fundamental insight into the biophysical nature of the ER.

This chapter was adapted from the manuscript, still in preparation, "A quantitative mapping of diffusion in the endoplasmic reticulum", by Matteo Dora, Christopher J. Obara, Tim Abel, Jennifer Lippincott-Schwartz, and David Holcman.

4.1 Introduction

One of the most notable aspects of eukaryotic life is the development of membrane-enclosed organelles [Baum and Baum 2020]. These compartments provide two important advantages to cells. First, they segregate unique chemical environments that allow incompatible biochemical reactions to occur simultaneously, and second, they provide extensive two-dimensional surfaces (membranes) to facilitate the efficiency of biological reactions [Angermann et al. 2012]. The largest organelle in eukaryotes is the ER. It is a major site of cellular translation, protein folding, and quality control [Rapoport, Li, and Park 2017; Ron and Walter 2007], the major cel-

¹I must add that, despite its relative brevity, this project required a significant effort, in both the analysis and the experimental part. I cannot thank enough Christopher Obara for his marvellous, indefatigable work on the photoactivation experiments discussed in this chapter, and the exceptional, passionate support he provided along the route.

lular calcium store and sink in most cells [Berridge 2002], and the primary site of lipid synthesis and regulation [Van Meer, Voelker, and Feigenson 2008]. In addition to these crucial cellular functions, it also makes contact with and regulates the biology of essentially every other cellular compartment through specific tethering molecules [Prinz, Toulmay, and Balla 2020]. Thus, understanding how molecules are spatially distributed over time within the network is important for understanding how their functions are achieved.

The ER is a single, immense compartment that divides the nucleoplasm from the cytoplasm and stretches to the furthest reaches of the cell as a structurally complex membrane-enclosed network [Goyal and Blackstone 2013]. By volume, the compartment is relatively small, making up significantly less than 25 percent of the total cell size, but it contains more than 50 percent of all the cellular membrane [Heinrich et al. 2021; Valm et al. 2017]. The resulting high surface area to volume creates a compartment with a large degree of variability in local environment for both luminal and membrane-associated molecules. Despite occupying a small proportion of the cell volume at any moment, the network explores more than 97 percent of it every 15 minutes through dynamic structural rearrangements [Valm et al. 2017].

It remains a fundamental question how biological phenomena are spatially coordinated within the network, since there are no known ER-localized structural proteins or motors. The development of technologies for fluorescent tagging of molecules has opened the possibility to probe the dynamic nature of proteins within the ER [Ellenberg, Lippincott-Schwartz, and Presley 1999; Tsien 1998]. In particular, fluorescence recovery after photobleaching (FRAP) and fluorescence lifetime in photobleaching (FLIP) [Lippincott-Schwartz, Altan-Bonnet, and Patterson 2003] have provided significant insight into the way many components of the compartment behave [Lippincott-Schwartz, Snapp, and Phair 2018]. These studies have suggested that the composition of the organelle is strikingly dynamic, equilibrating on a timescale of seconds [Cole et al. 1996; Nehls et al. 2000]. However, these studies are limited to timescales of tens of seconds to minutes (FLIP) or single spatial regions (FRAP), and the ability to simultaneously observe high-speed ensemble dynamics across the entire cell has proven elusive. The development of photoactivatable fluorophores [Patterson and Lippincott-Schwartz 2002] has allowed this type of experiment in other systems [Hirschberg et al. 1998; Nichols et al. 2001], but most photoactivatable fluorophores do not fold correctly or provide low signal to noise ratio (SNR) in the high redox environment of the ER lumen [Costantini et al. 2015; Kaberniuk et al. 2018].

Single molecule tracking can provide information about very high speed dynamics in the most peripheral regions of cells [Manley et al. 2008]. However, it is confounded by the need for model fitting [Sun et al. 2022], and it cannot deal with the more complex three-dimensional environment that most of the cell volume is made up of. We have shown that high speed single molecule tracking data in the periphery suggests there are interesting transport phenomena over very fast time scales [Dora and Holcman 2020; Holcman et al. 2018], but we require a tool to link

this data to what is seen over larger spatial and temporal scales.

In this chapter, we have developed a combined experimental and analysis pipeline to surmount these issues. We have used photoactivation of protein-linked organic dyes that work well within the ER lumen to track the dynamic distribution of ER-localized molecules. This approach supports high speed imaging while also collecting entire cells. We then coupled this to a novel, machine learning-assisted analysis pipeline to account for the convolutions of the complex organelle structure. Collectively, this pipeline allows the extraction of quantitative information about the molecular behaviour within the ER at the whole-cell scale.

4.2 Results

In order to track ER protein dynamics across entire cells, we utilized photoactivatable organic dyes that can be directly coupled to a protein of interest by genetic fusion to a HaloTag [Encell et al. 2012; Grimm et al. 2016; Los et al. 2008]. These dyes are functional in the ER lumen and do not induce the misfolding and ER stress-related pathways common with fluorescent proteins not optimized for the high redox environment within this compartment [Costantini and Snapp 2013]. Once labelled with the dye, proteins of interest can be fluorescently marked at specific locations within the ER structure, and their distribution over time can be monitored by high speed microscopy. We performed experiments in cells also expressing a non-perturbative GFP-derivative targeted to the ER lumen [Costantini et al. 2015], to provide a simultaneously collected map of the underlying structure of the ER for downstream analysis purposes.

Briefly, COS-7 cells were cotransfected with a GFP-derivative targeted to the ER lumen to label the structure and a HaloTag fused to the protein construct of interest. We then used high speed spinning disk microscopy (at 10 Hz acquisition rate) using dual cameras to simultaneously collect the distribution of the photoactivated species of the dye and the fluorescence of the GFP label, resulting in two stacks of 2D images at 100 ms intervals. Photoactivation was carried out continuously at single, diffraction limited spot (photoactivation region of interest (ROI)), and the diffusion of molecules away from that spot was visible in the distribution of the converted dye molecule.

Ensemble diffusion governs dynamics of ER luminal content. We first set out to evaluate the dynamics of lumen-resident proteins that are freely floating (i.e. without membrane anchors or known specific binding partners). We assumed a globally isotropic but not necessarily diffusive motion, since local directional motion may be present [Holcman et al. 2018]. We used a HaloTag targeted to the ER lumen with a signal sequence and a KDEL retention motif as a model probe for free-floating luminal content (HaloTag-KDEL). Photoactivation was carried out at a single, diffraction-limited point (photoactivation (PA) source) while simultaneously

CHAPTER 4. QUANTITATIVE ANALYSIS OF ER PROTEIN DYNAMICS

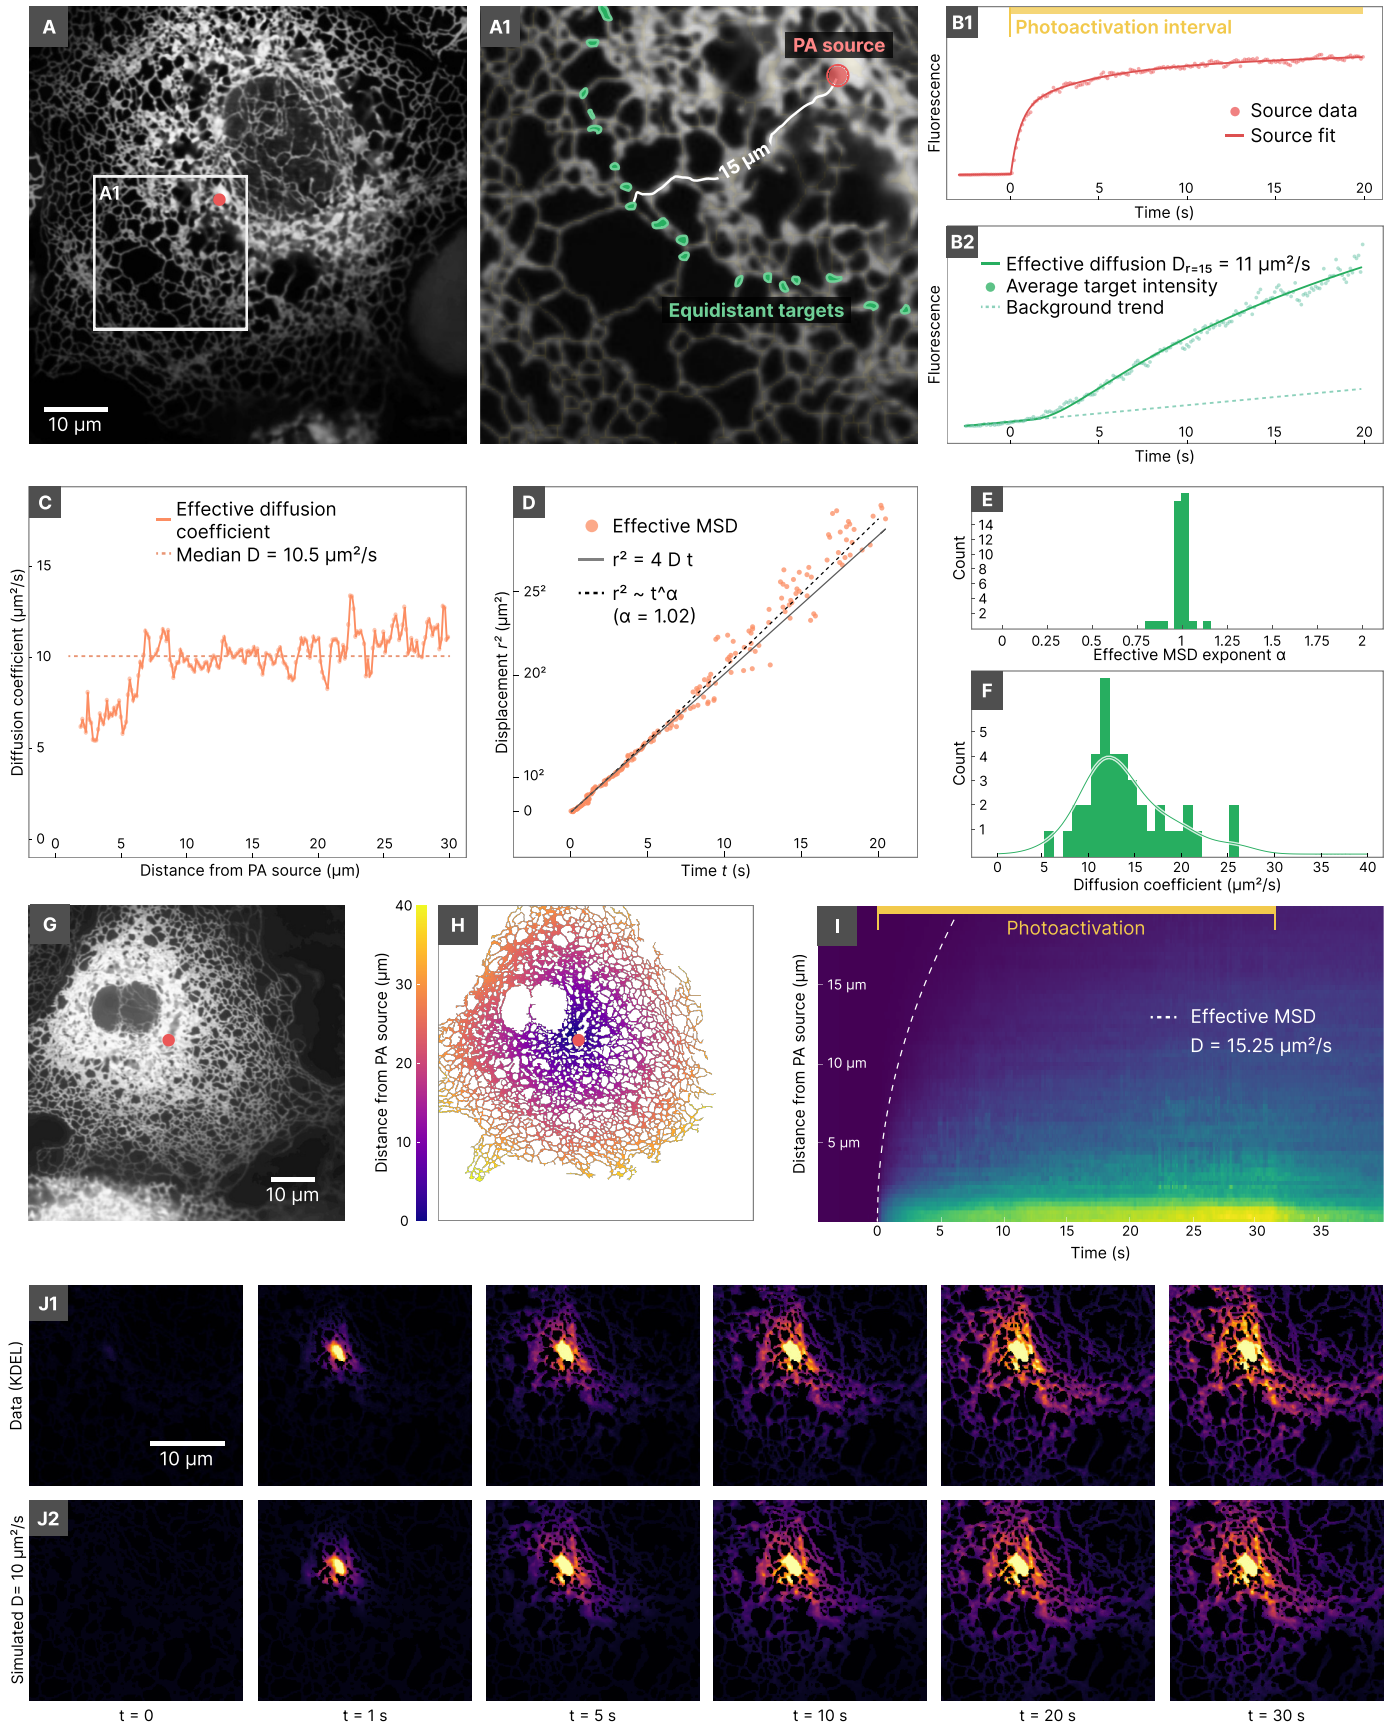


Figure 4.1 Brownian diffusion governs large scale ER luminal dynamics. **A** The ER structure of a COS-7 cell. The red dot indicates the location of photoactivation. Green regions correspond to ER domains that are equidistant ($d = 15 \mu\text{m}$) from the photoactivation region. **B1** Data and analytical fit of the fluorescence curve inside the photoactivation region. **B2** Average fluorescence of the target regions at distance $15 \mu\text{m}$ from the photoactivation source and fit of effective diffusion model ($D = 11.5 \mu\text{m}^2 \text{s}^{-1}$). **C** Effective diffusion coefficient as a function of the distance d from the photoactivation region, obtained from separate fits at each distance. **D** Effective mean-squared displacement (mean-squared displacement (MSD)). **E** Distribution of the effective MSD exponent of each cell, α , obtained from ensemble average over the cell ($N = 41$). **F** Distribution of the mean cellular diffusion coefficients obtained by ensemble estimate for HaloTag-KDEL in the ER lumen ($N = 41$). **G** Example of ER in a COS-7 cell. The photoactivation region is denoted by the red dot. **H** Structure of the ER network of the cell in G with colour scale denoting distances from the photoactivation source. **I** Kymograph of the average fluorescence in time of the cell in G, binned over the distance from the photoactivation source through the structure. The white dotted line indicates the bounds of the diffusion model that was fit. **J1** Experimental snapshots at different intervals from the start of photoactivation. **J2** Simulated snapshots built by evaluating the effective diffusion model on the ER shape obtained by experimental data.

collecting both the ER structure and the photoactivated signal using dual camera spinning disk microscopy (figure 4.1 A).

Proteins within the ER lumen are constrained to compartment by the bounding membrane, so they are not free to move in any arbitrary direction. Thus, analysis of the distribution patterns must account for the underlying organelle structure, or model fitting will produce artefacts [Sun et al. 2021]. To address this, we first segmented the ER by analysing the GFP channel to extract the network structure (see *Methods*, section 4.4.1 and figure 4.6). Regions of the ER were characterized by their distance from the PA source as measured through the underlying structure of ER (figure 4.1 A, inset).

Using this approach, the estimated arrival time of fluorescent molecules can be coupled with the distance from the PA source and the amount of photoactivation to produce an effective diffusion model. To do this, we first quantified the photoactivated fluorescence over time at the PA source by fitting a sum of saturating exponentials (figure 4.1 B1, and see *Methods*, section 4.4.3). This allowed us to constrain the effective diffusion model based on the measured intensity of the PA source. Once the source parameters have been determined, we performed a maximum-likelihood fit of the fluorescence intensity model to extract the effective diffusion coefficients corresponding to the average fluorescence measured in the equidistant points (figure 4.1 B2). The fluorescence intensity model also considers the possible background photoactivation (figure 4.1 B2, dashed line) due to inefficient photoactivation of the dye by the laser used to image the GFP (see *Methods*, section 4.4.3). Note that this approach can be modified to deal with sources of anisotropy if the model fit is not good, as would be predicted by synchronous directed flows or other micron scale sources of anisotropy (see below). We then repeated this procedure for every distance value in the structure, obtaining distance-dependent estimations of the effective diffusion coefficient (figure 4.1 C). In the immediate proximity of the PA source this approach can badly estimate the effective diffusion coefficient as a result of transient photoactivation by the activating laser,

but the measured value converges over distance to its true value (figure 4.1 C).

In order to evaluate the validity of a diffusion model, we regrouped the data in each cell along its time axis, creating an effective MSD curve describing the time evolution of the average distance travelled by a protein released from the PA source (figure 4.1 D). The exponent α of the MSD curve allowed us to evaluate possible deviations from standard diffusion ($\alpha = 1$) such as subdiffusion ($\alpha < 1$) or superdiffusion ($\alpha > 1$). For each cell, we determined the MSD exponent α by fitting the effective MSD data with $c t^\alpha$ (figure 4.1 D, dashed line). We found that population statistics of HaloTag-KDEL condition had an average MSD exponent $\alpha = 1$, thus suggesting that luminal dynamics can be well described by standard diffusion over this spatiotemporal scale (figure 4.1 E). We then defined a characteristic effective diffusion coefficient for each cell by considering the median value of the diffusion estimates between distance $2.5 \mu\text{m}$ and $20 \mu\text{m}$ from the PA source, where the data is most robust (figure 4.1 C, dashed line). We found for the HaloTag-KDEL condition a population diffusion coefficient $12.14 \mu\text{m}^2 \text{s}^{-1}$ (figure 4.1 F), which is very similar to previously reported estimates based on FRAP and single-particle tracking (SPT) [Holcman et al. 2018; Nehls et al. 2000; Snapp et al. 2006].

For organelles with less complex morphology, changes in the distribution of fluorescent molecules over time are often effectively represented with a kymograph, i.e. a time-space plot of the fluorescence observed along one dimension over time [Jakobs, Dimitracopoulos, and Franze 2019; Mangeol, Prevo, and Peterman 2016; Zhou, Brust-Mascher, and Scholey 2001]. In the ER, this is challenging due to the circuitous and dynamic nature of paths through the structure and the relatively low number of fluorescent molecules present in any isolated structural component. As a visualization tool, we used the distances calculated between the PA source and every point of the segmented ER structure (figure 4.1 G-H), and averaged the fluorescence values at each distance to create an equivalent kymograph describing the time-evolution of the fluorescence as a function of the distance from the PA source (figure 4.1 I). This representation shows concentrated photoactivated signal at the bottom (close to the PA source) and a gradual spreading through the structure moving upwards (towards distant regions) as the time of photoactivation increases. As a qualitative evaluation of the diffusion model fit to this cell, the time evolution of the fluorescence is well contained within the bounds predicted by pure diffusion at this effective diffusion coefficient.

To further confirm that our approach is appropriate, we simulated purely diffusive motion on the experimentally-obtained ER morphology (see *Methods*, section 4.4.3). We found that the simulated distribution could qualitatively reproduce the experimental photoactivation data with reasonable fidelity, suggesting that pure diffusion models are a good fit for ER luminal content at this spatiotemporal scale (figure 4.1 J).

Heterogeneity in apparent luminal mixing dynamics. The data presented thus far is analysed assuming a reasonable degree of isotropy in the system, since loca-

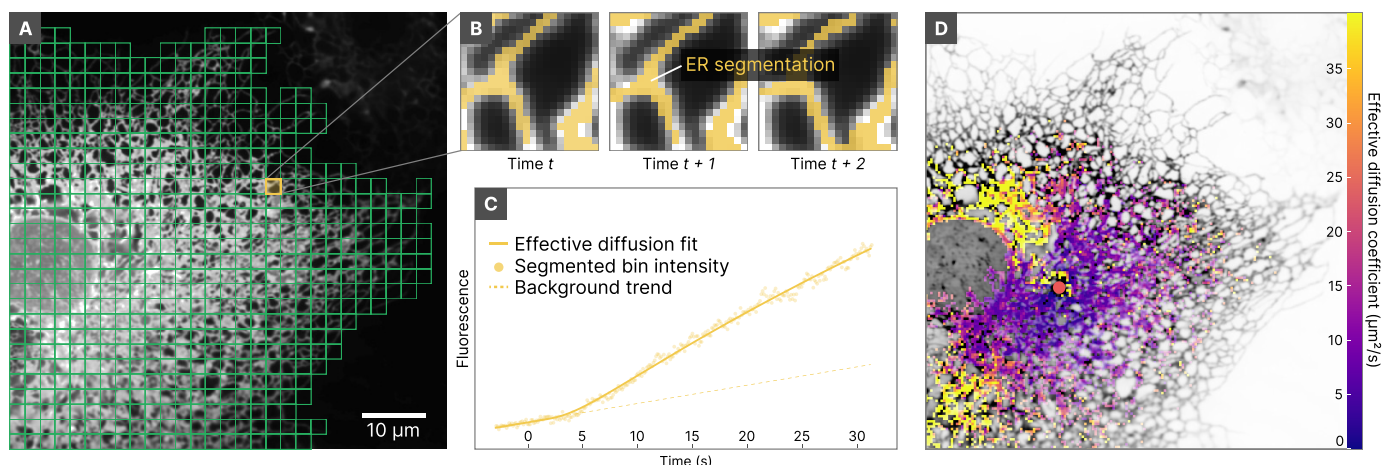


Figure 4.2 Heterogeneity in luminal mixing dynamics. **A** Subdivision of the ER using a grid of square bins. (A large bin size is used for easier visualization.) **B** An example of ER-associated, segmented pixels belonging to a given bin as they are tracked through time. **C** Average fluorescence over time for segmented pixels belonging to the bin shown in B. Dotted line shows the model fit. **D** Example of spatial map describing the observed mean arrival time, represented by a space-dependent effective diffusion coefficient (higher is faster). The red dot indicates the PA source. The grid bin size is 400 nm (3×3 camera pixels).

tions at similar distances from the PA source are averaged together. While this approach generally fits the data very well, we observed cells where the fit of the model was poor due to apparent anisotropy in the mean arrival times of fluorescent proteins.

To evaluate this heterogeneity, we decided to perform a spatially resolved model fit weighted by the amount of ER structure present. We first subdivided the field of view with a square grid of size 400 nm (figure 4.2 A). We then calculated the average fluorescence of the segmented ER portion contained in each grid bin for each frame in the temporal stack (figure 4.2 B). We discarded bins which did not contain ER structure for at least half of the time frames. We thus obtained a time evolution curve of the photoactivated fluorescence for each non-empty bin (figure 4.2 C). After fixing the PA source parameters as described above, we estimated an effective diffusion coefficient for each grid bin (see *Methods*, section 4.4.3), obtaining spatial maps of the effective diffusion coefficients (figure 4.2 D). Regions characterized by higher diffusion coefficients are associated with faster arrival of luminal proteins. We observed occasional cases of strong local heterogeneity in the dispersal of fluorescent proteins, represented by regions exhibiting faster arrival time with respect to the cell median mixing time (figure 4.2 D). Such regions seem to be localized in denser perinuclear regions of the ER, close to the nuclear envelope.

To ascertain whether this heterogeneity could be explained by an increased reshaping of the ER morphology in the affected regions, we quantified the morphology variation and tubule motion (see *Methods*, section 4.4.1). We found no association between higher morphological reshaping and regions with increased mixing

dynamics, suggesting that motion of the ER structure is unlikely to be the source of the observed anisotropy in the redistribution timescales.

Despite the variation in mean arrival times of fluorescent proteins, these cells still did not have MSD coefficients significantly different from 1 when considering the full cell data. Thus, this heterogeneity is likely to have to do with uneven connectivity within the structure rather than large scale directed motion (see *Discussion*). This is consistent with recently published three-dimensional electron microscopy volumes, showing significant variation in the mean luminal cross-section size in the perinuclear region of some cells [Heinrich et al. 2021].

Effects of membrane anchoring on ER protein diffusion. Early work in the literature showed that viscous drag of domains anchored in membranes is much higher than the drag on soluble domains. Consequently, membrane proteins generally show significantly lower effective diffusion coefficients than proteins freely floating within the lumen [Lippincott-Schwartz, Snapp, and Kenworthy 2001]. To examine the effect of this in our system, we compared the diffusion of HaloTag freely floating in the ER lumen (HaloTag–KDEL condition) and the same HaloTag anchored to the ER membrane (HaloTag–Sec61 β condition). The membrane anchoring was achieved by fusing HaloTag genetically to a minimal membrane targeting domain from Sec61 β that is known not to interact with the translocon in living cells [Nixon-Abell et al. 2016].

We evaluated both space-dependent and isotropic diffusion coefficients under both conditions. We found no significant difference between the two conditions regarding heterogeneity of the mixing dynamics, which were at least qualitatively comparable and isotropic in the majority of cells (e.g. figure 4.3 A1–A2, B1–B2). However, in agreement with the literature, the HaloTag–Sec61 β moved significantly less far through the structure for the same acquisition time (figure 4.3 A–B). As a result, spatial diffusion maps for the HaloTag–Sec61 β condition could typically only be evaluated in a smaller region around the PA source (up to 15 μm), where the SNR was sufficiently high with respect to the baseline. Estimation of the isotropic diffusion characteristics was consistent with published literature [Sun et al. 2021] and the previous results with HaloTag–KDEL (figure 4.3 C), showing that membrane dynamics was well described by standard Brownian diffusion (MSD exponent $\alpha = 1.02$, figure 4.3 D). Population level statistics for the effective diffusion coefficient (figure 4.3 D) confirmed that membrane proteins are diffusing significantly slower than luminal proteins in our system (membrane $D_{\text{Sec61}\beta} = 3.02 \mu\text{m}^2 \text{s}^{-1}$, compared to luminal $D_{\text{KDEL}} = 12.14 \mu\text{m}^2 \text{s}^{-1}$, p-value $< 1 \times 10^{-3}$).

Effect of protein size on mixing dynamics. Since the drag of a transmembrane anchor had such a dramatic effect on the apparent diffusion of molecules in the ER in our system, we wondered if our approach could be used to more finely quantify the effect of other sources of viscous drag. The Stokes–Einstein formula predicts that the mean diffusion coefficient of freely diffusing molecules should be in-

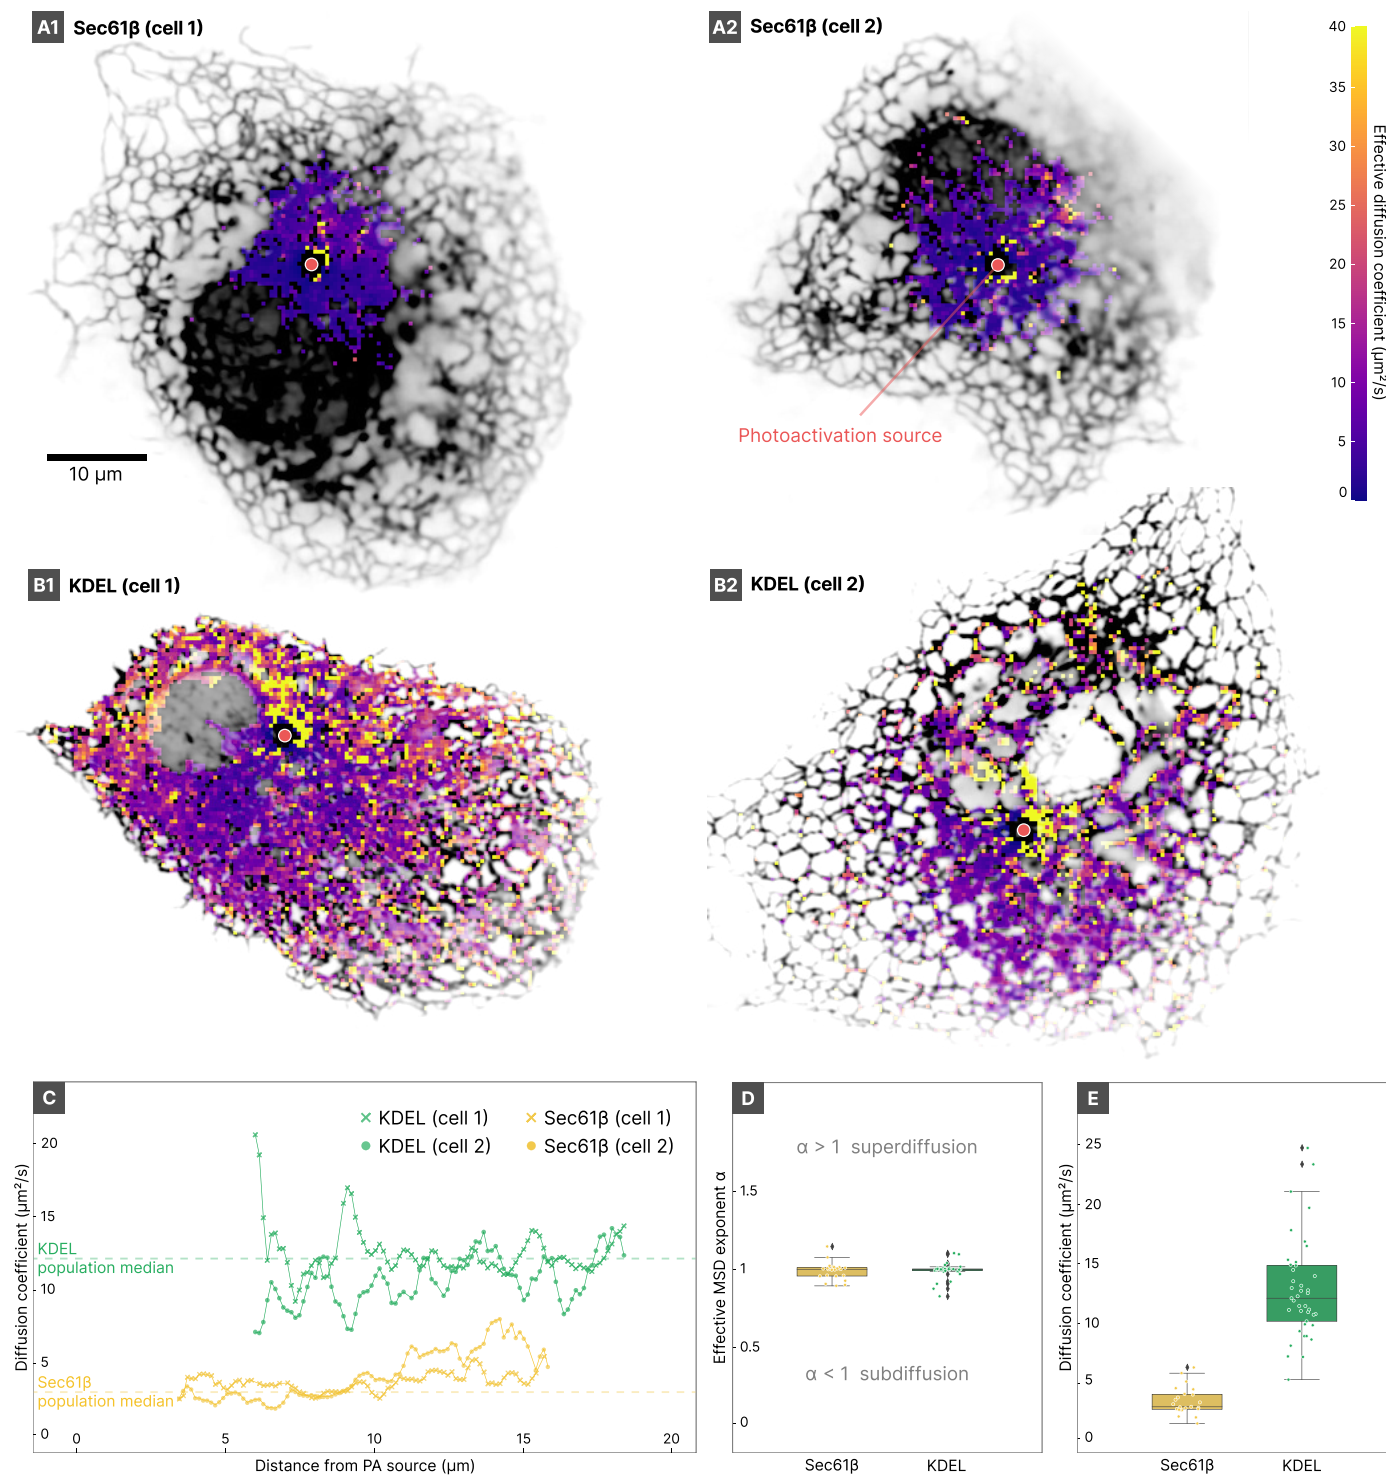


Figure 4.3 Comparison of ER protein dynamics in the lumen and in the membrane. **A** Examples of spatial diffusion map for HaloTag-Sec61 β . **B** Representative examples of the spatial diffusion map of HaloTag-KDEL. **C** Results for the isotropic diffusion fit (averaged over all targets at fixed distance) for cells A1–2 and B1–2. **D** Population statistics of the effective MSD exponent of HaloTag-KDEL and HaloTag-Sec61 β conditions. **E** Population statistics of the effective diffusion coefficient D_{eff} of HaloTag-KDEL and HaloTag-Sec61 β conditions.




Figure 4.4 Diffusion in the ER lumen but not the membrane is protein size-dependent. **A** Examples of spatial diffusion map for cells of 3×HALOTAG–KDEL. **B** Results for the isotropic diffusion fit (average over all targets at fixed distance) for the cells shown in A1 and A2 (3×HALOTAG–KDEL), compared to an example of HaloTag–KDEL. **C** Population statistics of the effective MSD exponent of for HaloTag–Sec61 β , 3×HaloTag–Sec61 β (control), HaloTag–KDEL, and 3×HALOTAG–KDEL. **D** Population statistics of the effective diffusion coefficient D_{eff} for the same conditions.

versely proportional to the radius of the molecule. We tested our ability to resolve this relationship by generating a larger version of a free-floating molecule in the ER lumen. Briefly, we genetically fused a signal sequence to three codon optimized concatemers of the HaloTag linked by flexible linkers and C-terminally fused to a KDEL retention sequence (3×HaloTag–KDEL).

As predicted for a system dominated by diffusional mobility, the photoactivated 3×HaloTag–KDEL explored a slightly smaller proportion of the cell in a defined time window than the single HaloTag–KDEL construct (compare, for example, figure 4.4 A with figure 4.3 B). Although a few cells did show some degree of anisotropy in arrival times of fluorescent molecules, the majority of the data fit well to an isotropic model. Again, we found that the dispersion of the 3×HaloTag was most effectively fit with a model for simple Brownian motion (effective MSD exponent $\alpha_{3\times}=1.01$, figure 4.4 B–C). (figure 4.4 A). However, in agreement with physical characteristics of diffusion, we observed decreased mixing speed for the 3×HaloTag–KDEL compared to single HaloTag–KDEL ($D_{3\times} = 8.22 \mu\text{m}^2 \text{s}^{-1}$, $D_{1\times} = 12.14 \mu\text{m}^2 \text{s}^{-1}$, p-value $< 1 \times 10^{-3}$, see figure 4.4 D). Notably, as a control, the same 3×HaloTag fused to the transmembrane targeting domain of Sec61 β had no detectable effect on its diffusion characteristics, since the viscous drag of the membrane on the membrane anchor is dominating (figure 4.4 C–D).

Dynamics of misfolded proteins in the ER lumen. One of the major functions of the ER is to serve as the major site of translation, folding, and quality control for membrane and secreted proteins [Rapoport, Li, and Park 2017; Ron and Walter 2007]. Studies using FRAP have demonstrated that during homeostasis the binding events of unfolded proteins to ER chaperones are transient and manifest as a slowed effective diffusion coefficient [Snapp et al. 2006]. However, if proteins are damaged or misfolded for too long, they can be removed from the ER and degraded through a process known as endoplasmic reticulum associated degradation (ERAD) [Ron and Walter 2007]. Studies with FRAP have suggested that under these conditions, proteins may become effectively immobilized [Lai et al. 2012; Shen et al. 2005]. Since our approach can accurately discern the small differences associated with protein size and provides a spatial-dependent estimate of diffusion that is not achievable with FRAP, we sought to try to observe how a misfolded protein would work in our system.

To this aim, we took advantage of the fact that the protein CD3 δ is obligately misfolded and degraded in normal tissue culture cells [Bernasconi et al. 2010; Bonifacino et al. 1989]. Truncation experiments have shown that luminal domain alone is sufficient for misfolding and degradation when expressed as a soluble construct [Bernasconi et al. 2010; Mehnert, Sommer, and Jarosch 2010]. We fused the signal sequence and soluble domain of CD3 δ to HaloTag (CD3 $\delta\Delta$ –HaloTag) to generate a photoactivatable version of a model misfolded protein. We then performed our full analysis pipeline, and analysed the dispersion of CD3 $\delta\Delta$ –HaloTag from the PA source.

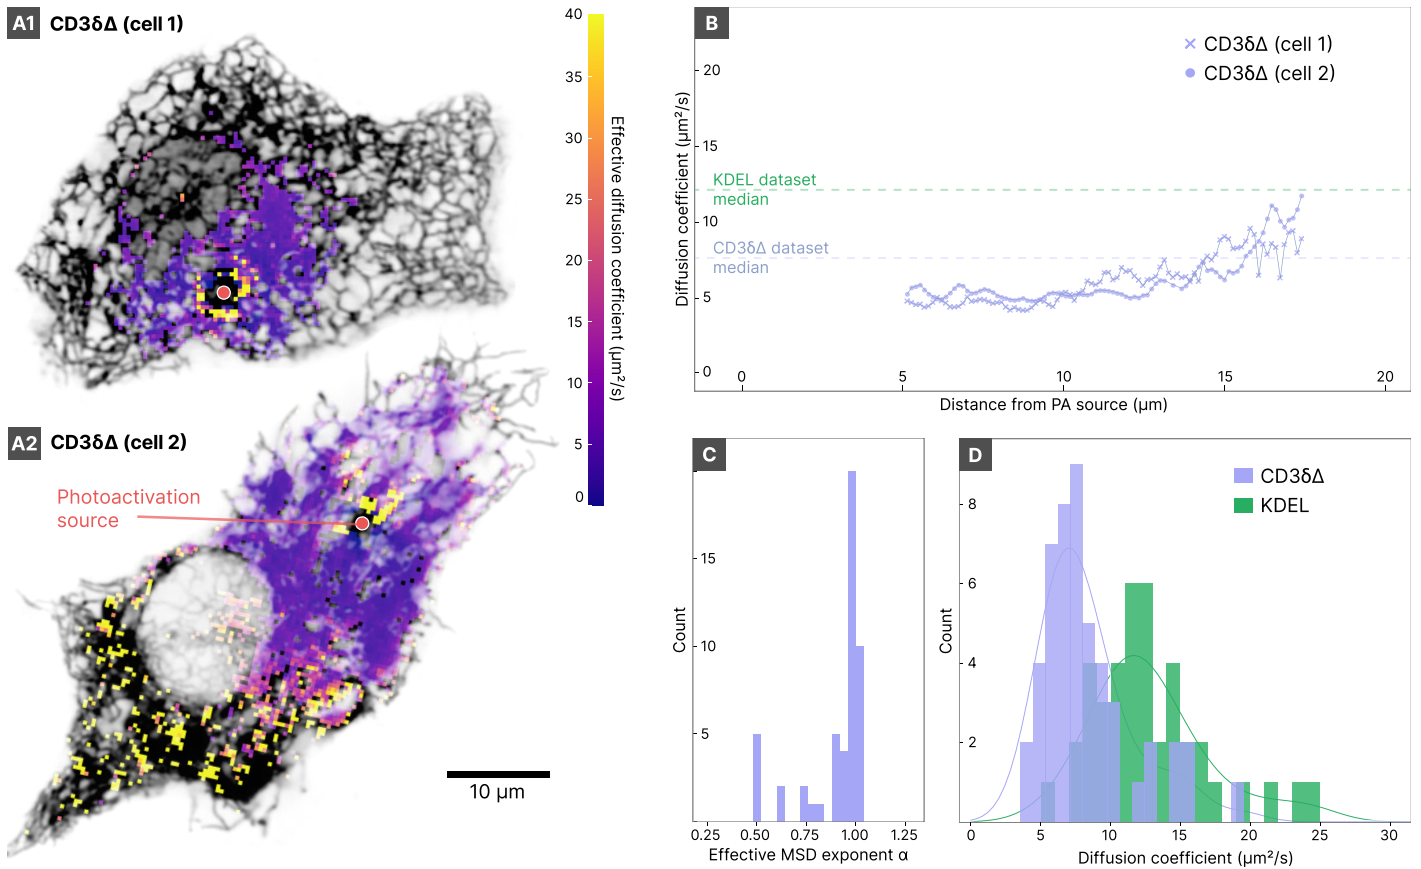


Figure 4.5 Misfolded luminal proteins show slowed effective diffusion and confined motion. **A** Representative examples of spatially defined diffusion of CD3δΔ-HaloTag. **B** Isotropic diffusion fit for the cells shown in A1 and A2 (CD3δΔ-HaloTag). **C** Distribution of the effective MSD exponent for the CD3δΔ-HaloTag condition. **D** Distribution of effective diffusion coefficient D_{eff} for CD3δΔ-HaloTag (lavender), compared to HaloTag-KDEL condition.

Examination of cells expressing the misfolded construct showed that a significant fraction of the protein was still dynamic within the system, allowing us to perform the full analysis pipeline (figure 4.5 A). Like the other constructs, most of the CD3δΔ-HaloTag-expressing cells fit well with an isotropic model. However, when analysed as a function of distance from the PA source, many CD3δΔ-HaloTag-expressing cells showed a trend towards higher diffusion at distances further away from the PA source (figure 4.5 B, see Discussion). The mean effective diffusion coefficient is significantly reduced for the misfolded protein ($D_{\text{CD3}\delta\Delta} = 7.57 \mu\text{m}^2 \text{s}^{-1}$, $D_{\text{KDEL}} = 12.14 \mu\text{m}^2 \text{s}^{-1}$, p-value $< 1 \times 10^{-3}$, see figure 4.5 D). Since this is predicted for a protein interacting with ER luminal chaperones, our data is consistent with that shown by FRAP [Snapp et al. 2006].

4.3 Discussion

In this chapter, we have combined high speed imaging with point-based photoactivation to examine the dynamic properties of proteins within the ER network. We have combined this with a machine learning-based segmentation pipeline and a spatially-defined model fitting procedure that can extract quantitative information about motion within the ER network. This approach supports evaluation of the primary mode of molecular motion at high speed across entire cells, and using this tool we have made a variety of observations about the physical factors that govern molecular behaviour within the ER. Many are consistent with the literature and suggest a dynamic, diffusion-dominated model for describing the distribution and mixing of proteins in the ER. However, a few observations are more surprising, and we discuss them briefly below.

Despite the evidence we have previously presented supporting locally directed motion of ER luminal proteins using SPT [Holcman et al. 2018], we were unable to detect any evidence in this study for superdiffusive motion in the ER lumen. We can think of three potential sources for this inconsistency. First, SPT is currently only able to be used at the outer periphery of cells, where our spatially defined model approach is the least robust. It may be that we could not observe a significant contribution of the ER flow to our data because it is simply not present in enough of the structure to be detected when averaging. Second, the spatial and time scales of locally-directed luminal flows are still a subject of some debate, but they may exist over scales that are beneath the resolving power of spinning disk microscopy as we have implemented it here. Thus, they are only observed in SPT-based approaches. Third, ER flows are not universal or only occur in some situations. Either way, resolving this issue will require more data collected by both SPT and our approach, ideally in the same system simultaneously. This will be an attempted target of future work.

Another surprising observation made in this study is that in some cells, there is significant anisotropy in the arrival times of fluorescent proteins. This was true to some extent for all the conditions, though it was always a small fraction of the total cells. In general, poor model fitting of this nature might be a predicted result of non-diffusive motion, but there are a few reasons we do not feel this is the most parsimonious explanation in this circumstance. First, these cells do not have effective MSD exponent fits that are significantly greater than 1, as would be predicted if such significant regions of the cell were undergoing directed motion. Second, the location of the regions is almost always in the perinuclear region immediately juxtaposed to the nuclear envelope. This region of ER in tissue culture cells is astoundingly structurally complex and cannot be resolved by light microscopy [Heinrich et al. 2021; Nixon-Abell et al. 2016; Xu et al. 2021]. Thus, it may be that in these complex regions our segmentation becomes a poor predictor of actual connectivity. Either way, it suggests that the ER of cells is not guaranteed to be evenly connected, which may carry important implications for many biological sorting processes. A

search for proteins that may affect this connectivity will be a focus of future work.

Although it was not universal, many cells showed a slight upwards trend in effective diffusion coefficient at distances quite far from the PA source. This was notable in the CD3 δ Δ examples shown in figure 4.5, but was present in some cells in each of the other conditions as well. We do not have a clear explanation for this phenomenon, but one possibility is that, as we increase the distance from the PA source (which was nearly always placed in the dense central regions of the cell) and go towards the ER periphery, the connectivity is reduced, making the 2D effective diffusion model inexact as the dimensionality gets closer to 1 (e.g. for the case of a molecule moving along a single tubule). Moreover, there is building data that ER tubules at the extreme periphery can be very thin [Terasaki 2018; Wang et al. 2021]. As the tubes get narrower, they may funnel molecules more efficiently to the end than in thicker tubes with more space for lateral motion. This will be an interesting direction for SPT experiments in the future.

It is an interesting result of this study that the fusion of the small misfolded domain of CD3 δ Δ slows effective HaloTag diffusion by a similar effect to that of linking three HaloTags as concatemers. A weakness of our ensemble approach to analysing diffusion is that a molecule with multiple distinct forms only reports the average of these states. It is impossible to say from this data what proportion of the molecules are in confined states in any given cell, but future experiments combining SPT or carefully designed drug treatments may be able to parse the contributions of chaperone binding from degradation-associated confinement in the net slowing of the CD3 δ Δ dispersal.

In this work, we have developed a quantitative approach to monitor high-speed protein dynamics within the ER. We see that ER proteins in the lumen and the membrane are largely governed by bulk diffusion, and there are many predictable aspects of membrane topology and size on the resulting molecular motion.

4.4 Methods

Fluorescence microscopy images are acquired in low-light conditions, since the power of the excitation laser and exposure interval are constrained by the fluorescence saturation rate, the risk of photodamaging the sample, and the acquisition rate for dynamical imaging. Fluorescent images are thus characterized by high level of noise. In particular, the main source of disturbance is represented by Poisson noise (or *shot noise*) originating from the discrete number of photons hitting the sensor [Foi et al. 2008]. This is combined with other signal-independent noise sources, due to thermal vibration (e.g. dark current) or caused by the electronics (i.e. noise due to the amplification, conversion, and transmission process, which is usually collectively indicated as *readout noise*). These last noise sources do not depend on the photon flux, and can be modelled effectively as a single aggregated white noise source. Thus, denoising of fluorescence images requires handling Poisson-

Gaussian noise. Various techniques have been developed to solve this problem, following mainly three approaches: variance stabilizing transformation followed by additive white noise removal [Foi et al. 2008], algorithms specifically designed for Poisson-Gaussian mixtures such as PURE-LET [Luisier, Blu, and Unser 2011], and deep learning models based on convolutional neural networks (CNNs) [Lehtinen et al. 2018; Zhang et al. 2017]. A recent work, benchmarking the most popular methods on a dataset of fluorescence images, highlighted the superior performances of CNN-based models [Zhang et al. 2019]. The main drawback of CNNs is that they usually require training on a set comprising both noisy and clean images. However, in the last few years various deep learning models have been proposed which allow training without clean data [Batson and Royer 2019; Lehtinen et al. 2018; Ulyanov, Vedaldi, and Lempitsky 2018]. Another drawback of CNNs is the phenomenon of pareidolia, i.e. the tendency to impose a meaningful interpretation to random or ambiguous patterns. CNNs are, by their nature, biased towards certain structures which can be intrinsic or derived from the training data [Mitchell 2017; Ulyanov, Vedaldi, and Lempitsky 2018]. While this is the very reason of their effectiveness, it can also be the cause of unwanted visual artefacts: a CNN model trained on ER images will develop a tendency to see ER even where there is none. For these reasons, in our analysis of photoactivation data, we have taken a mixed approach: we relied on a deep-learning model to facilitate the segmentation of the ER morphology, but followed a more conservative approach when extracting quantitative measures from the photoactivation curves, considering a statistical model of the raw data.

4.4.1 Pre-processing and segmentation of ER morphology

To ensure reliable segmentation of the ER structure, we preprocess the data of the morphology channel (green fluorescent protein (GFP)-KDEL dye) in several steps (see fig. 4.6). These preprocessing steps are needed because the ER structure and density can vary significantly. Regions of the ER close to the nuclear envelope are characterized by denser structures (membrane sheets or dense matrices of tubules [Nixon-Abell et al. 2016; Terasaki, Chen, and Fujiwara 1986]) which are not resolvable with confocal microscopy and thus appear as continuous membrane arrangements. At the cell periphery, the ER is formed by a network of sparse tubules with diameter 60–100 nm [Nixon-Abell et al. 2016; Shibata, Voeltz, and Rapoport 2006]. These structural differences determine the quantity of fluorescent proteins and worsen the image contrast, making it challenging to segment simultaneously the entire ER structure [Pain et al. 2019].

Frame denoising. We first denoised the morphology data to reduce the level of noise and improve the detection of finer structures. To this aim, we adapted the Noise2Noise technique [Lehtinen et al. 2018] to train a U-Net convolutional neural network [Ronneberger, Fischer, and Brox 2015] for noise removal of single frames. We trained U-Net model using pairs of noisy realizations generated by extracting

128×128 px tiles from time-adjacent frames. One noisy realization is kept as a reference while the other is processed by the U-Net model. The loss is then calculated as the mean-squared error (MSE) between the reference noisy frame and the model output. This approach relies on the assumption that the acquisition rate is fast enough to make changes between consecutive frames negligible. The assumption is reasonable as most of the ER structure is steady for the acquisition interval used (100 ms). To minimize the bias created by the time evolution of the ER network, both the order of the training samples and the order of the two images forming a couple were randomized in the repeated training epochs. We keep 10% of the samples for validation, excluding them from training. The training is performed on batches of size 4, using the Adam optimization algorithm [Kingma and Ba 2017] with parameters $\beta_1 = 0.9$, $\beta_2 = 0.99$, no weight decay and initial learning rate $\gamma = 2.75 \times 10^{-4}$. The learning rate was then gradually reduced with a factor 0.75 when the MSE loss calculated on the validation set did not improve over 5 epochs. To prevent overfitting, training was stopped when the loss over the validation set did not improve for 15 consecutive epochs, selecting the model with minimal loss.

Sharpening and structure enhancement. To enhance the resolution of the finer ER structure, we deconvolved the morphology frames to create sharper images. Image deconvolution is an inverse problem of the type $\mathbf{y} = \mathbf{A}\mathbf{x}$, consisting in recovering the sharp image \mathbf{x} from a blurred version \mathbf{y} , with \mathbf{A} representing the point spread function (PSF) operator associated to the optical system. Since the problem is ill-posed, the solution is found by adding regularization constraints on \mathbf{x} such as total variation [Getreuer 2012]. Recently, a sparsity constraint in the deconvolution procedure has been shown to successfully increase the spatio-temporal resolution in fluorescence microscopy [Zhao et al. 2021]. Based on these approaches, we defined a spatio-temporal deconvolution problem with total variation and sparsity regularization terms, to provide good reconstruction of edges, and a smoothness constraint in time dimension to account the gradual structural movements. We define the optimization problem as

$$\arg \min_{\mathbf{x}} \left(\frac{\mu}{2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \lambda_t \|\mathbf{x}\|_1 + \lambda_v \sum_{n=1}^N \|\nabla_{\mathbf{v}_n} \mathbf{x}\|_1 + \lambda_2 \|\nabla_t^2 \mathbf{x}\|_1 \right) \quad (4.1)$$

where \mathbf{y} represents a sequence of frames, \mathbf{A} is the PSF operator, $\|\cdot\|_1$ and $\|\cdot\|_2^2$ represent respectively the ℓ_1 and ℓ_2 norms, $\nabla_{\mathbf{v}_n}$ is the directional derivative along \mathbf{v}_n , and ∇_t^2 is the second derivative in the time dimension. Multiple spatial directions $\mathbf{v}_n = (\sin(2\pi n/N), \cos(2\pi n/N), 0)$ in space (x, y, t) to reduce staircase artefacts. To ensure correct enforcement of the sparsity constraint, each frame in \mathbf{y} was preprocessed by taking a 2D discrete wavelet transform and dampening the approximation coefficients by a factor 0.5. The operator \mathbf{A} was assumed to be a Gaussian PSF with appropriate standard deviation. The solution \mathbf{x} was computed with the split Bregman method [Goldstein and Osher 2009]. Parameter values used in the optimization are reported in table 4.1.

Parameter	Value
PSF σ	1.27 px
μ	2
λ_1	0.1
λ_v	0.05 / N
λ_t	1
Split Bregman inner iterations	2
Split Bregman outer iterations	5

Table 4.1 Parameter values for morphology deconvolution. See equation 4.1

Segmentation of the network structure. The ER network was segmented from the sharpened stack as follows. First, the frames were clipped at zero and rescaled between 0 and 1 and local contrast was enhanced by contrast limited adaptive histogram equalization (CLAHE) [Heckbert 1994] for five iterations. To detect tubular structures, the frames were first transformed with a Sato tubeness filter [Sato et al. 1998] with $\sigma = 100$ nm and then binarized by locally adaptive Niblack thresholding [Niblack 1985] with window size 127 px ($\simeq 17 \mu\text{m}$) and weight 0.1. Empty background regions, which could create bad detections with purely local thresholding, were excluded by setting a minimum global threshold. Continuous surfaces corresponding to dense matrices and cisternæ were obtained by a similar thresholding procedure, followed by a morphological opening operation with a disk of radius 250 nm. The thresholded tubules and sheets were then merged in a single frame sequence. To preserve details in very close tubular structure, local minima were identified with a h-minimum transform (with threshold 1×10^{-3}) before and set to zero [Pain et al. 2019]. To fix occasional flickering of isolated pixels, a morphological closing operation was applied along the time axis of the thresholded stack. Finally, the network structure was extracted by thinning procedure [Zhang and Suen 1984]. To preserve dense matrices and cisternæ, these regions were again extracted by morphological opening and added back to the skeletonized network structure.

Estimation of morphology reshaping. The amount of local morphological reshaping was quantified in two ways. First, by estimating the relative variance of the denoised morphology frames along the time stack, highlighting regions of the network undergoing more variation in time. Predictably, most of the variation was localized on the tubule edges, due to small oscillations (figure 4.7 B, blue shade). Note that in our method, such small oscillations do not affect the estimate of the dynamics, as only the pixels clearly belonging to the ER structure are considered.

Second, we estimated motion of tubules and other structures by optical flow using Farnebäck’s algorithm [Farnebäck 2003], then calculated the mean norm of the velocity vectors over the photoactivation time to reveal regions characterized by more active reshaping of the network structure (figure 4.7 B, red shade).

4.4.2 Noise model of the photoactivatable channel

Estimation of the noise model parameters. We considered a Poisson-Gaussian noise model with clipping as described in [Foi 2009a; Foi et al. 2008]. The original unclipped noise model which describes the observed pixel value $z(x)$ at position x is defined by

$$\begin{aligned} z(x) &= y(x) + \sigma(y(x)) \xi(x) \\ \sigma^2(x) &= ay(x) + b \end{aligned} \quad (4.2)$$

where $y(x)$ is the true pixel value, $\xi(x)$ is zero-mean and unit variance random noise, and $\sigma(y)$ is the y -dependent standard deviation function characterizing the Poisson-Gaussian noise. Note that this model assumes that the photon count y is large enough, such that the Poisson term can be approximated with a normally distributed noise with signal-dependent variance. Imaging sensors are also characterized by a bias level μ , i.e. an artificial offset added to the collected charge to reduce clipping. This offset value must be subtracted from the readout to correctly model the signal-dependent part of the noise, e.g. by considering the variance function $\sigma(y - \mu)$ or by redefining $y \leftarrow y - \mu$. If the sensor is not able to capture the full dynamic range of the image (particularly in the case of fluorescence imaging, where frames are intrinsically underexposed), one can define a clipped observation model as:

$$z_c(x) = \max\{0, \min\{z(x), 1\}\}. \quad (4.3)$$

The goal is to reconstruct the model parameters a, b from the video data. We applied the technique described in [Foi et al. 2008] on all frames, segmenting the full video data in level sets, excluding the tubules boundaries obtained from the segmentation of the morphology channel, to estimate $(\hat{y}_i, \hat{\sigma}_i)$ pairs. We then performed a maximum likelihood fit of the model parameters using the L-BFGS algorithm.

4.4.3 Estimation of diffusion coefficients

Effective diffusion model. To extract quantitative features describing the protein dynamics, we consider an effective 2D diffusion model. The approximation with a 2D space is appropriate given that the ER structure in COS-7 cells can, for most of its surface, be described by a planar graph [Nixon-Abell et al. 2016]. While it is possible to derive a reaction-diffusion system which models a constant-rate photoconversion in the source ROI (the derivation of such a model is given in appendix 4.A), this approach requires knowledge (or fit) of the photoconversion rate parameter. Moreover, due to the reshaping of the ER, the ROI exposed to the photoactivation laser may change constantly, altering the effectiveness of photocon-




Figure 4.6 Morphology processing pipeline. **A** Data denoising with Noise2Noise method. From the stack of frames (**A1**), pairs of noisy tiles are extracted (**A2**). The UNet model is trained using the mean-squared error between the model prediction and the other noisy frame (**A3**). **B1** Example of raw morphology frame. **B2** Frame denoised with the UNet model. **B3** Sharpened and enhanced frame with the deconvolution of eq. 4.1. **B4** Final segmentation of the ER structure.

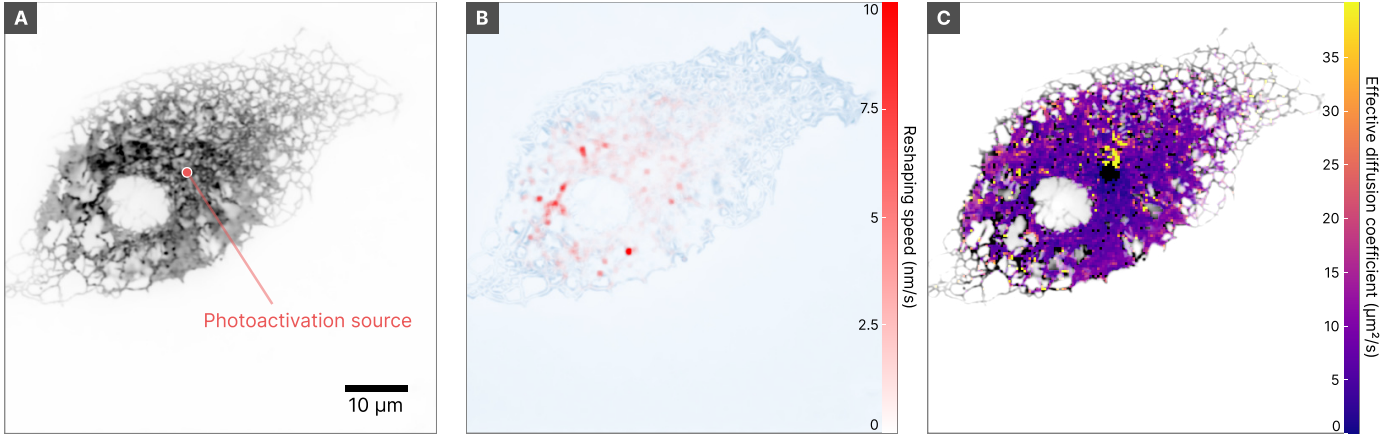


Figure 4.7 Estimation of morphology reshaping. **A** Cell morphology. **B** Reshaping intensity estimated by optical flow (red) and relative variance (blue shade). **C** Comparison with the diffusion map for the same cell.

version. To avoid these problems, we used a hybrid approach which allows to preemptively tune the diffusion model based on the data. To this aim, we considered a diffusion model in which we impose a time-dependent Dirichlet boundary condition (see appendix 4.4.3 for details) corresponding to the fluorescence measured close to the photoactivation ROI. We denote by $\varphi(t)$ the protein concentration measured at distance R from the centre of the photoactivation ROI. Given an analytical expression for $\varphi(t)$, we can calculate the solution of the radial diffusion equation $u(r, t)$ (see appendix 4.4.3, eq. 4.35):

$$u(r, t) = \int_0^t d\tau \varphi(t - \tau) \partial_t v(r, t), \quad (4.4)$$

where $v(r, t)$ is the solution of the auxiliary diffusion problem with a constant boundary condition ($\varphi(t) = 1$), which can be explicitly expressed in Laplace domain as

$$\hat{v}(r, s) = \frac{K_0(\sqrt{s/D}r)}{s K_0(\sqrt{s/D}R)}, \quad (4.5)$$

where K_0 is the modified Bessel function of order zero. It is convenient to express the convolution in equation 4.4 as a product in the Laplace domain:

$$\hat{u}(r, s) = \hat{\varphi}(s) \frac{K_0(\sqrt{s/D}r)}{K_0(\sqrt{s/D}R)}, \quad (4.6)$$

calculating the $u(r, t)$ by inverting the Laplace transformation.

In principle, we can choose any analytical form of $\varphi(t)$ that well describes the fluorescence curve at distance R from the centre of the photoactivation ROI observed in the experimental data. We considered here the particular case where $\varphi(t)$ can be

expressed as a sum of exponentially saturating functions:

$$\varphi(t) = \sum_n c_n (1 - e^{-\lambda_n t}), \quad (4.7)$$

$$\hat{\varphi}(s) = \sum_n c_n \left(\frac{1}{s} - \frac{1}{s + \lambda_n} \right). \quad (4.8)$$

Using this definition for φ combined with eq. 4.6 we can write

$$\hat{u}(r, s) = \sum_n c_n \left(\frac{1}{s} - \frac{1}{s + \lambda_n} \right) \frac{K_0(\sqrt{s/a} r)}{K_0(\sqrt{s/a} a)}, \quad (4.9)$$

from which $u(r, t)$ can be obtained by inverting the Laplace transform numerically.

Fluorescence intensity model. To properly describe the fluorescence data, we need to take into account variations which not related to the diffusion process. These are mainly constituted by the background intensity and the random photoconversion of proteins outside the photoactivation region due to leaks from the other channel and acquisition laser. We describe these by an offset and a linear trend which can be observed before the photoactivation laser is activated. Thus, we describe the fluorescence intensity $f(\mathbf{x}, t)$ at position \mathbf{x} as the sum of a linear polynomial $c_1 t + c_0$ and the diffusion process $u(r, t)$:

$$f(\mathbf{x}, t) = c_2(\mathbf{x}) u(|\mathbf{x} - \mathbf{x}_0|, t) + c_1(\mathbf{x}) t + c_0(\mathbf{x}) \quad (4.10)$$

where \mathbf{x}_0 is the location of the PA ROI, $c_0(\mathbf{x})$ describes the background luminosity, $c_1(\mathbf{x})$ is the slope of the linear trend of random photoactivation, and $c_2(\mathbf{x})$ is a proportionality constant representing the density of photoactivated proteins. Note that, in general, the coefficients c_0 , c_1 , and c_2 are space-dependent.

Maximum-likelihood fit of the diffusion model. We consider the intensity data obtained from the raw images in the form of points $x_i = (r_i, z_i, t_i)$, where r_i represents the distance from the photoactivation source and z_i its measured fluorescence at time t_i from the beginning of photoactivation. We then use the noise model defined in equation 4.2, which has been previously estimated according to the procedure described above, to estimate the likelihood of the fluorescence model of equation 4.10 depending on the x_i samples. The noise variance for a point $x_i = (r_i, z_i, t_i)$ is given by

$$\sigma_i^2 = a(z_i - \mu) + b, \quad (4.11)$$

where μ is the imaging sensor bias level, and the log-likelihood is obtained as

$$\log \mathcal{L}(\theta) = \sum_i 2 \log \sigma_i + \frac{(z_i - f(x_i, t_i; \theta))^2}{\sigma_i^2} \quad (4.12)$$

with θ is the set of parameters c_0, c_1, c_2, D . The maximum likelihood parameter is obtained by solving the problem

$$\hat{\theta} = \arg \max_{\theta} \log \mathcal{L}(\theta), \quad (4.13)$$

which we solve numerically by L-BFGS. The same maximum likelihood estimation can be extended to ensemble points obtained by averaging multiple pixel values in the photoactivation stack. The noise variance of such aggregated points is simply

$$\hat{\sigma}_i^2 = \frac{a(\tilde{z}_i - \mu) + b}{N}, \quad (4.14)$$

where N is the number of averaged pixels and \tilde{z}_i is average fluorescence value. By maximum-likelihood estimation it is then possible to recover, beyond D , the coefficients \tilde{c}_0, \tilde{c}_1 , and \tilde{c}_2 corresponding respectively to the average background intensity, average slope of the linear trend, and average density over the ensemble.

In the case of isotropic diffusion estimation, these ensemble points are built by binning pixels in each frame based on their distance from the PA source (with bin size corresponding to the microscope resolution of 133 nm). Only pixels belonging to the segmented ER structure are considered. Ensemble points characterized by the same distance r are grouped to build a time series $(r, \tilde{z}_i, t_i)_{i=1,2,\dots}$. We note that, due to changes in the morphology of the ER, the pixels belonging to a given distance bin are not invariant across frames. This frame by frame adjustment allows cancelling at least part of the effect of the morphology reshaping.

For the estimation of space-dependent effective diffusion, ensemble points are built by dividing the field of view in equal bins using a squared grid (figure 4.2 A). For each bin, average fluorescence intensity and median distance from the source are evaluated in each frame (considering only the pixels belonging to the segmented ER structure, figure 4.2 B). The ensemble points obtained by this procedure are then used to estimate the effective diffusion by maximum likelihood (figure 4.2 C), obtaining a diffusion coefficient for each bin of the grid (figure 4.2 D).

4.A Reaction-diffusion modelling of photoactivation

We consider the dynamics of proteins moving onto a 2D surface which are photoactivated by a laser beam hitting the sample on a disk of radius R . We can describe such system by the following reaction-diffusion equation:

$$\begin{cases} \partial_t u(r, t) - \alpha \left(\partial_r^2 u(r, t) + \frac{1}{r} \partial_r u(r, t) \right) = f(r, t) & t > 0 \\ u(r, t = 0) = 0, \end{cases} \quad (4.15)$$

where $f(r, t)$ describes the process of photoactivation of inactive particles by the laser beam with rate κ inside the disk of radius R :

$$f(r, t) = \begin{cases} \kappa (1 - u(r, t)) & r < R \\ 0 & r \geq R \end{cases} \quad (4.16)$$

We solve eq. 4.15 separately for the two cases $r < R$ and $r \geq R$. In the case $r \geq R$, we have the homogeneous partial differential equation:

$$\partial_t u(r, t) - \alpha \left(\partial_r^2 u(r, t) + \frac{1}{r} \partial_r u(r, t) \right) = 0. \quad (4.17)$$

By taking the Laplace transform, we obtain the modified Bessel equation

$$z^2 \partial_z^2 \hat{u} + z \partial_z \hat{u} - z^2 \hat{u} = 0, \quad \text{with } z = \sqrt{s/\alpha} r, \quad (4.18)$$

which has general solution

$$\hat{u}_{r \geq R}(r, s) = c_1(s) I_0(\sqrt{s/\alpha} r) + c_2(s) K_0(\sqrt{s/\alpha} r). \quad (4.19)$$

In the case $r < R$, we need to solve the nonhomogeneous equation:

$$\partial_t u(r, t) - \alpha \left(\partial_r^2 u(r, t) + \frac{1}{r} \partial_r u(r, t) \right) = \kappa (1 - u(r, t)). \quad (4.20)$$

By taking Laplace transform, we obtain

$$z^2 \partial_z^2 \hat{u} + z \partial_z \hat{u} - z^2 \hat{u} = -\frac{\kappa}{s(s+\kappa)} z^2, \quad \text{with } z = \sqrt{\frac{s+\kappa}{\alpha}} r. \quad (4.21)$$

We can write the general solution of eq. 4.21 as the sum of the solutions of the corresponding homogeneous equation and a particular solution ψ_{part} . The corresponding homogeneous equation is a Bessel modified equation with solutions $I_0(z)$ and $K_0(z)$. We search a particular solution constant in r and find $\psi_{\text{part}} = \frac{\kappa}{s(s+\kappa)}$. The solution for the case $r < R$ is then:

$$\hat{u}_{r < R}(r, s) = c_3(s) I_0 \left(\sqrt{\frac{s+\kappa}{\alpha}} r \right) + c_4(s) K_0 \left(\sqrt{\frac{s+\kappa}{\alpha}} r \right) + \frac{\kappa}{s(s+\kappa)}. \quad (4.22)$$

Putting the two solutions together, we can write the solution of eq. 4.15 in Laplace domain:

$$\hat{u}(r, s) = \begin{cases} \hat{u}_{r < R}(r, s), & 0 \leq r < R \\ \hat{u}_{r \geq R}(r, s), & r \geq R. \end{cases} \quad (4.23)$$

We now determine the coefficients c_1, c_2, c_3, c_4 . First, we require the solution to be bounded for all values of r : we set $c_4 = 0$ to prevent the solution to diverge at $r = 0$, and $c_1 = 0$ to prevent the solution to diverge for $r \rightarrow \infty$. To determine the remaining coefficients we impose the continuity of the solution and its first derivative at $r = R$:

$$\hat{u}_{r < R}(R, s) = \hat{u}_{r \geq R}(R, s), \quad (4.24)$$

$$\partial_r \hat{u}_{r < R}(r, s)|_{r=R} = \partial_r \hat{u}_{r \geq R}(r, s)|_{r=R}. \quad (4.25)$$

Finally, we get

$$c_2 = \frac{\kappa I_1 \left(\sqrt{\frac{s+\kappa}{\alpha}} R \right)}{s \sqrt{s+\kappa} \left(\sqrt{s} K_1 \left(\sqrt{\frac{s}{\alpha}} R \right) I_0 \left(\sqrt{\frac{s+\kappa}{\alpha}} R \right) + \sqrt{s+\kappa} K_0 \left(\sqrt{\frac{s}{\alpha}} R \right) I_1 \left(\sqrt{\frac{s+\kappa}{\alpha}} R \right) \right)}, \quad (4.26)$$

$$c_3 = - \frac{\kappa K_1 \left(\sqrt{\frac{s}{\alpha}} R \right)}{\sqrt{s}(s+\kappa) \left(\sqrt{s} K_1 \left(\sqrt{\frac{s}{\alpha}} R \right) I_0 \left(\sqrt{\frac{s+\kappa}{\alpha}} R \right) + \sqrt{s+\kappa} K_0 \left(\sqrt{\frac{s}{\alpha}} R \right) I_1 \left(\sqrt{\frac{s+\kappa}{\alpha}} R \right) \right)}. \quad (4.27)$$

4.B Effective diffusion model

We consider a model of classical diffusion in a planar space. We assume that the recorded luminosity is proportional to the density of photoactivated particles in the given region. Let us call $\varphi(t)$ the density of fluorescent particles in the source region S . Standard diffusion theory predicts that the density of photoactivated particles at any point in space is described by the following Cauchy problem:

$$\begin{cases} \partial_t u(\mathbf{x}, t) - D \nabla_{\mathbf{x}}^2 u(\mathbf{x}, t) = 0, & \mathbf{x} \notin S \\ u(\mathbf{x}, t) = \varphi(t), & \mathbf{x} \in \partial S \\ u(\mathbf{x}, t = 0) = 0, \end{cases} \quad (4.28)$$

where we denoted the diffusion coefficient by D , we imposed $\varphi(t)$ as Dirichlet boundary condition on the boundary of the source region S , and assumed a zero initial condition. By Duhamel's theorem, if v is the solution of the auxiliary problem with a time independent boundary condition

$$\begin{cases} \partial_t v(\mathbf{x}, t) - D \nabla_{\mathbf{x}}^2 v(\mathbf{x}, t) = 0, & \mathbf{x} \notin S \\ v(\mathbf{x}, t) = 1, & \mathbf{x} \in \partial S \\ v(\mathbf{x}, t = 0) = 0, \end{cases} \quad (4.29)$$

then we can build a solution for equation 4.28 as a convolution between the auxiliary solution $v(r, t)$ and the time-varying boundary term:

$$u(\mathbf{x}, t) = \int_0^t \varphi(t - \tau) \partial_{\tau} v(\mathbf{x}, \tau) d\tau. \quad (4.30)$$

We consider this problem in a 2D plane with radial symmetry where the source S is a disk of radius a and diffusion occurs in the space $r > a$. Expressing the Laplace operator in polar coordinates, the auxiliary Cauchy problem becomes:

$$\begin{cases} \partial_t v(r, t) - D \left(\partial_r^2 v(r, t) + \frac{1}{r} \partial_r v(r, t) \right) = 0, & r > a \\ v(r, t) = 1, & r = a \\ v(r, t = 0) = 0. \end{cases} \quad (4.31)$$

By taking the Laplace transform of the diffusion equation, we obtain a modified Bessel equation of order 0:

$$z^2 \partial_z^2 \hat{v}(r, s) + z \partial_z \hat{v}(r, s) - z^2 \hat{v}(r, s) = 0, \quad \text{with } z = \sqrt{s/D} r \quad (4.32)$$

where $\hat{v}(r, s) = \mathcal{L}\{v(r, t)\}$ indicates the Laplace transform of $v(r, t)$. The only non-diverging solution of the modified Bessel equation is

$$\hat{v}(r, s) = A(s) K_0(\sqrt{s/D} r) \quad (4.33)$$

where $K_0(z)$ is the modified Bessel function of the second kind (order 0) and $A(s)$ is a coefficient determined by imposing the boundary condition at $r = a$. The solution is thus

$$\hat{v}(r, s) = \frac{K_0(\sqrt{s/D}r)}{s K_0(\sqrt{s/D}a)}. \quad (4.34)$$

To find the solution of the time-varying problem, we need to apply Duhamel's theorem (eq. 4.30):

$$u(r, t) = \int_0^t \varphi(t - \tau) (\mathcal{L}^{-1} \{s \hat{v}(r, s)\}|_{\tau}) d\tau, \quad (4.35)$$

where we used the differentiation property $\mathcal{L} \{\partial_t f\} = s \mathcal{L} \{f\}$.

4.C Denoising of photoactivated channel

In the case of the photoactivated target channel, we are interested in cleaning the images while preserving the dynamical evolution in light intensity caused by photoactivation. Thus, we cannot apply the Noise2Noise model as presented for the morphology channel and we have to follow a more conservative approach. We thus employ the method described in [Foi 2009a] to denoise clipped noisy images, adapting its implementation for the specific case of videos. The denoising procedure can be described by the following steps:

1. Estimate the clipped Poisson-Gaussian model parameters [Foi et al. 2008];
2. Find an optimal a variance stabilizing transformation (vst) [Foi 2009b];
3. Perform the vst;
4. Denoise the resulting video with the V-BM3D algorithm [Dabov, Foi, and Egiazarian 2007; Ehret and Arias 2020];
5. Invert the VST and perform the declipping.

Variance stabilizing transformation. We use a variance stabilizing transformation to map the Poisson-Gaussian model to a unit variance Gaussian model. Based on the noise model parameters a, b we find an optimized variance stabilizing function following the recursive integral refinement method presented in [Foi 2009b]. We find the inverse transformation by linear interpolation. If $b \ll a$ (i.e. the Gaussian component and the clipping are negligible, with Poisson-only noise prevailing) we use the analytical approximation of the generalized Anscombe transform and its inverse [Makitalo and Foi 2011]. When performing the inverse transformation, we take into account the systematic bias introduced by taking the expectation on the variance-stabilized values, compensating it as described in [Foi 2009a].

Gaussian denoising. We use the V-BM3D algorithm [Dabov, Foi, and Egiazarian 2007] to denoise the variance stabilized video, using the implementation provided in [Ehret and Arias 2020].

PART II

Neural time series

CHAPTER 5

Denoising neural time series

The material in this chapter is the result of a recent collaboration with Vincent Villette from Stéphane Dieudonné’s group at École Normale Supérieure.¹ After an internal talk I gave at ENS, Vincent contacted me to see if the processing and denoising techniques I developed for the analysis of endoplasmic reticulum (ER) dynamics—which we have seen in chapter 4—could be applied to the optical recording data he was working on. The idea was that if we could improve the signal to noise ratio (SNR) by means of mathematical methods, we would have been able to obtain the same signal quality by recording less. In particular, the extra gain in SNR could be exploited to expand the acquisition capabilities to record simultaneously from multiple cells. What follows is an overview of this ongoing² project.

A broad variety of techniques have been proposed to study neural circuits in behaving animals, binding the genetic characterization of cell types and the properties of their electrical activities [Lecoq, Orlova, and Grewe 2019; Lipovsek et al. 2021; Tremblay, Lee, and Rudy 2016]. However, to be able to resolve the neural code, the area of investigation has to be framed at the good scale, i.e. the functional unit of the neural circuit. Among the plethora of recording techniques, covering all scales from the single cell to an entire population, the intermediate scale of neural microcircuits has emerged as one of the most suitable to understand activity patterns and neural code in the brain [Feldt, Bonifazi, and Cossart 2011]. Microcircuits represent specific connectivity patterns between neurons in a limited region, in general associated to a particular functional role. Thus, if single-cell recordings provide no information about interactions, and in population-level recordings individual activities are often inextricably mixed, the microcircuit scale provides a compromise between the two: the possibility of accessing single activities, and at the same time the presence of clear and structured patterns of interaction.

In the study of neural microcircuits, one can take the advantage of the fact that signals can be monitored directly and simultaneously by recording the neuronal membrane potential. The techniques that allow recording of membrane potential fluctuations and neuronal spiking in genetically identified cells are basically limited to three: visually guided patch clamp [Jouhanneau and Poulet 2019], optotag-

¹I really wish to thank Vincent Villette for his incredible dedication in the project and for sharing with me part of his extraordinary talent. I also want to thank Stéphane Dieudonné and Jonathan Bradley for their patient support and their masterful feedback, without which this work would not have been possible.

²As I am writing, Vincent is performing experiments on new mice to confirm the preliminary results outlined in this chapter.

ging coupled with juxtacellular recording and labelling [Ding et al. 2022], and optical voltage recordings [Bando et al. 2019]. The yield of the first method is already low in anaesthetized conditions (except for very specialized laboratories) and becomes almost null in awake, behaving mice, while the second does not allow simultaneous multicellular recording and is typically limited to one or two neurons per brain. For this reason, we focus here on the third technique.

The recent development of genetically encoded voltage indicator (GEVI), coupled with optical methods, allowed performing single-trial tracking of spikes in awake, behaving, head-fixed mice for several minutes (figure 5.1 A) [Villette et al. 2019], overcoming the methodological limitations discussed above. While single-photon based techniques have been recently improved [Xiao et al. 2021], the axial extent of the excitation light makes it impossible to prevent mixing of the signal of close by neurons. This is particularly problematic when recording activities in a densely labelled preparation (for example, the one shown in figure 5.1 B). This limitation can be avoided by multiphoton techniques, such as those used to monitor calcium dynamics with genetically encoded calcium indicator (GECI) in neuronal populations, given their capability to ensure source separation in both lateral and axial dimensions.

In contrast to GECIS, to provide a good signal to noise ratio (SNR), GEVIS need to be addressed at the membrane. This translates to a substantial reduction in the number of fluorophores contributing to the signal. In addition, the photonic property of the multiphoton light source (polarized laser [Bloxham et al. 2021]) reduces the number of excitable molecules in the approximately spherical shape of a neuronal cell body (figure 5.1 B). The technique thus requires very efficient methods to expand the excitation volume to contain a high density of GEVI molecules. This has been recently implemented in awake, head-fixed, behaving mice [Villette et al. 2019] thanks to the ultrafast local volume of excitation (ULOVE) method (figure 5.1 C), which allows recording of spiking and subthreshold membrane potential with high temporal resolution. The technique can be used to record activity from distinct volumes, possibly in different cells, at a relatively high sampling rate (figure 5.1 D). However, it is currently limited to simultaneous recording of two cells, or three at best, as the dwell-time cannot be reduced too much without also compromising the signal to noise ratio. For the study of microcircuits, it is thus essential to extend this technique to enable recording of multiple cells, all while maintaining a sufficiently high temporal resolution and a reasonable SNR.

To overcome this limit, two approaches may be considered. First, one may decide to refine the excitation pattern, concentrating the excitation in a smaller volume (figure 5.1 E) as to collect a higher photon flux. Second, one can post-process the signal to improve the SNR by mathematical means. The first method allows working with a more densely labelled cell population, which is a prerequisite to record several neurons at the same time, but also increases the sensitivity to artefacts induced by animal motion, which then need to be reduced by optimized surgeries and behavioural paradigms. The second method is the topic of this chapter.

The idea is that we can collect fewer volumes—ideally, just one per cell—and remove the photon-counting noise post-hoc via a denoising procedure, ensuring satisfactory SNR. Thus, instead of recording multiple volumes for each cell, we can distribute them across multiple cells, potentially multiplying the sampling capability.

Given the recentness of these optical developments, no specifically designed denoising method has been proposed in the literature. We note that the random access multiphoton microscopy (RAMP) ultrafast local volume of excitation (ULOVE) strategy is not equivalent to imaging, since it does not capture an image by reconstructing pixel by pixel the signal collected with a photomultiplier tube (PMT), but instead produces what we call an *optical recording*, i.e. a temporal trace where each data point is the sum of the signal resulting from a specific excitation pattern. While imaging by camera sensors or slow scanning multiphoton based imaging (using PMT) provide spatiotemporal information, ULOVE optical recording provides mainly temporal information and only partial spatial information (due to scanning along a spatial axis, figure 5.1 E).

However, similarly aimed denoising techniques have been employed for single-photon voltage imaging and calcium imaging, either to assist cellular segmentation or to improve the SNR. We can classify these techniques in two main families: methods based on matrix factorization, and deep learning methods. For the first family, non-negative matrix factorization (NMF) has long been used to perform automatic cell segmentation on calcium imaging recordings [Maruyama et al. 2014], improving on earlier PCA-ICA methods [Mukamel, Nimmerjahn, and Schnitzer 2009]. This technique consists in finding a low-rank matrix decomposition with non-negative components [Lee and Seung 1999]. Since then, several techniques have been developed to improve source separation, scalability, and speed [Giovannucci et al. 2019; Keemink et al. 2018; Pnevmatikakis 2019]. Voltage signals, in particular, are characterized by limited amplitude and SNR. For this reason, significant effort has been devoted in improving the factorization techniques and in making them efficient for single-photon voltage imaging data [Buchanan et al. 2018; Cai et al. 2021; Giovannucci et al. 2021; Oltmanns et al. 2020; Xie et al. 2021]. State of the art algorithms use penalized matrix decomposition (PMD) [Witten, Tibshirani, and Hastie 2009] in order to separate signals from uncorrelated noise. Such PMD-based techniques have been shown to improve SNR of functional imaging data by a factor 2 to 4 [Buchanan et al. 2018].

On the other hand, deep learning methods have also been applied to denoise calcium imaging data [Lecoq et al. 2021; Soltanian-Zadeh et al. 2019]. Interestingly, Lecoq et al. proposed to use deep learning techniques on several types of data, including recently developed multichannel electrophysiological probes (Neuropixels). The SNR characteristics and voltage fluctuations recorded by these devices are comparable to those obtained by RAMP-ULOVE optical recording. With a deep learning based denoising method, Lecoq et al. report a decrease of the root-mean-square error by a factor 1.7, allowing an increase of the number of neurons per probe by

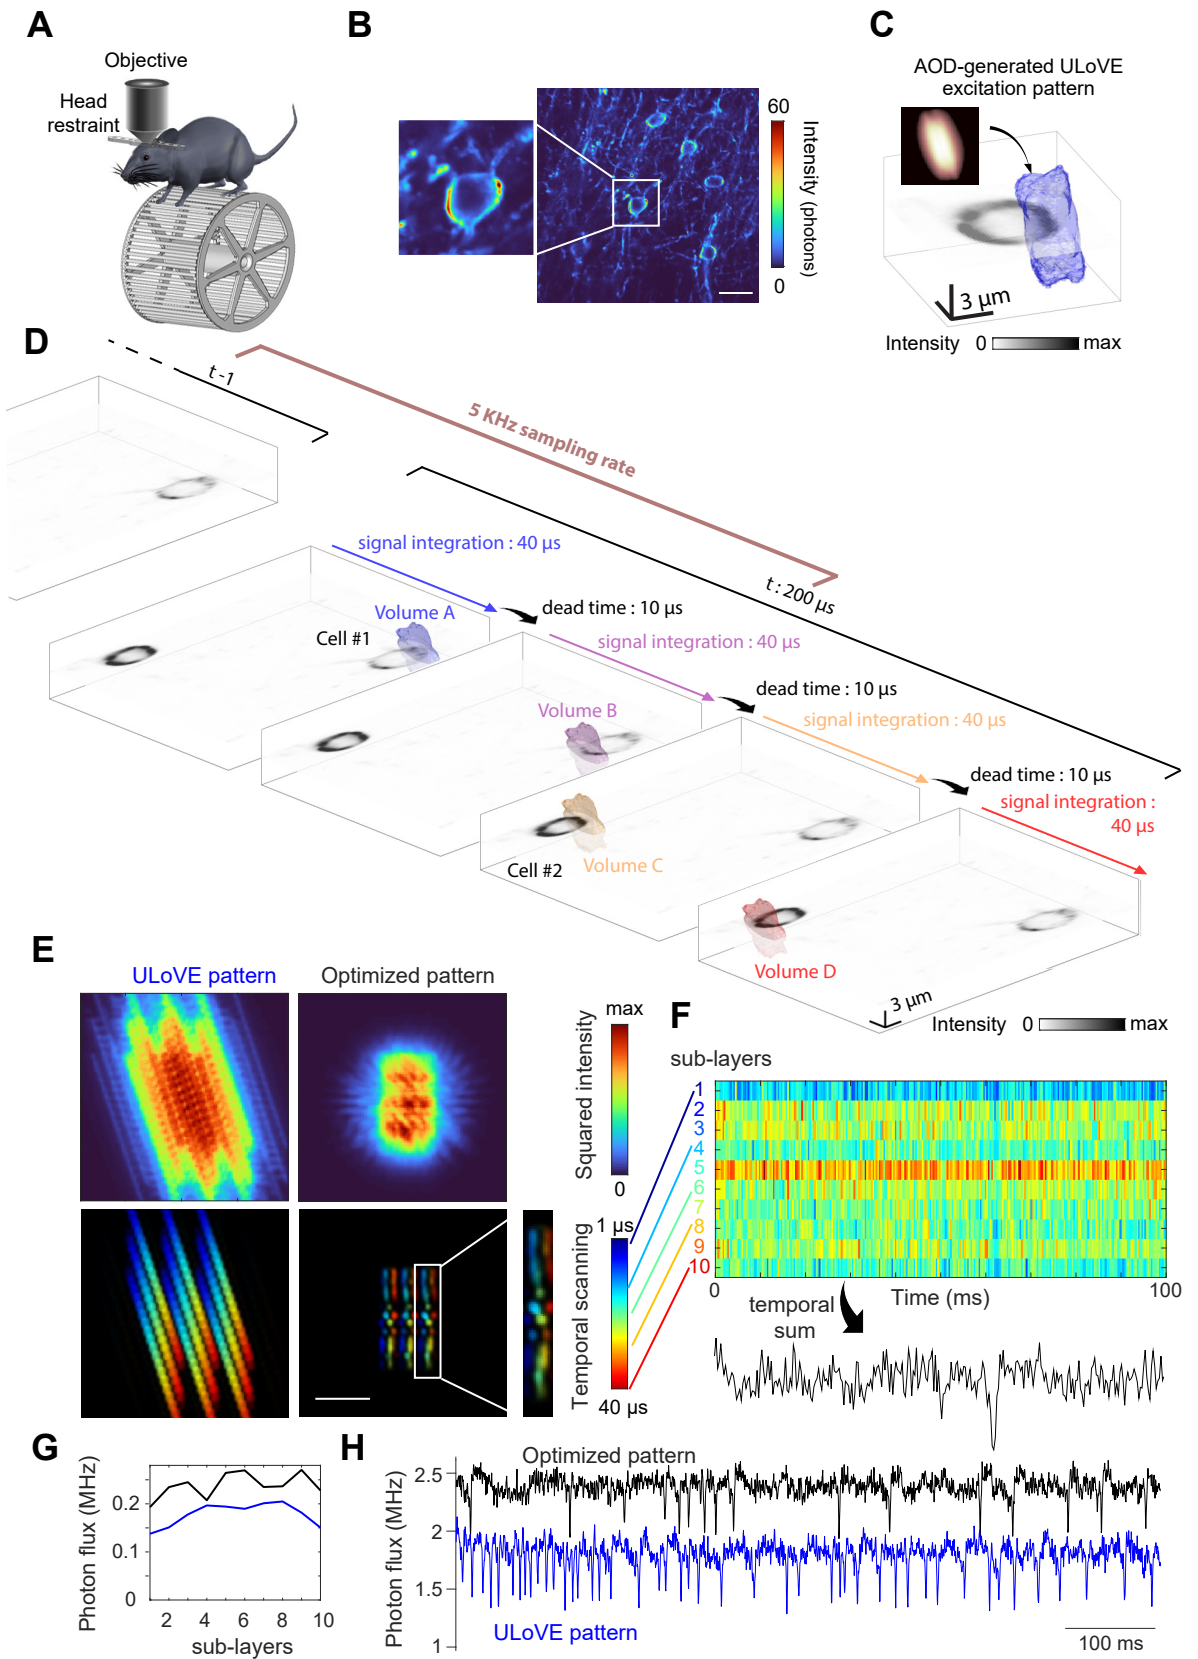


Figure 5.1 Experimental principles of optical voltage recordings on awake, head-fixed, behaving mice. **A** Scheme of the head-fixed mouse on a frictionless, non-motorized rotating wheel. **B** Cerebellar interneurons express JEDI-2P, a fast GEVI (scale bar denotes 20 μm). The inset shows a single molecular layer interneuron (MLI) soma, highlighting the effect of polarized excitation light on the membrane GEVI molecules. **C** A 3D representation of the ultrafast local volume, covering a part of the cell membrane. **D** Scheme of the acquisition protocol for simultaneous recording of two cells with two volumes each. The total sampling rate is 5 kHz. **E** Top projection of the squared intensity (top) and temporal microscanning at focal plane (bottom) for uLOVE [Villette et al. 2019] and optimized patterns (unpublished), obtained from simulations (work of Alberto Lombardini). The scale bar denotes 5 μm . The inset shows a single microscanned line of the optimized pattern. Note that temporal microscanning is visible but due to multiplexing (microscanning of 3 lines), spatiotemporal data are interleaved within each layer. **F** The 2D spatiotemporal data is turned into a single vector by summing over fragments. **G** Profiles over fragments indicate that photon flux is higher for the optimized (black) versus the uLOVE pattern (blue). **H** Traces represents the temporal sum of the optimized (black) versus the uLOVE (blue) pattern. They are obtained serially from the same molecular layer interneuron. *Courtesy of Vincent Villette. A, C, and D are copied and adapted from the JEDI-2P article by Liu, Lu, Gou, Villette, et al. (currently in revision).*

25.5% [Lecoq et al. 2021]. Those encouraging results lead us to consider spatiotemporal based denoising methods.

The uLOVE technique uses two principles: multiplexing and scanning. The scanning is a temporal shift across space. Thus, by subdividing one volume into temporal fragments of time while the excited light is steered through the cell membrane, one can effectively record spatiotemporal data (figure 5.1 E–F). Information about the spatiotemporal structure can then be used to separate the low-rank correlated signals from the temporally uncorrelated noise. As mentioned above, another possibility is to increase the photon flux by concentrating the local volume of excitation (optimized pattern, figure 5.1 E). This comes at the cost of destroying the spatial structure of the pattern. A comparison of the two patterns is presented in figure 5.1 G: while the uLOVE pattern shows a bell-shaped photon flux along the volume layers (blue curve), corresponding to the membrane structure, the optimized pattern is characterized by a more irregular flux (black). Interestingly, we thought we could exploit these methodological variations to compare the gain of a 2D spatiotemporal denoising versus a 1D, temporal-only denoising.

5.1 Overview of the denoising problem

Before delving in the technicalities of the denoising method, we present a theoretical overview of the denoising problem applied to optical recording data. In the following, we introduce a model for the data and the noise, and briefly examine the properties of the statistical estimators that could allow us to recover the signal.

5.1.1 Data model

Optical recordings of neuronal activity with GEVI are acquired by sequentially scanning portions of the neuron soma with special excitation patterns (figure 5.1 C–D). As explained above, the scanning consists in shifting the excitation pattern along

a spatial axis during acquisition (figure 5.1 E). By subdividing the scanning process in small time intervals, one thus achieve subvolume resolution as the volume is acquired in the form of separate spatial layers corresponding to a partial sweep along the scanning axis (figure 5.1 F). Given that the time required to scan a full volume is relatively short ($40 \mu\text{s}$ for a 10-layer volume, as the one shown in figure 5.1 F), we consider the acquisition of the subvolume layers to be instantaneous. We thus describe the recording of a volume of M layers by the M -channel time series $\mathbf{x}(t)$ defined as

$$\mathbf{x}(t) = (x_1(t), x_2(t), \dots, x_M(t))^T \quad (5.1)$$

where the channels $x_1(t), x_2(t), \dots, x_M(t)$ correspond to the light intensity recorded in each of the M layers. In practice, we only have discrete time samples, and we can represent the acquired data by the data matrix

$$\mathbf{X} = \begin{bmatrix} x_{11} & x_{12} & \cdots & x_{1N} \\ x_{21} & x_{22} & \cdots & x_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ x_{M1} & x_{M2} & \cdots & x_{MN} \end{bmatrix} \in \mathbb{N}^{M \times N} \quad (5.2)$$

with elements $x_{ij} = x_i(t_j)$ representing photon counts, i.e. a matrix where each row represents a channel of the time series with N samples, and each column is a photograph of the M layers composing a volume. Note that, in practice, N is much larger than M .

We consider the recorded signal \mathbf{X} to be the sum of the membrane potential $y(t)$ and some random noise $\zeta(t)$ contaminating the signal³. We can model each channel as

$$x_m(t) = \underbrace{k_m y(t)}_{\text{membrane potential}} + \underbrace{\zeta_m(t)}_{\text{noise}}, \quad (5.3)$$

where the factors $k_m \geq 0$ describe the spatial distribution of the true membrane potential across the layers, which can depend on the amount of membrane or concentration of voltage-sensitive indicator (GEVI). In matrix form,

$$\mathbf{X} = \mathbf{Y} + \mathbf{Z} \quad (5.4)$$

where $\mathbf{Y} = \mathbf{k} \otimes \mathbf{y}$ is a rank-1 matrix representing the neural signal and $\mathbf{Z} = [\zeta_{ij}]_{ij}$ is the noise component. Note that in general the spatial distribution described by \mathbf{k} should also be time-dependent, but here we consider that the time variation is negligible. This assumption is valid because variation of the spatial evolution is mostly caused by movement artefacts, whose timescale is much longer ($\sim 40\text{--}200$ ms) than that of the neural patterns we are interested in studying (~ 5 ms). Thus, we can always analyse the data by splitting it in relatively short windows such that the spatial variation is negligible. In the following, we will briefly study traditional statistical approaches to estimate $y(t)$ depending on the type of noise, beyond simple averaging of the channels.

³In practice, the data is contaminated by other sources of fluorescence such as the signal from neighbouring neurons or neuropil. Our aim here is, however, to improve the SNR to produce a high quality recording. We thus do not consider the unmixing problem for now, given that it could be more easily solved once the noise is suppressed. Empirical evidence seems to support the idea that treating denoising and unmixing problems separately can prove more effective than trying to solve the two problems at once [Buchanan et al. 2018].

5.1.2 Noise model and best estimators

We consider here a general definition for the noise term $\zeta(t)$ of equation 5.3 which allows approximating real noise sources in a flexible manner:

$$\zeta(t) = \sigma(y(t)) \xi(t), \quad \text{with } \xi(t) \sim \mathcal{N}(0, 1) \quad (5.5)$$

where σ is a positive real function of the signal y . This general definition allows modelling noise sources with signal-dependent variance (heteroskedastic) or constant variance (homoskedastic), depending on the chosen form of $\sigma(y)$. As we have seen in chapter 4, a good choice for $\sigma(y)$ when working with optical sensors is the following [Foi et al. 2008]:

$$\sigma^2(y) = ay + b. \quad (5.6)$$

This form allows us to model a mixture of signal-independent white noise with variance b , modelling instrumental noise such as the sensor dark current, and a signal-dependent Poisson-like⁴ noise, which models the photon-counting noise.

Estimators for homoskedastic noise

We start by considering the simplest case, that of homoskedastic noise with constant, signal-independent variance $\sigma^2 = b$ (with $a = 0$). Let us start by assuming that the coefficients $\mathbf{k} = (k_1, k_2, \dots, k_M)$ are known. In principle, we can get an estimate of \mathbf{y} simply by averaging over the channels $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_M$ (up to normalization), such as:

$$\hat{\mathbf{y}}_{\text{naive}}(t) = \frac{1}{\sum_m k_m} \sum_m \mathbf{x}_m(t) \quad (5.7)$$

But can we define a better estimator? Admittedly, we should take advantage of the dependency between the channels that results from relation 5.3 to improve the estimation of $y(t)$. One way to proceed is maximum likelihood estimation, i.e. to look for the estimator $\tilde{\mathbf{y}}$ that maximizes the likelihood of the observation \mathbf{X} . The intuitive result of such procedure is

$$\hat{\mathbf{y}}_{\text{ML}} = \frac{\mathbf{X}^T \mathbf{k}}{\|\mathbf{k}\|_2^2}, \quad (5.8)$$

i.e. we obtain $\hat{\mathbf{y}}_{\text{ML}}$ by projecting the data on the axis in the direction of \mathbf{k} . In practice though, the coefficients \mathbf{k} are unknown, and thus the projection axis is undetermined. Nevertheless, one can use the same maximum likelihood estimation for an estimator $\hat{\boldsymbol{\theta}}_{\text{ML}} = (\hat{\mathbf{k}}_{\text{ML}}, \hat{\mathbf{y}}_{\text{ML}})$ to simultaneously determine the two parameter vectors. In appendix 5.A we give an explicit derivation of this estimator, showing that it can be obtained by singular value decomposition (SVD) of the data matrix \mathbf{X} . More precisely, given the SVD factorization

$$\mathbf{X} = \sum_{m=1}^M \sigma^{(m)} \mathbf{u}^{(m)} \otimes \mathbf{v}^{(m)}, \quad (5.9)$$

⁴Assuming the number of photons is high enough that one can reliably approximate the Poisson process with a normal distribution as we have done in equation 5.5.

the maximum likelihood estimates $\hat{\mathbf{k}}_{\text{ML}}$ and $\hat{\mathbf{y}}_{\text{ML}}$ correspond respectively to the left and right singular vectors $\mathbf{u}^{(1)}$ and $\mathbf{v}^{(1)}$ of the data matrix \mathbf{X} associated to the largest singular value $\sigma^{(1)}$ (up to normalization). This problem is in fact akin to a low-rank approximation problem of the matrix \mathbf{X} , specifically we look for the rank-1 matrix $\hat{\mathbf{Y}}$ minimizing the Frobenius norm $\|\hat{\mathbf{Y}} - \mathbf{X}\|_{\text{F}}$.

Estimators for heteroskedastic noise

We now consider the case of $a > 0$, i.e. a signal-dependent noise variance. Again, if the coefficients \mathbf{k} are known, the inter-channel correlation can be exploited to find an estimate of \mathbf{y} . We proceed similarly to the homoskedastic case to obtain the maximum likelihood estimator $\hat{y}_{\text{ML}}(t)$. In the case $a > 0, b = 0$ we can obtain the explicit form:

$$\hat{y}_{\text{ML}}(t_n) = a \frac{M}{2 \sum_m k_m} \left(\sqrt{1 + \frac{4}{a^2 M^2} \left(\sum_m k_m \right) \left(\sum_m \frac{x_{m,n}^2}{k_m} \right)} - 1 \right). \quad (5.10)$$

While the result is not as clean and intuitive as in the homoskedastic case, it can still be seen as a (non-linear) projection on the axis along \mathbf{k} . Sadly, the general solution for $a > 0, b > 0$ cannot be easily expressed in closed form.

As to the case where the coefficients \mathbf{k} are also undetermined, the problem becomes quickly intractable due to interdependence between \mathbf{k} and \mathbf{y} , so that little progress can be made in the calculations. To my knowledge, solution of such a problem has to be obtained through numerical optimization, for example via the expectation–maximization scheme [Leeb 2021]. We should note, however, that optimization algorithms may be tricky to employ in practice as they provide no guarantee of convergence towards the absolute maximum. One risks falling into a local extremum, which can be represented for example by the naive average estimator defined in equation 5.7. It may thus be tricky to correctly initialize the algorithm so that it converges to a non-trivial solution.

Optimal estimators

One may wonder whether we can do any better than the maximum-likelihood estimators presented above. In other words, based on the simple model defined in equations 5.3 and 5.4, what is the maximum amount of information we can exploit to reconstruct the signal \mathbf{y} ?

A well-known result in mathematical statistics is the Cramér–Rao bound, which allows defining a lower bound to the variance of an estimator. In particular, the Cramér–Rao bound of an unbiased estimator $\boldsymbol{\theta}$ is given by:

$$\text{Cov}[\boldsymbol{\theta}] \geq \mathbf{J}^{-1}(\boldsymbol{\theta}), \quad (5.11)$$

where $\mathbf{J}^{-1}(\boldsymbol{\theta})$ is the inverse of the Fisher information matrix defined as

$$[\mathbf{J}(\boldsymbol{\theta})]_{i,j} = \mathbb{E} \left[\left(\frac{\partial}{\partial \theta_i} \log f(\mathbf{X} | \boldsymbol{\theta}) \right) \left(\frac{\partial}{\partial \theta_j} \log f(\mathbf{X} | \boldsymbol{\theta}) \right) \middle| \boldsymbol{\theta} \right], \quad (5.12)$$

with $f(\mathbf{X} | \theta)$ representing the likelihood of the observed data \mathbf{X} .

In our case, considering the estimator \hat{y} with known coefficients \mathbf{k} and $\sigma^2(y) = ay + b$, the data likelihood is given by

$$f(\mathbf{X} | \theta) = \prod_{m=1}^M \prod_{n=1}^N \frac{1}{\sqrt{2\pi\sigma^2(y_n)}} \exp\left(-\frac{(x_{m,n} - k_m y_n)^2}{2\sigma^2(y_n)}\right), \quad (5.13)$$

and the resulting Fisher information is an $N \times N$ matrix with elements $J_{i,j}$ defined as

$$J_{i,j} = \sum_{m=1}^M \frac{k_m^2 (a^2 + 2ak_m y_j + 2b)}{2(ak_m y_j + b)^2} \delta_{i,j}. \quad (5.14)$$

We can express the Cramer–Rao bound for the special cases described above.

Let us start from the homoskedastic case ($a = 0$). From equations 5.11 and 5.14 we derive the lower bound

$$\text{Var} [\hat{y}(t)] \geq \frac{b}{\|\mathbf{k}\|_2^2}. \quad (5.15)$$

It is straightforward to verify that the variance of the projection estimator of equation 5.8 is indeed equal to this lower bound, meaning that the maximum likelihood estimator is—in this case—the minimum-variance unbiased estimator. To better understand how the projection estimator achieves optimality, we can compare it to the variance of the naive estimator (equation 5.7):

$$\text{Var} [\hat{y}_{\text{naive}}(t)] = \frac{Mb}{\left(\sum_{m=1}^M k_m\right)^2}. \quad (5.16)$$

Let us assume that the average of the coefficients k_m is 1, i.e. $\sum_{m=1}^M k_m = M$. In this case, $\text{Var}[\hat{y}_{\text{naive}}(t)] = b/M$, which is the classical result for the averaging of M independent and normally distributed random variables with variance b . Note that this result does not depend on how the total value M is partitioned among the k_m terms. Instead, the variance of the projection estimator \hat{y}_{ML} depends on the distribution of the k_m . The worst case corresponds to the smallest value of $\|\mathbf{k}\|_2^2$: this is obtained when the signal is uniformly distributed across the channels, i.e. $k_m = 1/M$. Vice versa, the best case is obtained when the signal is fully concentrated in one channel, e.g. $k_m = M \delta_{0,m}$. The corresponding variances for the maximum-likelihood estimator are

$$\text{Var}[\hat{y}_{\text{ML}}^{(\text{worst})}(t)] = b/M, \quad \text{for } k_m = \frac{1}{M}; \quad (5.17)$$

$$\text{Var}[\hat{y}_{\text{ML}}^{(\text{best})}(t)] = b/M^2, \quad \text{for } k_m = M \delta_{0,m}. \quad (5.18)$$

This makes the difference between the naive and optimal estimator more intuitive. Given that the noise variance is constant for all $x_m(t)$, the signal to noise ratio depends only on k_m and is proportional to $k_m^2 b$. While the naive estimator averages over all the channels independently of the SNR, the projection estimator exploits

the difference in SNR by giving more weight to the channels with higher SNR, thus producing a better estimation of the signal $y(t)$. In the limit for which just one k_m is different from zero, information about $y(t)$ is concentrated in one channel, while the remaining $M - 1$ channels only contain noise. In this case, summing all the channels simply means adding more noise to the estimate. The projection operator, instead, will only pick the channel containing the actual signal and will discard the others.

This is a demonstration of why taking averages over samples with different SNR is not a good idea. This is particularly relevant for the case of optical recording we are treating here, since the fluorescent proteins are concentrated on the cell membrane, representing a fraction of the recorded volume. In this case, giving appropriate weight to the different layers can improve the estimation of the membrane potential. We can make a back of the envelope calculation to understand what gain can be obtained by this approach in a practical situation. In the optical recordings described above, we normally have $M \approx 10$ layers per volume (see figure 5.1 F). In the optimal case of equation 5.18, this would translate to the reduction of the noise variance by a factor 10 when compared to the naive average estimator. This is however an ideal limit. As more realistic case, we can consider the signal to be uniformly spread among half of the layers, with the remaining ones only containing noise. This situation would result in $\|\mathbf{k}\|_2^2 = 2M$, corresponding to a reduction of the variance by a factor 2 with respect to averaging.

We now turn to the heteroskedastic noise case with $a > 0, b = 0$. Simplifying equation 5.14, we obtain the Cramér–Rao bound:

$$\text{Var} [\hat{y}(t)] \geq \frac{a y(t)}{\|\mathbf{k}\|_1} \frac{1}{1 + \frac{a M}{\|\mathbf{k}\|_1} \frac{1}{2y(t)}}, \quad (5.19)$$

where $\|\mathbf{k}\|_1 = \sum_{m=1}^M |k_m|$. We should note, however, that the noise model in equation 5.6 can only be used to model photon-counting noise when $y(t)$ is sufficiently large so that the Poisson distribution can be approximated by a normal distribution (indicatively $y(t) \gtrsim 30$). We can thus see that for $a \approx 1$ this effectively converges to the Cramér–Rao bound of the Poisson distribution, i.e.

$$\text{Var} [\hat{y}(t)] \geq \frac{y(t)}{\|\mathbf{k}\|_1}. \quad (5.20)$$

In this case, the average estimator (equation 5.7) is the minimum-variance unbiased estimator. In fact, the noise variance is proportional to the signal and the proportional weighting, which we have seen to be fruitful in the homoskedastic, does not bring any advantage here. While biased estimators such as the one defined in equation 5.10 may provide some improvement over the bound, in practice the advantage is questionable as the bias will also contribute to the mean-squared error (MSE) of the estimation. Unfortunately this means that little progress in denoising can be made in the heteroskedastic case under the simple model of equation 5.4, unless

more assumptions regarding the structure of the data are made. Note that in what we discussed so far, we made no assumptions regarding the characteristics of $y(t)$.

5.1.3 Current approaches

In the previous section we have shown how, for homoskedastic noise, the best estimate of $y(t)$ is obtained by weighting accordingly the different channels based on their SNR either by projection (when the proportionality \mathbf{k} is known) or SVD (when proportionality must also be estimated). We have also seen how this approach does not work well in the case of heteroskedastic noise. Most of the current approaches [Buchanan et al. 2018; Giovannucci et al. 2019; Pnevmatikakis et al. 2016] tackle this problem by implicitly considering that the noise variance is almost constant in time. This is appropriate if the signal $y(t)$ makes only relatively small oscillations around a baseline, so that the noise variance can be approximated by $\sigma^2(y(t)) \approx a\bar{y} + b$ where \bar{y} is the signal baseline. As we mentioned before, this assumption is particularly relevant if the signal is analysed on relatively short time windows. In this case, one can rescale the signals to make the noise variances equal across all channels, effectively tracing back the problem to the homoskedastic model we described above. The estimate of the signal $y(t)$ is then obtained via some form of matrix factorization, similarly to what we have described in section 5.1.2 (see equation 5.9). Since simple SVD truncation can only provide limited reduction of the mean-squared error (MSE) [Vu, Chunikhina, and Raich 2021], in particular when the noise level is high, state of the art methods exploit additional constraints to improve the estimation. Examples of such constrained methods are non-negative matrix factorization (which constraints matrix elements to be non-negative) [Maruyama et al. 2014] or the more general penalized matrix decomposition [Witten, Tibshirani, and Hastie 2009]. The latter approach is typically represented by an optimization problem of the form [Buchanan et al. 2018; Witten, Tibshirani, and Hastie 2009]

$$\arg \min_{\mathbf{u}^{(m)}, \mathbf{v}^{(m)}} \left\| \mathbf{X} - \sum_{m=1}^M \sigma^{(m)} \mathbf{u}^{(m)} \otimes \mathbf{v}^{(m)} \right\|_{\text{F}}^2, \text{ subject to } g_1(\mathbf{u}^{(m)}) \leq c_1, g_2(\mathbf{v}^{(m)}) \leq c_2, \quad (5.21)$$

where g_1, g_2 represent penalties on the spatial and temporal components respectively. These penalty terms are used to impose regularity constraints such as sparsity (using ℓ^1 norm), smoothness (e.g. using derivatives), or total variation. Total variation and trend filtering penalties have been reported to be particularly effective in both calcium and voltage imaging [Buchanan et al. 2018].

5.2 Methods

In this section, we present a new approach to the denoising of optical voltage recordings, taking a different route with respect to the methods discussed above. Our approach is composed of two parts. Firstly, we estimate the parameters of the noise,

similarly to what we have done in chapter 4, assuming both constant and signal-dependent contributions. We will show in section 5.3 that the signal-dependent component is indeed dominant in our data. Then, instead of normalizing the signals assuming homoskedasticity of the noise, we perform variance stabilization using the generalized Anscombe transform [Makitalo and Foi 2011]. Secondly, inspired by image denoising techniques, we expand the block-matching and 3D filtering (BM3D) algorithm [Dabov et al. 2007] to work with both multi- and single-channel time series. This approach relies on finding similar patterns in the signal (block-matching) which are then grouped together and transformed to obtain a sparse representation. The sparsified signal is then filtered by wavelet thresholding to eliminate the noise and transformed back in the original domain. Instead of imposing directly regularity constraints as in penalized matrix factorization, this method exploits the presence of recurrent patterns in the signal such as spikes or up and down states. Regularity is still part of the process, but in an indirect way, depending on the choice of the sparsifying transform.

5.2.1 Estimation of the noise parameters and variance stabilization

We consider the noise model of equation 4.2. To estimate the parameters a, b , we proceed as discussed in [Foi et al. 2008], considering the data matrix \mathbf{X} as an image where each pixel represents a photon count at given time and volume layer. First we decompose the data by discrete wavelet transform (1-level) obtaining approximation coefficients \mathbf{W}_{app} and detail coefficients \mathbf{W}_{det} . Because of its irregularity, most of the noise will be concentrated in the detail coefficients. We define bins (*level sets*) by further smoothing \mathbf{W}_{app} using a Gaussian kernel ($\sigma = 2$ px) and then rounding each element to the nearest integer:

$$\mathbf{B} = \lfloor \mathbf{G} \otimes \mathbf{W}_{\text{app}} \rfloor \quad (5.22)$$

where \mathbf{G} is the Gaussian kernel, \otimes indicates the convolution operation, and $\lfloor \cdot \rfloor$ denotes elementwise integer rounding. We define a level set \mathcal{S}_k as the group of pixels with the same bin number k in \mathbf{B} , i.e. $\mathcal{S}_k = \{(i, j) \mid B_{i,j} = k\}$. Then, for each level set we estimate the mean by averaging over the elements of \mathbf{W}_{app} and the (unbiased) standard deviation over elements \mathbf{W}_{det} [Foi et al. 2008]:

$$\hat{y}_k = \frac{1}{N_k} \sum_{i,j \in \mathcal{S}_k} [\mathbf{W}_{\text{app}}]_{i,j}, \quad (5.23)$$

$$\hat{\sigma}_k = \frac{1}{\kappa_{N_k}} \sqrt{\text{Var}_{i,j \in \mathcal{S}_k} [\mathbf{W}_{\text{det}}]_{i,j}} \quad \text{with } \kappa_n = \sqrt{\frac{2}{n-1} \frac{\Gamma(\frac{n}{2})}{\Gamma(\frac{n-1}{2})}}, \quad (5.24)$$

where N_k is the number of elements in \mathcal{S}_k . We can then use the points $(\hat{y}_k, \hat{\sigma}_k)$ to estimate the noise parameters a, b based on the relation $\sigma = ay + b$ (figure 5.2 A). We do this by maximum-likelihood (see [Foi et al. 2008, section III.D] for details), solving the optimization problem numerically with the L-BFGS-B algorithm [Zhu et al. 1997].




Figure 5.2 Noise modelling and parameters estimation. **A** The signal-dependent noise relation $\sigma^2(y) = ay + b$ is estimated according to the procedure described in section 5.2.1. Dashed line represents the maximum likelihood estimate. **B** Distribution of the signal-dependent noise coefficient a for a set of test cells. **C** Distribution of the signal-independent noise variance b .

Once the noise parameters have been estimated, we perform variance stabilization via the generalized Anscombe transformations $\mathcal{A}_{a,b}(x)$ [Makitalo and Foi 2011]:

$$\mathcal{A}_{a,b} = \frac{2}{a} \sqrt{ax + \frac{3}{8}a^2 + b}. \quad (5.25)$$

We transform the data $x_{i,j}$ into the variance-stabilized $\tilde{x}_{i,j} = \mathcal{A}_{a,b}(x_{i,j})$ with unitary variance. After denoising, we will transform the signal back to the original domain by the closed-form approximation of the unbiased inverse of the Anscombe transformation [Makitalo and Foi 2013]:

$$\mathcal{A}_{a,b}^{-1}(x) = a \left(\frac{1}{4}x^2 + \frac{1}{4}\sqrt{\frac{3}{2}}\frac{1}{x} - \frac{11}{8}\frac{1}{x^2} + \frac{5}{8}\sqrt{\frac{3}{2}}\frac{1}{x^3} - \frac{1}{8} - \frac{b}{a^2} \right) \quad (5.26)$$

5.2.2 Block matching and filtering for neural time series

Block-matching and 3D filtering (BM3D) [Dabov et al. 2007] is considered a state of the art method in image denoising, but has also been generalized as a flexible tool in image deblurring and deconvolution [Danielyan, Katkovnik, and Egiazarian 2012]. In its essence, the block-matching and 3D filtering (BM3D) approach is analogous to non-local means algorithms. It works by grouping together similar patches (block-matching) and then combining them to reconstruct a clean version of the image. The difference with respect to non-local mean algorithms is the use of a 3D transform to obtain sparse representations of the groups of patches. Instead of simple averaging over similar patches, noisy components are removed via wavelet shrinkage and the patches are then recombined by weighted averaging based on their sparsity (collaborative filtering). In a second step, the estimation is refined by applying a wavelet-domain Wiener filter whose coefficients are calculated using the estimate obtained from the first wavelet-shrinkage step.

While the original BM3D algorithm was developed to denoise images taking advantage of spatial correlation, we adapt here the same methodology for time series. There is one important conceptual difference in the application of method. The reason of BM3D effectiveness has been attributed to the fact that the collaborative filtering step assigns more weight to uniform patches (because they are more sparse) and tends to neglect the patches containing edges [Lebrun 2012]. This is very relevant in the context of natural images, which are generally well described by a union of approximately uniform regions separated by clear edges. Our results show that the effectiveness of the method extends well beyond this model, as far as the chosen transform produces sparse patches. In our case, the signal can be represented as slowly oscillating baseline combined with sequences of spikes (see for example figure 5.1 H). In this context, patches containing a spike discontinuity are not necessarily weighted less since they can be represented in a sparse way if the appropriate basis is considered.

Based on these intuitions, we define the method as follows. We consider the signal from a single volume $\tilde{\mathbf{X}}$ after variance stabilization (equation 5.25) where rows represent layers and columns time samples (as described in section 5.1.1). While in general we treat the data matrix as a 2D image, we treat the two dimensions differently. Also, we consider that in general the data may contain only a single channel, i.e. $\mathbf{X} \in \mathbb{N}^{1 \times N}$.

We introduce differences with respect to the classical BM3D algorithm in the sparsifying transform and the block matching procedures. We then proceed as in the original BM3D to estimate the denoised signal in two steps: first we build an estimate by collaborative filtering with wavelet shrinkage, then we refine the result by an empirical wavelet-domain Wiener filter.

Sparsifying transform. We expect the spatial and temporal dimension to have different smoothness and sparsity properties, we define the initial 2D transform as a composition of separable 1D transforms over the spatial and time dimension. We denote this transform by the operator $T_{2D} = T_s T_t$, where T_s and T_t are respectively the spatial and temporal 1D transforms. Since we expect the spatial dimension to be fairly smooth, we use discrete cosine transform (DCT) along the spatial axis. Along the temporal axis, to obtain a sparse representation of the spikes, we use a discrete wavelet transform (DWT) with small number of vanishing moments (see section 5.2.3 for details regarding the choice of the transform). The final sparse representation is obtained by transforming a group of similar blocks along the grouping axis, which we define by the operator T_g . Similarly to BM3D we use a Haar transform along the grouping axis.

Block matching. We consider rectangular patches, with different dimension in the spatial and temporal axes chosen as described in section 5.2.3. We define the blocks $b_k \in \mathbb{R}^{N_s \times N_t}$ as patches of the original data $\tilde{\mathbf{X}}$, where N_s and N_t indicate the block size in the spatial and temporal axis respectively. Note that the blocks can be

overlapping. By $\mathcal{B}_X = \{b_1, b_2, \dots\}$ we denote the set of all distinct blocks that can be extracted from the data \mathbf{X} . (For sake of simplicity, in the following we drop the notation $\tilde{\mathbf{X}}$ and use \mathbf{X} to indicate the variance-stabilized data.)

To define groups of blocks, we first define the distance between two blocks as in the original BM3D algorithm [Dabov et al. 2007]:

$$d(b_i, b_j) = \frac{\|\text{ht}_{\lambda_1}(\mathbf{T}_{2D}b_i) - \text{ht}_{\lambda_1}(\mathbf{T}_{2D}b_j)\|_2^2}{N_s N_t}, \quad (5.27)$$

where $\text{ht}_{\lambda}(\cdot)$ denotes a hard-thresholding operation with threshold λ (i.e. the operation which sets to zero all elements smaller than the threshold). This hard-thresholding operation is needed to ensure good results even when the noise level is high. Groups of blocks are created by first considering a reference block b_r and then finding all blocks b_k within distance $d(b_r, b_k) < \tau$. We add here an additional constraint with respect of traditional block-matching by considering only those blocks which are local minima of the distance function (along space and time dimensions). This step helps to better preserve the shape of spikes in our data and minimize distortions, as we do not want to group together blocks which are close in terms of the distance d but with a shift in the location of the spike peak. We thus define the group of blocks based on the reference b_r as:

$$\mathcal{G}_r = \left\{ b_k \in \mathcal{B}_X \mid d(b_r, b_k) < \tau \wedge d(b_r, b_k) = \min_{b_i \in \mathcal{N}(b_k)} d(b_r, b_i) \right\}, \quad (5.28)$$

where $\mathcal{N}(b_k)$ indicates the set of neighbouring blocks with respect to b_k . In practice, we keep just the N_g closest blocks.

Collaborative and Wiener filtering. We perform the filtering procedure as in the original BM3D. For completeness, we briefly summarize the main steps. Given a group of similar blocks \mathcal{G} , we define the tensor \mathbf{G} formed by stacking the blocks $b \in \mathcal{G}$ along a new dimension, such that $\mathbf{G} \in \mathbb{R}^{N_s \times N_t \times N_g}$. Noise suppression is performed by applying the sparsifying transform to the group, followed by hard thresholding:

$$\tilde{\mathbf{G}}^{(\text{HT})} = \text{ht}_{\lambda_2}(\mathbf{T}_g \mathbf{T}_{2D} \mathbf{G}), \quad (5.29)$$

where the tilde indicates that the tensor is expressed in the transformed domain. Border effects can be reduced by applying a window (e.g. a Kaiser window) to the blocks before the 2D transform. After thresholding, the transform is inverted to obtain the denoised blocks

$$\mathbf{G}^{(\text{HT})} = \mathbf{T}_g^{-1} \mathbf{T}_{2D}^{-1} \tilde{\mathbf{G}}^{(\text{HT})}, \quad (5.30)$$

each block is then mapped back to its original location to build a denoised estimate of the signal $\hat{\mathbf{Y}}^{(1)}$. This procedure is applied repeatedly for groups created starting from different reference blocks, chosen sequentially to cover the full signal. Since the blocks are overlapping, multiple estimates are combined by weighted averaging based on the sparsity of the transformed group, with weights defined as

$w_k = 1/\max(|\tilde{\mathbf{G}}_k^{(\text{HT})}|_0, 1)$, where $|\tilde{\mathbf{G}}_k^{(\text{HT})}|_0$ indicates the number of non-zero elements in the transformed group (sometimes called ℓ^0 norm). We note that if a window function was applied before the transform, the same window has to be applied the weights. The same procedure is applied for a second time, substituting hard thresholding with Wiener filtering. Optionally, different transformations may be employed for the second step. The block matching is performed on the initial estimation $\hat{\mathbf{Y}}^{(1)}$. Each group is then filtered using the empirical Wiener coefficients calculated on the initial estimation:

$$\mathbf{W} = \frac{\|\mathbf{T}_g \mathbf{T}_{2D} \mathbf{G}^{(1)}\|_2^2}{\|\mathbf{T}_g \mathbf{T}_{2D} \mathbf{G}^{(1)}\|_2^2 + 1}, \quad (5.31)$$

$$\tilde{\mathbf{G}}^{(\text{Wi})} = \mathbf{W} \odot (\mathbf{T}_g \mathbf{T}_{2D} \mathbf{G}), \quad (5.32)$$

$$\mathbf{G}^{(\text{Wi})} = \mathbf{T}_g^{-1} \mathbf{T}_{2D}^{-1} \tilde{\mathbf{G}}^{(\text{Wi})}, \quad (5.33)$$

where \mathbf{W} are the empirical Wiener filter weights, $\mathbf{G}^{(1)}$ is the group of blocks expressed on $\hat{\mathbf{Y}}^{(1)}$, \mathbf{G} is the group expressed on the raw data \mathbf{X} , and \odot indicates elementwise multiplication (Hadamard product). Similarly to the first step, the blocks are then aggregated by weighted averaging with weights $w = 1/\|\mathbf{W}\|_2^2$ determined for each group, leading to the final estimate $\hat{\mathbf{Y}}$, which is then transformed via the inverse Anscombe transform (equation 5.26) to recover the signal in the original signal space.

5.2.3 Choice of parameters

Block size. We determine the size of patches by considering the scales associated with the spatial and temporal dimension separately. For the temporal scale, the detailed features are represented by spikes with a timescale of ~ 3 ms (FWHM < 1 ms). We choose a temporal block size of 10 ms to be able to enclose single spikes. For the spatial scale, we expect the layers to be relatively stable, only allowing small oscillations to movement of the animal. We thus consider a block size of 8 out of 10 layers to compensate for these small oscillations. When working with single-channel data, we simply skip the spatial transform by considering blocks with size 1 in the spatial axis (i.e. $T_s = 1$).

Wavelets. The central point of our approach is the denoising and preservation of spike patterns in the signal. On the one hand, this is motivated by the biological significance of the spikes. On the other hand, spikes are the part of the signal which is more difficult to separate from noise, since the low frequency oscillations have relatively high SNR and can often be recovered by simple low-pass filtering. Thus, in the selection the wavelets to be used in the sparsifying transform, we focus on providing a good isolation of the spikes. We do this by studying optical recording where spike locations have been found with traditional methods [Deneux et al. 2016]. We average the detected spikes to obtain spike epochs with high SNR, and

systematically decompose these archetypal spike epochs by discrete wavelet transform (DWT) using different orthogonal wavelets. Then, we rate the efficacy of the decomposition by evaluating the sparseness of the obtained representation via ℓ^1 norm (table 5.1). The results lead us to the selection of the symlet with 3 vanishing moments for the collaborative filtering step, and of the coiflet with 1 vanishing moment (in the scaling function) for the Wiener filtering step. The choice of different wavelets for the two steps is taken to reduce possible artefacts [Lebrun 2012].

Wavelet	ℓ^1 norm
Haar	613.50
Symlet 2	578.16
Symlet 3	487.58
Coiflet 1	470.05

Distance threshold. For the collaborative filtering step, we define the minimum distance between blocks in a group τ assuming that the blocks are not overlapping and that noise variance is unitary. Based on this, we define $\tau = 2q/N_s N_t$ where q is the 0.95 percentile of the χ^2 distribution with $N_s \times N_t$ degrees of freedom. For the Wiener filtering step, we divide this value by 4, assuming that noise has been significantly reduced in the first denoising step.

Table 5.1 Sparsity of the spike epochs in the wavelet domain.

Hard thresholding. The choice of the threshold λ_2 (equation 5.29) of the hard thresholding step constitutes a compromise between noise suppression and preservation of the signal. That is, a high value will remove more noise, but at the cost of partially suppressing the signal that we want to recover. Traditionally, the value $\lambda_2 = 2.7$ has been considered to give the best results in the denoising of natural images [Dabov et al. 2007; Lebrun 2012]. It is however an empirical choice. In our case, we found the slightly more conservative choice $\lambda_2 = 2.5$ to be better suited, in particular to allow better preservation of the spike shape.

Validation and selection of remaining parameters. To verify the effectiveness of the parameters, we create simulated validation data starting from a single-channel electrophysiological recording of the neuronal activity. This time series is characterized by high SNR and sampling frequency. After a simple spike detection by traditional filtering and thresholding, we rescale the signal so that the relative spike amplitude ($\Delta f/f$) matches the one we observe in optical recordings of cerebellar interneurons (a difference of 15–20 % with respect to the baseline). We then rescale the signal to create multiple channels with the same baseline of a reference optical recording. We thus obtain a clean, multichannel time series with statistics which are very similar to those of optical recording data. Then, we artificially contaminate this simulated data by Poisson noise and apply the denoising algorithm. We compare the denoised data with the ground truth to define the best parameters of the algorithm, in particular by examining the MSE both on the full signal and limited to spike epochs only.

We found that limiting the search of similar blocks in a window of $\pm 20\,000$ samples did not impact the denoising performance, while it significantly decreased the computational cost. Similarly, we limited the number of blocks in a group to 16, without observing significant impact on the result. The hard thresholding in the definition of block distances (equation 5.27) gave best results for $\lambda_1 = 0.5$. In the

sparsifying transform, we obtained best results for discrete cosine transform (DCT) along the spatial axis. While wavelets with a high number of vanishing moments resulted in an increased reduction of the MSE over the full signal, we confirmed that the symlet 2 and coiflet 1 wavelets discussed above provided among the best results when considering the MSE of spike epochs only, yielding at the same time a good trade-off with respect to the global MSE.

5.2.4 Validation

To validate the effectiveness of the denoising algorithm, we recorded multiple volumes in the same cell. We apply the full algorithm (sections 5.2.1 and 5.2.2) to the signal of a single volume, using the sum over the remaining volumes to calculate a reference signal which we use to evaluate the denoising performance. Recording from multiple volumes also allows us to reproduce the conditions of recording single-volumes from multiple cells, which is the final aim of our method. We indicate the denoised data by \mathbf{X}_{den} and the reference signal by \mathbf{X}_{ref} , defined as follows:

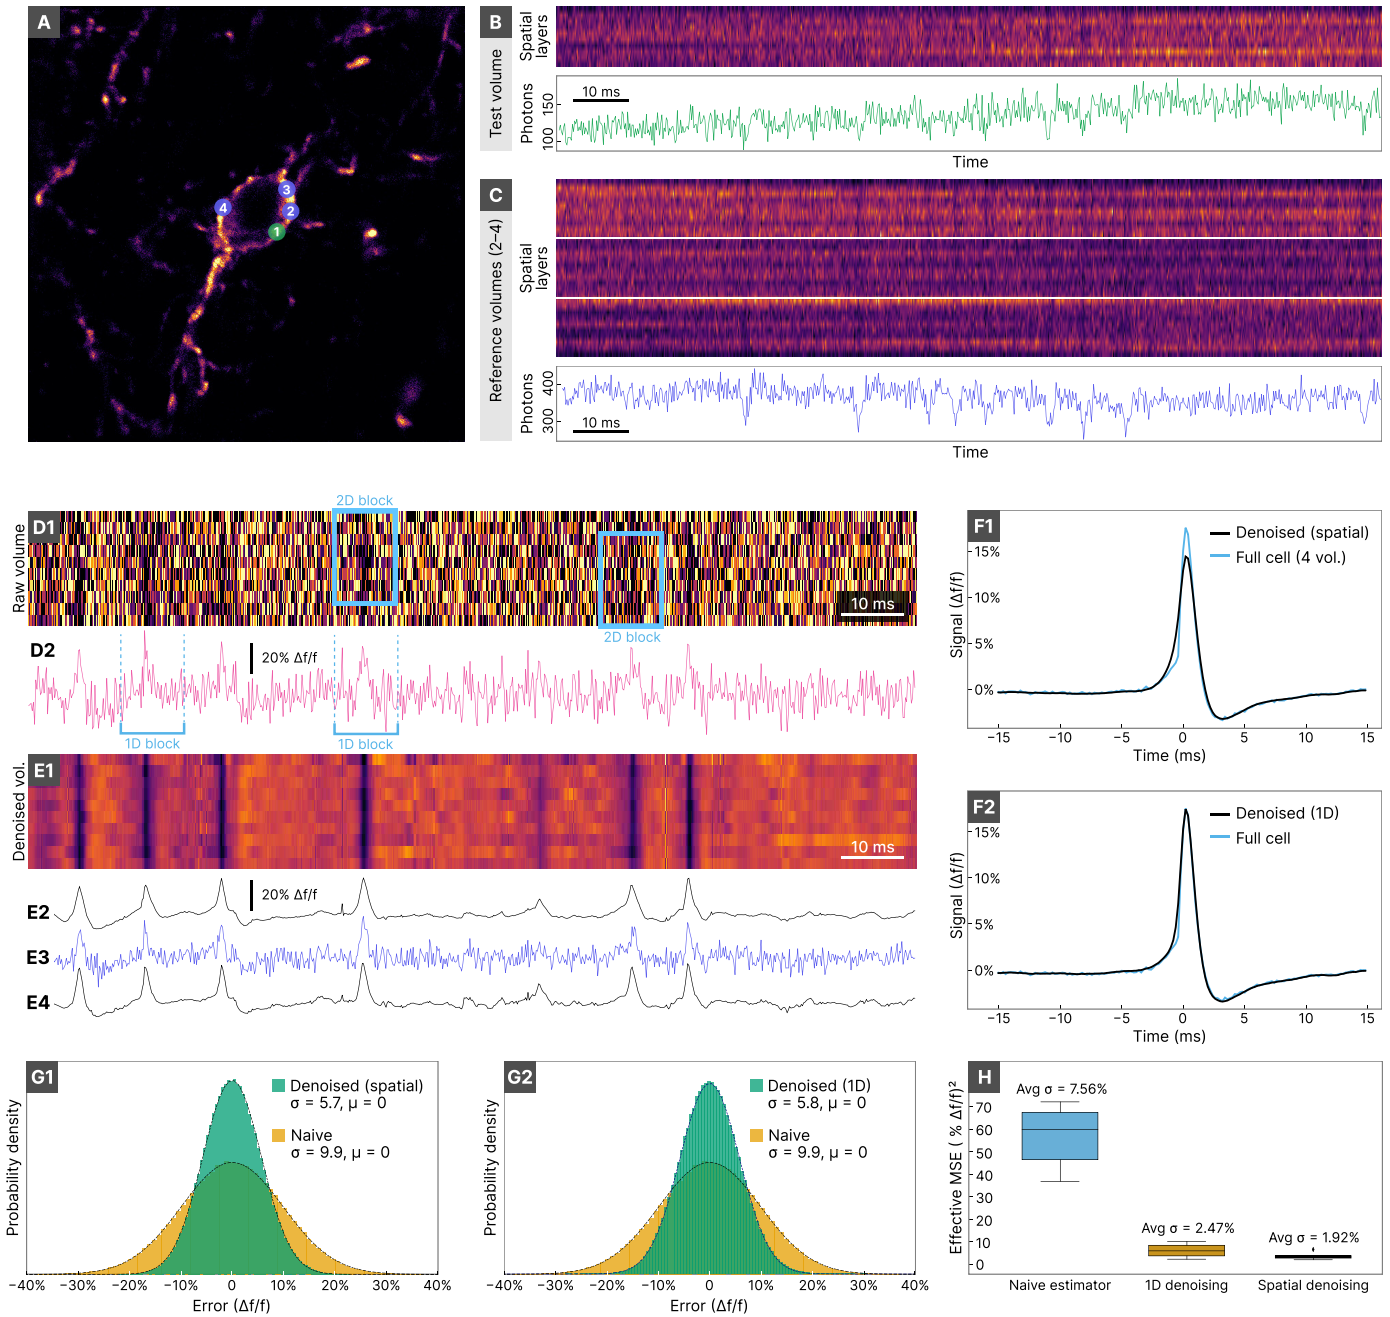
$$\mathbf{X}_{\text{den}} = \mathcal{D}(\mathbf{X}^{(1)}), \quad (5.34)$$

$$\mathbf{X}_{\text{ref}} = \sum_{k=2}^{N_{\text{vol}}} \mathbf{X}^{(k)}, \quad (5.35)$$

where \mathcal{D} indicates the denoising operation.

We post-process both signal estimations as in [Villette et al. 2019] to remove the photobleaching and obtain the relative $\Delta f/f$ representation: first the data channels are summed to obtain a single trace \mathbf{y} and the baseline \bar{y} is estimated by fitting a sum of exponentials; then, the relative trace $\mathbf{y}_{\%}$ is obtained by dividing by the baseline value elementwise, i.e. $\mathbf{y}_{\%} = (y_1/\bar{y}_1, y_2/\bar{y}_2, \dots)$; the signal $\mathbf{y}_{\%}$ is high-pass filtered at 15 Hz to remove animal movement induced oscillations; finally, the sign is inverted to have the spike peaks in the positive direction, obtaining $\mathbf{z} = -\text{LP}_{15\text{ Hz}}(\mathbf{y}_{\%})$.

Figure 5.3 Preliminary results of the denoising method. **A** MLI cell with 4 recorded volumes. Volume 1 is used for testing, while volumes 2–4 are kept for validation. **B** Spatial data of the test volume (brighter colour indicates higher photon count), and correspondent 1D signal obtained by summing all channels. **C** Spatial data of volumes 2–4, and correspondent 1D time series. **D1** Test volume data normalized over baseline ($\Delta f/f$). Examples of 2D blocks, as used in the denoising algorithm, are highlighted in blue. **D2** Post-processed trace of the test volume. Examples of 1D blocks are highlighted in blue. **E1** Spatial denoising of the test volume. Spikes are clearly visible (dark vertical patterns). **E2** Post-processed spatial denoising (obtained by summing the denoised data E1 over space). **E3** Post-processed trace obtained by summing all volumes. **E4** Post-processed trace obtained by 1D denoising of the test trace in D2. **F1** Average spike epoch of spatially denoised signal (black), compared to the full cell average (blue). Spatial denoising attenuates the spike peak and causes some distortion in the rising phase. **F2** Average spike epoch of 1D denoised signal (black), compared to the full cell average (blue). The spike peak is preserved, and only little distortion is present in the raising phase. **G** Distribution of the error with respect to the reference signal for naive estimator (yellow) and denoised (green) with spatial method (G1) and 1D method (G2). **H** Effective MSE for the naive average estimator, 1D denoising, and spatial denoising.



We denote the post-processed denoised and reference signals as \mathbf{z}_{den} and \mathbf{z}_{ref} respectively. Likewise, we consider the signal \mathbf{z}_{nai} obtained by post-processing the raw single-volume data $\mathbf{X}^{(1)}$, corresponding to the naive estimator defined in equation 5.7.

We define the MSE between the denoised signal and the reference as

$$\text{MSE}(\mathbf{z}_{\text{den}}, \mathbf{z}_{\text{ref}}) = \frac{1}{N} \|\mathbf{z}_{\text{den}} - \mathbf{z}_{\text{ref}}\|_2^2, \quad (5.36)$$

where N is the signal length. In principle, we do not know the true signal \mathbf{y} (and its post-processed version \mathbf{z}), thus we cannot calculate exactly the MSE with respect to it. Yet, we can produce an estimate of $\text{MSE}(\mathbf{z}_{\text{den}}, \mathbf{z})$ based on our knowledge of the noise model, if we assume that the denoise operation \mathcal{D} returns an approximately normally distributed estimate. We can verify this by looking at the distribution of the error $\varepsilon = y_{\text{den}} - y_{\text{ref}}$ (figure 5.3). Under this assumption, and with some approximation, calculation is straightforward. We have that

$$y_{\text{den}} \sim \mathcal{N}(c_1 y + c_1 \mu, c_1^2 \sigma_{\text{den}}^2), \quad y_{\text{ref}} \sim \mathcal{N}(c_2 y, a c_2 y + b), \quad (5.37)$$

where c_1, c_2 are arbitrary proportionality factors with respect to the source y and μ is the bias of the denoising estimator \mathcal{D} . If we know consider the signals normalized by the baseline, we have

$$z_{\text{den}} = \frac{y_{\text{den}}}{\bar{y}_{\text{den}}} \approx \frac{y_{\text{den}}}{c_1 \bar{y}}, \quad z_{\text{ref}} = \frac{y_{\text{ref}}}{\bar{y}_{\text{ref}}} \approx \frac{y_{\text{ref}}}{c_2 \bar{y}}, \quad (5.38)$$

from which we can estimate the distribution of the error $z_{\text{den}} - z_{\text{ref}}$:

$$z_{\text{den}} - z_{\text{ref}} \sim \mathcal{N}\left(\frac{\mu}{\bar{y}}, \frac{\sigma_{\text{den}}^2}{\bar{y}^2} + \frac{a c_2 y + b}{c_2^2 \bar{y}^2}\right). \quad (5.39)$$

Using the approximation $c_2 y \approx c_2 \bar{y} \approx \bar{z}_{\text{ref}}$ we obtain an expression for the MSE:

$$\text{MSE}(\mathbf{z}_{\text{den}}, \mathbf{z}_{\text{ref}}) \approx \frac{\mu^2}{\bar{y}^2} + \frac{\sigma_{\text{den}}^2}{\bar{y}^2} + \frac{a \bar{z}_{\text{ref}} + b}{(\bar{z}_{\text{ref}})^2}, \quad (5.40)$$

which allows us to estimate the effective MSE of the denoised result with respect to the true signal \mathbf{z} :

$$\text{MSE}(\mathbf{z}_{\text{den}}, \mathbf{z}) \approx \text{MSE}(\mathbf{z}_{\text{den}}, \mathbf{z}_{\text{ref}}) - \frac{a \bar{z}_{\text{ref}} + b}{(\bar{z}_{\text{ref}})^2}. \quad (5.41)$$

We can verify that these approximations hold by using the same approach for the naive estimate \mathbf{z}_{nai} , for which we can calculate the variance analytically.

5.3 Preliminary results

We tested the method on cerebellar MLIS and Golgi cells recorded in awake, behaving mice. We first estimated the noise model parameters as described in section 5.2.1 (figure 5.2). We found that noise in the optical recordings is characterized

by parameter $a \approx 1$ (figure 5.2 B) and relatively small signal-independent variance b (figure 5.2 C). These results indicate that the noise is compatible with a purely Poissonian model, meaning that photon counting represents the main source of noise.

We acquired signal from 4 up to 7 volumes in the same cell (figure 5.3 A). A single volume (figure 5.3 B) is used to test the denoising method while the remaining volumes (figure 5.3 C) are kept as a reference to evaluate the denoising performances. We tested two variations of the denoising algorithm: spatial and 1D denoising. For the spatial denoising, we applied the adapted block-matching method described in section 5.2.2 to the volume data using 2-dimensional blocks of size 8 px in the spatial dimension and 10 ms in the time dimension (figure 5.3 D1), recovering a 2D denoised estimate of the volume (figure 5.3 E1). For the 1D denoising, we first sum the volume data along the spatial axis to obtain a single-channel signal (figure 5.3 D2), and then apply the denoising method on this 1-dimensional time series using 1D blocks of length 10 ms. In figure 5.3 we show examples of spatial (figure 5.3 E1–E2) and 1D denoising (figure 5.3 E4), compared to the reference trace obtained by summing over all volumes (figure 5.3 E3).

We found that the spatial denoising tended to smoothen the signal more significantly, in particular making spikes less prominent and causing some distortion of the spike shape when compared to the full cell average (figure 5.3 F1). However, we observed that the 1D denoising better preserved the spike prominence and shape (figure 5.3 F2).

We then proceeded to evaluate the performance of the denoising algorithm on the post-processed traces as described in section 5.2.4. We first examined the distribution of the error between the denoised and the reference signal to evaluate possible bias introduced by the estimation. We found that the error distribution was well described by a zero-mean Gaussian for both spatial and 1D denoising (figure 5.3 F1–F2, green coloured distribution), indicating the denoising method does not introduce significant bias in the estimation. This matched the assumptions made in section 5.2.4, allowing us to calculate the effective MSE for the different estimators. Analysis of the recording from 7 different cells showed that the 1D denoising reduced the MSE by a factor close to 10, and spatial denoising by a factor 15 with respect to naive averaging (figure 5.3 H). In practice, this translates to standard deviation values on the post-processed traces of $\sigma_{1D} = 2.47\%$ for 1D denoising and $\sigma_{\text{spatial}} = 1.92\%$ for spatial denoising, compared to $\sigma_{\text{naive}} = 7.56\%$ for simple averaging over the volume.

Finally, in figure 5.4 we show a visual examples for different excitation patterns and cell types. The 1D denoising method significantly reduces the noise both in recordings based on the optimized excitation pattern (figure 5.4 A) and the ULOVE pattern (figure 5.4 B). However, since the spatial information is discarded, the optimized pattern (which is characterized by a more efficient excitation due to increased confinement) presents the best qualitative results in terms of spike prominence. In figure 5.4 C, we present an example of 1D denoising for a Golgi cell, show-




Figure 5.4 Examples of 1D denoising and comparison by cell and excitation pattern. **A** Optical recording from a MLI cell with optimized excitation pattern. **A1** Single volume trace. **A2** Denoised trace (compared to original signal in light grey). **A3** Trace from full cell averaging (7 volumes). **B** MLI with ULoVE excitation pattern. **B1** Single volume trace. **B2** Denoised trace. **B3** Full cell trace (7 volumes). **C** Golgi cell with optimized excitation pattern. **C1** Single volume trace. **C2** Denoised trace. **C3** Full cell trace (7 volumes).

ing how the method adapts to different spiking rates and subthreshold patterns.

The complete denoising procedure, from noise estimation to the final estimate, has been implemented in a Python toolbox which can work with both 1D and 2D data, and that allows arbitrary choice of block sizes, transforms, and other parameters described in section 5.2.3. To provide good performances, the code is fully parallelized, allowing us to process a 5 kHz 1D signal in a time roughly equivalent to the duration of the recording (about 90 s for a 2-minute recording). These characteristics make the technique ready to use in practical applications.




Figure 5.5 Ensemble recording. **A** Ensemble of 4 molecular layer interneuron cells recorded simultaneously. **B** Raw, untreated signals of each cell. **C** Recovered signals after 1D denoising.

5.4 Conclusions and future work

While spatial denoising presented the largest reduction of MSE, we attribute this to increased smoothening of the signal which improves noise reduction for low frequency oscillations. In practical applications, the preservation of the spike prominence and shape is more important than absolute noise reduction, thus we consider 1D denoising to be more suitable for neural time series analysis. In fact, a slightly higher noise level in the high frequencies is tolerable if compensated by a higher fidelity in the reconstruction of the spike shapes. This conclusion also seems to indicate that more focused excitation patterns (characterized by low spatial information, but higher photon flux) should be preferred when we are interested in reliable denoising of the spikes. It also suggests that, for the specific experimental setting analysed here, the temporal dimension provides more rich and useful information about the structure of the signal with respect to the spatial one.

In summary, our technique shows that it is possible to reduce the MSE of optical recordings approximately by factor 10. In other words, a higher noise level in the recordings can be compensated by the denoising, allowing a reduction of the total number of layers acquired for each cell for the same SNR in the results. Pre-processing of the raw data with our method can thus allow us to record and analyse ensembles of multiple cells, all while preserving current acquisition protocols and without the need to change successive steps in the analysis pipeline such as spike detection or unmixing. In figure 5.5 we show a qualitative example of the

application of our technique to an ensemble of four molecular layer interneurons in awake, behaving mice. The denoising post-processing of the original traces provides highly improved SNR, allowing us to clearly distinguish the action potentials (figure 4.1 B–C).

We consider these results very promising, and we consider that this kind of pre-processing could allow us to perform a more in depth study of MLIs functional and structural properties in cortical microcircuits. There are, however, some issues that require further work. First, the proposed approach needs to be validated on a larger set of recordings, comprising data from different animals, to verify its general application. In particular, it would be important to understand the limits of the block matching approach for cells characterized by complex, overlapped spike patterns. Second, a systematic comparison should be made with current methods used in calcium imaging. While the application of such methods is not directly comparable to our experimental setting, exploration of the differences between these techniques could provide a clearer view of how spatial and temporal structures affect the denoising. Moreover, some properties of the denoising method remain to be explored. In particular, future work should include an analysis of how the denoising performances vary depending on the photon flux. If clarified, this information could help in adapting the experimental protocol, ensuring optimal photon flux and sampling rate per cell. Last, the main direction of our future work remains the investigation of neural microcircuits, to reveal and quantify interactions and characteristics of MLIs and other cortical neurons in behaving animals, thanks to the possibilities offered by this new denoising technique.

5.A Maximum-likelihood estimator for the case of a white noise source

We consider here the case of white noise (i.e. Gaussian) with constant variance for all the M channels. We thus define the ζ_{ij} from equation 5.4 to be i.i.d. random variables

$$\zeta_{ij} \sim \mathcal{N}(0, \sigma^2). \quad (5.42)$$

We calculate the reconstruction \tilde{y} by maximum likelihood estimation (MLE) in the parameter space $\theta = (y_1, \dots, y_N, k_1, \dots, k_M)$. Given the model described by equation 5.3 in the case of white noise, the probability of observing x_{ij} given θ is:

$$p(x_{ij}|\theta) = \frac{1}{\sqrt{2\pi\sigma}} \exp\left(-\frac{(x_{ij} - k_i y_j)^2}{2\sigma^2}\right). \quad (5.43)$$

Since the x_{ij} are independent, the likelihood $\mathcal{L}(\theta)$ is the product of the $p(x_{ij}|\theta)$:

$$\mathcal{L}(\theta) = \prod_{i,j} p(x_{ij}|\theta), \quad (5.44)$$

and the maximum likelihood estimate $\tilde{\theta} = (\tilde{y}, \tilde{k})$ is given by

$$(\tilde{y}, \tilde{k}) = \arg \max_{y,k} \mathcal{L}(\theta). \quad (5.45)$$

Maximizing $\mathcal{L}(\theta)$ is equivalent to maximize $\log \mathcal{L}\theta$:

$$\log \mathcal{L}(\theta) = - \sum_{i,j} (x_{ij} - k_i y_j)^2 + \text{const.} \quad (5.46)$$

To find the extremum points, we impose zero derivative with respect to θ , resulting in the system:

$$\begin{cases} \frac{\partial \log \mathcal{L}}{\partial y_n} = 0 & \text{for } n = 1, 2, \dots, N \\ \frac{\partial \log \mathcal{L}}{\partial k_m} = 0 & \text{for } m = 1, 2, \dots, M \end{cases} \quad (5.47)$$

We also add the condition $\|y\|^2 = 1$ to guarantee the uniqueness of the solution. We obtain the relations

$$y_n = \frac{1}{\|\mathbf{k}\|^2} \sum_{i=1}^M x_{in} k_i, \quad (5.48)$$

$$k_m = \sum_{j=1}^N x_{mj} y_j. \quad (5.49)$$

In vector form

$$\mathbf{k} = \mathbf{X}\mathbf{y}, \quad (5.50)$$

$$\mathbf{y} = \frac{1}{\|\mathbf{k}\|^2} \mathbf{X}^T \mathbf{k} \quad (5.51)$$

where we defined the matrix of observations \mathbf{X} with entries $(\mathbf{X})_{ij} = x_{ij}$, and \mathbf{X}^T denotes the transpose of \mathbf{X} . Substituting \mathbf{k} in the second equation, we obtain:

$$\mathbf{X}^T \mathbf{X} \mathbf{y} = (\mathbf{y}^T \mathbf{X}^T \mathbf{X} \mathbf{y}) \mathbf{y}. \quad (5.52)$$

This equation has N solutions, corresponding to the eigenvectors of $\mathbf{A} = \mathbf{X}^T \mathbf{X}$ which represents each an extremum of the likelihood $\mathcal{L}(\cdot)$. The solution corresponding to the absolute maximum is the one associated to the largest eigenvalue. In fact, considering a solution $\mathbf{y}^{(s)}$ with eigenvalue λ and substituting back in equation 5.46 we obtain:

$$\log \mathcal{L}(\mathbf{y}^{(s)}, \mathbf{k}^{(s)} = \mathbf{X}\mathbf{y}^{(s)}) = \lambda + \text{const.}, \quad (5.53)$$

so that the maximum likelihood is obtained for the largest λ . To summarize, denoting by $\mathbf{u}^{(1)}, \mathbf{u}^{(2)}, \dots, \mathbf{u}^{(N)}$ the eigenvectors of $\mathbf{A} = \mathbf{X}^T \mathbf{X}$ associated with eigenvalues $\lambda^{(1)} \geq \lambda^{(2)} \geq \dots \geq \lambda^{(N)}$, the maximum likelihood estimator is given by

$$\begin{aligned} \tilde{\mathbf{y}} &= \mathbf{u}^{(1)} \\ \tilde{\mathbf{k}} &= \mathbf{X}\mathbf{u}^{(1)} \end{aligned} \quad (5.54)$$

Equivalence to singular value decomposition. The solution of equation 5.52 can also be found by performing a singular value decomposition (SVD) of \mathbf{X} . In fact, SVD is the factorization of \mathbf{X} such that

$$\mathbf{X} = \sum_{k=1}^M \sigma^{(k)} \mathbf{u}^{(k)} \otimes \mathbf{v}^{(k)} \quad (5.55)$$

where $\mathbf{u}^{(k)}$ and $\mathbf{v}^{(k)}$ are the left and right singular vectors respectively and $\sigma^{(k)}$ are the associated singular values ($\sigma^{(1)} \geq \sigma^{(2)} \geq \dots \geq \sigma^{(M)}$). Truncation at $k = 1$ provides the approximation in the form $\tilde{\mathbf{X}} = \sigma^{(1)} \mathbf{u}^{(1)} \otimes \mathbf{v}^{(1)}$ which minimizes the Frobenius norm $\|\mathbf{X} - \tilde{\mathbf{X}}\|_F$ (and thus the squared error). In the case of white noise $\tilde{\mathbf{X}}$ is thus the maximum likelihood estimator and the solutions $\tilde{\mathbf{k}}$ and $\tilde{\mathbf{y}}$ would correspond respectively to the left and right singular vectors of \mathbf{X} which are associated to the largest singular value $\sigma = \sqrt{\lambda^{(1)}}$. From a general point of view, this problem can be formulated in the framework of matrix denoising, where the goal is to estimate a low-rank matrix \mathbf{Y} (in our case $\mathbf{Y} = \mathbf{k} \otimes \mathbf{y}$) from a matrix of noisy observations $\mathbf{X} = \mathbf{Y} + \mathbf{Z}$ where \mathbf{Z} is a full-rank matrix of noise [Leeb 2021].

CHAPTER 6

Artefact removal for clinical electroencephalography

In this chapter we turn to signal coming from the entire brain, recorded via electroencephalography (EEG). This is a project I started during the COVID-19 lockdowns, while data collection for my other more biologically-oriented work was put on hold, as a way to support the lab ongoing research on EEG during general anaesthesia. EEG devices have become very common in the clinical setting, especially in intensive care and the operating room, as they provide an accessible, non-invasive tool to monitor the state of the brain. Often, for reasons of speed and simplicity only few electrodes can be used and contamination of the EEG signal by artefacts is inevitable. In the following, I present a wavelet-based method to attenuate artefacts in single-channel EEG recordings in an adaptive way, based on wavelet coefficients statistics from clean portions of the signal. Empirical benchmarking on publicly available EEG datasets shows that this technique can achieve state of the art results for artefacts of different nature, without requiring manual tuning.

This chapter was adapted from “Adaptive single-channel EEG artifact removal with applications to clinical monitoring”, by Matteo Dora and David Holcman, in IEEE Transactions on Neural Systems and Rehabilitation Engineering (2022).¹

6.1 Introduction

Electroencephalography (EEG) is a non-invasive procedure to monitor brain activity by recording electrical signal from electrodes placed on the scalp of a patient. Compact, portable and easy-to-use EEG devices are now routinely used for clinical needs. For example, EEG is employed to monitor the depth of general anaesthesia [Purdon et al. 2015], a method pioneered by the introduction of the bispectral index (BIS) in the 90s [Sigl and Chamoun 1994]. Today, more than ten EEG monitoring devices for anaesthesia monitoring are available on the market [Fahy and Chau 2018]. Similar devices are used for automatic monitoring and scoring of sleep [Fiorillo et al. 2019] or for evaluating neurological disorders in intensive care units (ICUs) [André-Obadia 2018; Maas et al. 2017]. Outside the clinical setting, ambulatory electroen-

¹I must also point out that the method described here is part of the European patent “Computer-implemented method for assisting a general anaesthesia of a subject” (File No. EP21306053).

cephalography (AEEG) allows the acquisition of EEG data through a portable device which can be carried by the patient and can record up to 72 hours of activity. With the fast development of telemedicine, EEG recordings can even be sent in real-time for remote interpretation [Lemesle et al. 2015]. Tele-EEG systems are particularly useful for hospitals which do not dispose of resident neurophysiologists [Coates et al. 2012; Lasierra et al. 2009; Lemesle et al. 2015].

All these EEG monitoring applications are made possible by the little to no supervision needed to operate the device and the ease of installation even by non-technical staff. The downside of operating in these diverse and uncontrolled conditions is the possible contamination of the signal by extraneous sources. This is especially common for devices operating in intensive care units, epilepsy monitoring units, or the operating room [Tatum et al. 2011]. Since signal intensity in scalp EEG is weak (typically 20–100 μV amplitude in adults [Schomer and Da Silva 2012]), even small variations of the electric signal can produce visible artefacts on the EEG traces. Examples are eye movements and blinks, muscular or cardiac activity, motion of electrodes or cables, skin perspiration, and interferences caused by electrical devices such as pacemakers [Schomer and Da Silva 2012; White and Van Cott 2010].

While artefacts can be recognized and sampled out by visual inspection when the analysis is performed offline by practitioners, they pose a major problem in monitoring devices which operate in real-time and unsupervised. In most cases artefacts can be detected quite reliably by localizing non-physiological high voltages or frequencies. Typically, clinical monitoring devices include some artefact detection, and they obviate the problem by temporarily suspending the monitoring and resuming it only after the ratio between artefacted and clean signal drops below a given threshold. During anaesthesia, this interruption deprives the medical staff of relevant information and may delay decision-making. It is thus desirable to assure the continuity of monitoring even in the presence of artefacts.

Several methods have been proposed for EEG artefact removal. In this work, we focus on those methods which are suitable for single-channel recordings. In fact, EEG monitoring devices designed for ICUs or the operating room typically provide only one to four electrodes [Fahy and Chau 2018]. One of the most common techniques is blind source separation (BSS) via independent component analysis (ICA) [Sweeney, Ward, and McLoone 2012; Urigüen and Garcia-Zapirain 2015]. Based on the principle that brain activity is independent of artefactual sources, in ICA-based artefact removal the EEG signal is decomposed into a set of independent sources and the artefactual ones are selectively suppressed. This procedure is semi-automated as it requires identification of artefactual sources. The identification can be performed by visual inspection, but automated classification methods are also available [Jas et al. 2017; Nolan, Whelan, and Reilly 2010]. The quality of ICA decomposition is highly dependent on the preprocessing procedure [Klug and Gramann 2020]. Although good results can be achieved by applying high-pass (or band-pass) filters before ICA, this filtering procedure may also remove relevant informa-

tion [Klug and Gramann 2020; Pignat et al. 2013]. Although source separation techniques require multiple channels, they can be used on a single-channel in association with wavelet transform (WT) or empirical mode decomposition (EMD). The decomposition by WT or EMD allows to split a one-dimensional signal into multiple components which are then used as a multichannel input for BSS. Examples of this approach are WICA [Inuso et al. 2007], EMD-ICA [Mijović et al. 2010], and EMD-canonical correlation analysis (CCA) [Sweeney, McLoone, and Ward 2012]. However, one major drawback of source separation techniques such as ICA is that they cannot be applied on-line [Sweeney, Ward, and McLoone 2012]. Although some real-time applications have been proposed by recurrently computing ICA on a sliding window [Esposito et al. 2003; Mayeli et al. 2016], no efficient on-line algorithm is today available. Moreover, the performances of ICA are affected by the length of the signal [Onton et al. 2006]: on one hand, a too short signal can prevent reliable separation of sources; on the other hand, when the signal is too long, the properties of the sources can change in time leading to improper isolation of the artefacts. These limitations make ICA and similar BSS-based methods unpractical in the context of real-time continuous monitoring.

An alternative to ICA are wavelet-based methods. The effectiveness of the wavelet transform (WT) at capturing time-frequency patterns at multiple scales has made it a common technique in various noise removal tasks, with applications ranging from images [Averbuch, Lazar, and Israeli 1996], audio [Averbuch, Neittaanmäki, and Zheludev 2019], to physiological signals [Le Van Quyen and Bragin 2007; Unser and Aldroubi 1996; Worrell et al. 2012]. The WT decomposes a signal into multiple sets of coefficients representing time-frequency patterns of the signal. Denoising a signal is typically achieved by thresholding the wavelet coefficients, eliminating those corresponding to noise. Due to their flexibility, efficiency, and robustness, wavelet-based methods have been widely used for EEG artefact removal. Compared to blind source separation, the WT has a low computational cost, can be applied on-line, and does not require multiple channels. Wavelet-based methods can be very efficient at eliminating artefacts when the chosen wavelet basis allows good separation of the signal from noise [Urigüen and Garcia-Zapirain 2015]. Today, methods can combine wavelet decomposition with ICA [Azzarboni et al. 2002; Inuso et al. 2007; Mammone, La Foresta, and Morabito 2012].

Other methods, such as adaptive filtering [Correa et al. 2007], multichannel Wiener filtering [Somers, Francart, and Bertrand 2018], spatial filtering [Guarnieri et al. 2018], were not considered here as they are not suitable for single-channel EEG or require auxiliary signals which are generally not available in clinical monitoring applications.

In this chapter, we introduce a wavelet-based method called wavelet quantile normalization (WQN) which allows eliminating transient artefacts [Schomer and Da Silva 2012; Tatum et al. 2011] from single-channel EEGs. The method consists in normalizing the wavelet coefficients of artefacted intervals, using temporally adjacent uncontaminated signal as a reference. This approach is relevant for continu-

ous monitoring during general anaesthesia and coma, where we expect a temporal continuity of the power distribution of the signal. The parameters used in the WQN method are estimated from the EEG signal itself with no prior human intervention. In this work, we consider the segmentation of artefactual epochs to be known. This requirement can be alleviated by using an appropriate artefact detection algorithm [Jas et al. 2017; Nolan, Whelan, and Reilly 2010], similarly to those employed in BSS-based methods, making WQN artefact removal completely automated. Most EEG monitoring devices already include artefact detection algorithms so that automated analysis can be suspended when the signal quality is degraded.

This chapter is organized as follows. In section 6.2 we introduce the notation and the discrete wavelet transform (6.2.1), we briefly present the classical wavelet denoising approaches which we used to benchmark our method (6.2.2), and finally we present the wavelet quantile normalization method (6.2.3). In section 6.3, we describe the validation datasets and methodology: we briefly present the datasets used for benchmarking (6.3.1), the performance metrics (6.3.2), and the parameters used in the comparison (6.3.3).² In section 6.4 we discuss the results and present potential application to anaesthesia monitoring. Finally, we present our conclusions in section 6.5.

6.2 Methods

6.2.1 Discrete Wavelet Transform

We decompose the EEG signal x into time-frequency components using the discrete wavelet transform (DWT). The DWT coefficients are obtained by decomposing the signal onto an orthogonal basis obtained by dilating (scaling) and translating in time a wavelet function $\psi(t)$ and scaling function $\varphi(t)$ [Daubechies 1992; Jaffard, Meyer, and Ryan 2001]. The DWT can be computed efficiently with Mallat's pyramidal algorithm, implemented by means of a hierarchical set of subband filters [Mallat 1989]. The DWT decomposition of the signal $x(t)$ to the level M is defined as [Daubechies 1992]:

$$x(t) = \sum_n c_{M,n} \varphi_{M,n}(t) + \sum_{m=1}^M \sum_n d_{m,n} \psi_{m,n}(t), \quad (6.1)$$

where

$$\psi_{m,n}(t) = 2^{-m/2} \psi(2^{-m}t - n), \quad (6.2)$$

$$\varphi_{m,n}(t) = 2^{-m/2} \varphi(2^{-m}t - n), \quad (6.3)$$

and $c_{m,n} = \langle x, \varphi_{m,n} \rangle$ and $d_{m,n} = \langle x, \psi_{m,n} \rangle$ are the approximation and detail coefficients at level m . The scalar product is defined by $\langle f, g \rangle = \sum_{t \in \mathbb{Z}} f(t) g(t)$.

²The Python code implementing the WQN algorithm and the benchmarking procedure is available on Zenodo (<https://doi.org/10.5281/zenodo.4783449>).

The DWT conserves energy so that:

$$\|x\|_2^2 = \sum_n \|c_{M,n}\|_2^2 + \sum_{m=1}^M \sum_n \|d_{m,n}\|_2^2. \quad (6.4)$$

6.2.2 Denoising by Wavelet Thresholding

Wavelet denoising methods rely on the assumption that artefactual components can be well isolated when the signal is represented in a wavelet basis. Denoising is performed by identifying and removing the wavelet coefficients associated with artefacts. For that goal, the signal is first decomposed onto a wavelet basis via DWT. Then, the wavelet coefficients are either removed or attenuated by a thresholding procedure. Finally, the corrected signal is recovered by inverting the wavelet transform.

The main challenge of this approach is therefore the selection of the appropriate threshold. While this could be done by manual inspection of the wavelet coefficients, this is unpractical in most applications. Various approaches for the automatic selection of thresholds of wavelet denoising have been proposed [Donoho and Johnstone 1994a; Donoho and Johnstone 1994b; Donoho et al. 1995; Johnstone and Silverman 1997]. A common choice is the universal threshold [Donoho and Johnstone 1998] defined as:

$$\theta_U = \sigma \sqrt{2 \ln N}, \quad (6.5)$$

where σ is the standard deviation of the wavelet coefficients w and N is the number of coefficients. The value of σ can be estimated from the data using the median absolute deviation (MAD) estimator $\sigma = k \cdot \text{median}(|w|)$ with $k \approx 1.4826$ [Donoho and Johnstone 1994a].

Threshold values can be calculated separately for each DWT level. This procedure extends the universal threshold (equation 6.5) to signals which are only weakly stationary and are correlated in time [Johnstone and Silverman 1997]. This is relevant in EEG artefact removal applications where stationarity is not guaranteed [Bajaj et al. 2020; Chavez et al. 2018]. In that case, equation 6.5 becomes $\theta_m = \sigma_m \sqrt{2 \ln N_m}$ where σ_m and N_m refer to the wavelet coefficients at level m .

The thresholding procedure is usually carried out by applying a hard or a soft thresholding function on the wavelet coefficients w , defined as follows:

$$\lambda_{\text{hard}}(w) = \begin{cases} w, & \text{if } w < \theta \\ 0, & \text{otherwise} \end{cases} \quad (6.6)$$

$$\lambda_{\text{soft}}(w) = \begin{cases} w, & \text{if } w < \theta \\ \theta, & \text{otherwise,} \end{cases} \quad (6.7)$$

where θ is the chosen threshold (see figure 6.1).

An alternative approach is the surrogate-based artefact removal method (SUBAR) [Chavez et al. 2018], where time-varying thresholds are defined by considering the




Figure 6.1 Threshold functions.
A Soft thresholding (equation 6.7).
B Hard thresholding (equation 6.6).

statistical deviation computed over an ensemble of surrogate time series. The surrogates are randomly generated via the iterative amplitude-adjusted Fourier transform (IAAFT) algorithm [Schreiber and Schmitz 1996; 2000] to produce stationary time series which have the same distribution as the original signal and a similar power spectrum. This method effectively amounts to detecting deviations from stationarity in the EEG signal and correct them via thresholding of the wavelet coefficients with the function λ_{soft} (equation 6.7).

6.2.3 Wavelet Quantile Normalization

As an alternative to wavelet thresholding, we develop here a novel renormalization technique for the wavelet coefficients which allows to adaptively attenuate artefactual components. The proposed technique allows to flexibly remove artefacts of different nature without requiring manual adjustments. The main idea of the method is that in most monitoring contexts the brain rhythm does not change dramatically during short time intervals (e.g. < 1 s). For example, in clinical monitoring of anaesthesia brain activity is mostly controlled by hypnotic drugs. These drugs take effect gradually, producing brain waves patterns with a timescale of minutes [Purdon et al. 2015]. We can thus consider that the distribution of energy across frequency bands is continuous and that there are no abrupt changes in the power of a band. We do not expect such continuity for transient brain activity, such as the response to an external stimulus [André-Obadia 2018]. Applying the proposed method in those situations would be inappropriate as it would probably attenuate genuine brain activity. In the following, we shall focus on continuous monitoring of the brain, where changes in the activity happen gradually. Under this continuity condition, the statistical characteristics such as the distribution of energy across frequencies of a short EEG fragment will be similar to the signal immediately preceding or following it.

The WQN method works as follows. First, artefacts in the EEG signal are segmented (figure 6.2 A) with an appropriate detection algorithm [Islam, Ghorbanzadeh, and Rastegarnia 2021; Islam, Rastegarnia, and Yang 2016; Jas et al. 2017; Nolan, Whelan, and Reilly 2010; Sweeney et al. 2012]. For each artefact, we construct a

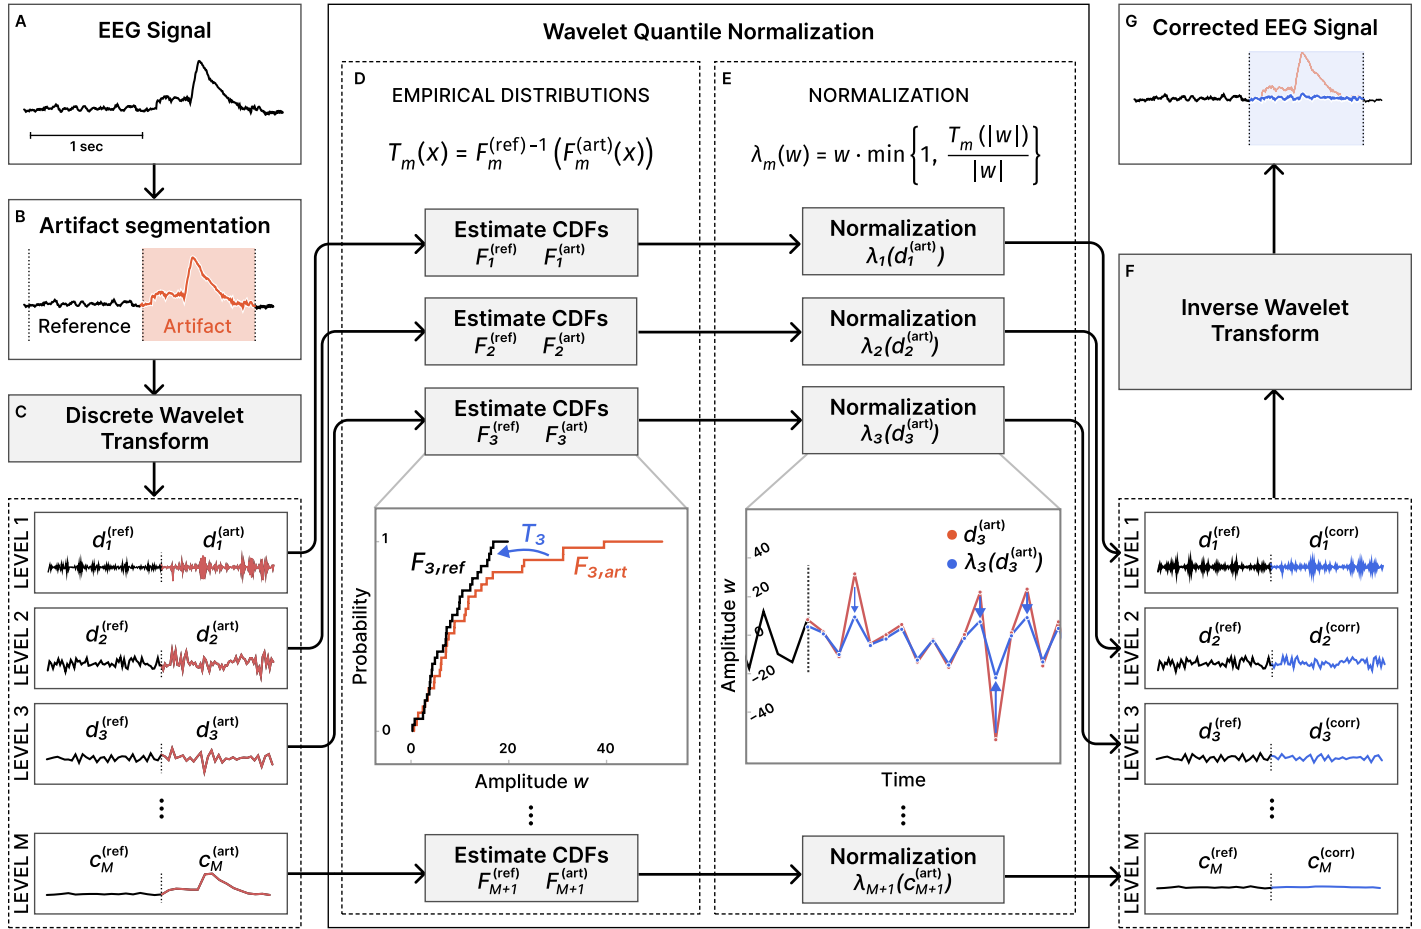


Figure 6.2 Diagram of proposed artefact removal method. **A** Input EEG signal (single-channel). **B** Artefact segmentation. **C** Decomposition with DWT on M levels. **D** Estimation of empirical cumulative density functions and transformation T_m (equation 6.9). **E** Normalization of wavelet coefficients (equation 6.10). **F** Reconstruction from corrected wavelet coefficients (equations 6.11 to 6.13). **G** Artefact-corrected EEG signal.

reference signal by considering some short portions of the EEG immediately before and/or after the occurrence of the artefact (figure 6.2 B). Both the artefact and the reference segments are then decomposed via an M -level DWT (figure 6.2 C), obtaining the coefficients $d_m^{(art)}$, $c_M^{(art)}$ and $d_m^{(ref)}$, $c_M^{(ref)}$. For uniformity of notation, we denote all coefficients as $w_m = d_m$ for $m = 1, \dots, M$ and $w_{M+1} = c_M$.

For each level $m = 1, \dots, M + 1$, we compute the empirical cumulative density function (CDF) $F_m^{(ref)}$, $F_m^{(art)}$ of the amplitude of the wavelet coefficients w_m^{ref} for reference and artefact (figure 6.2 D) defined by

$$F_m(x) = \frac{1}{N_m} \sum_{n=1}^{N_m} 1_{|w_{m,n}| < x}, \quad (6.8)$$

where $1_{|w| < x}$ is the indicator function which takes value 1 if $|w| < x$ and 0 otherwise. We can transform the wavelet coefficients w_m^{art} of the artefact segment so that their

distribution of amplitude will match that of the reference segment. This mapping is provided by the transformation T_m defined as

$$T_m(x) = F_m^{(\text{ref})-1} \left(F_m^{(\text{art})}(x) \right), \quad (6.9)$$

where $F_m^{(\text{ref})-1}$ indicates the inverse of $F_m^{(\text{ref})}$. T_m transforms the wavelet coefficients $w_m^{(\text{art})}$ so that their amplitude distribution (and thus equivalently their energy distribution) becomes identical to that of $w_m^{(\text{ref})}$ (see appendix 6.A). This transformation adaptively attenuates high-power artefactual coefficients. Since we are only interested in reducing the power of artefacts, we restrict the normalization of coefficients to prevent amplification. Thus, we define the normalization function

$$\lambda_m(w) = w \cdot \min \left\{ 1, \frac{T_m(|w|)}{|w|} \right\}, \quad (6.10)$$

which maps a coefficient w from $w_m^{(\text{art})}$ to its possibly attenuated value (figure 6.2 E). Due to the energy conservation of the wavelet transform (equation 6.4), attenuation of the wavelet coefficients leads to a reduction in the energy of specific time-frequency components of the original signal. We define the corrected coefficients as:

$$d_{m,n}^{(\text{corr})} = \lambda_m \left(d_{m,n}^{(\text{art})} \right), \quad m = 1, \dots, M; \quad (6.11)$$

$$c_{M,n}^{(\text{corr})} = \lambda_{M+1} \left(c_{M,n}^{(\text{art})} \right). \quad (6.12)$$

Finally, we reconstruct the corrected version of the artefacted segment by inverting the DWT as in equation 6.1, using the corrected coefficients (figure 6.2 F):

$$x^{(\text{corr})}(t) = \sum_n c_{M,n}^{(\text{corr})} \varphi_{M,n}(t) + \sum_{m=1}^M \sum_n d_{m,n}^{(\text{corr})} \psi_{m,n}(t). \quad (6.13)$$

This algorithm can be applied to every artefacted signal interval to recover a corrected EEG (figure 6.2 G). For real-time applications, this procedure can be applied on short (e.g. 1 s) EEG fragments, only requiring a small buffer containing the distribution of the reference coefficients.

6.3 Validation

6.3.1 Datasets for benchmarking

To validate the proposed method, we compare its performances on both unmodified and semi-simulated EEG recordings using publicly available datasets [Klados and Bamidis 2016; Sweeney et al. 2012; Zhang et al. 2020]. Each dataset comprises an artefacted EEG signal and a corresponding reference EEG signal used as ground truth. Artefact detection labels are given as input to the different algorithms.

Physiobank Motion Artefacts dataset [Goldberger et al. 2000; Sweeney et al. 2012]. This dataset contains 24 EEG recordings. The EEG signal was acquired by two closely placed electrodes producing highly correlated signals. The cable of one electrode was pulled at intervals of roughly 2 minutes to produce random artefacts, while the other electrode was kept still. Artefact removal can then be applied on the signal produced by the contaminated electrode, while the other channel is used as a ground truth reference. We filtered both channels in the range 0.1–100 Hz with a Butterworth forward-backward filter of order 4 and down sampled to 256 Hz. Labelling of the artefacts was performed by thresholding the difference between the reference and the artefacted electrode. Recordings 10, 20, 21 were partially excluded after visual inspection revealed artefacts in the reference channel.

Semi-simulated EEG and electrooculogram (EOG) [Klados and Bamidis 2016]. This dataset contains 55 recordings from 19 EEG electrodes (200 Hz) and an equal number of EOG signals. The EOG signals are combined linearly to the EEG signal [Klados and Bamidis 2016] to produce artefacted recordings. We modified the dataset to obtain 30 seconds recordings where the central 10 seconds are corrupted by EOG and marked as an artefact. Recording 36 was excluded after visual inspection.

Denoise-Net [Zhang et al. 2020]. This dataset includes thousands of 2-second epochs of EEG, EOG and electromyogram (EMG) recorded at 256 Hz. We combined the clean EEG epochs with EOG and EMG to produce simulated datasets with 3400 artefacted epochs. These epochs were produced by summing one second artefact signal (EOG, EMG, or EOG+EMG) to the second half of the clean epoch. The amplitude of the artefact was rescaled to obtain signal to noise ratio (SNR) values in the range -20 – 5 dB (figure 6.3). A limitation of this dataset is that the clean and artefacted epochs are independent and there is no guarantee that they belong to the same subject.

6.3.2 Performance metrics

We employed different metrics to evaluate the performances of the artefact removal methods. Benchmarks were run with Python 3.9.7 on a Linux workstation equipped with an Intel Xeon Gold 5218R CPU (2.10 GHz).

Change in signal to noise ratio. We define the metric Δ SNR as the change in the SNR before and after artefact removal. We only consider the signal epochs which were labelled as artefacted. The calculation of Δ SNR is defined as [Sweeney, McLoone,

and Ward 2012]:

$$\text{SNR}_{\text{before}} = 10 \log_{10} \frac{\text{Var}(x_0)}{\text{Var}(x^{(\text{art})} - x_0)}, \quad (6.14)$$

$$\text{SNR}_{\text{after}} = 10 \log_{10} \frac{\text{Var}(x_0)}{\text{Var}(x^{(\text{corr})} - x_0)}, \quad (6.15)$$

$$\Delta\text{SNR} = \text{SNR}_{\text{after}} - \text{SNR}_{\text{before}}, \quad (6.16)$$

where $\text{Var}(\cdot)$ is the variance estimator, x_0 is the ground truth signal, $x^{(\text{art})}$ is the artefacted signal, and $x^{(\text{corr})}$ is the signal obtained after artefact removal.

Normalized mean squared error. The normalized mean-squared error (NMSE), in decibels, is defined as:

$$\text{NMSE} = 10 \log_{10} \frac{\sum_i |x_i^{(\text{corr})} - x_{0,i}|^2}{\sum_i |x_{0,i}|^2}, \quad (6.17)$$

where x_i is the i -th sample of the signal x . Similarly to the definition of ΔSNR , we only consider signal epochs which are labelled as artefacted.

Change in correlation. We defined the change in correlation ΔR before and after the artefact removal as:

$$\Delta R = R_{\text{after}} - R_{\text{before}}, \quad (6.18)$$

where R_{before} and R_{after} are the Pearson correlation coefficients between the ground truth and the signal before and after artefact removal respectively.

Improvement in spectral coherence. We defined the improvement in coherence I_{coh} before and after the artefact removal as:

$$I_{\text{coh}} = \frac{C_{\text{after}} - C_{\text{before}}}{1 - C_{\text{before}}}, \quad (6.19)$$

where C_{before} and C_{after} denote the average magnitude squared coherence estimates in the 0–40 Hz band, calculated between the ground truth and the signal before and after artefact removal respectively.

Execution time. The execution time (see Table 6.2) was calculated using a sample of duration 30 s at 256 Hz taken from the Physiobank dataset [Sweeney et al. 2012]. Each method was repeatedly applied to the same EEG sample, selecting the fastest run out of 10. All methods except SUBAR were implemented in Python and included in the published code. The original Matlab implementation provided by the authors was used for the SUBAR method [Chavez et al. 2018].

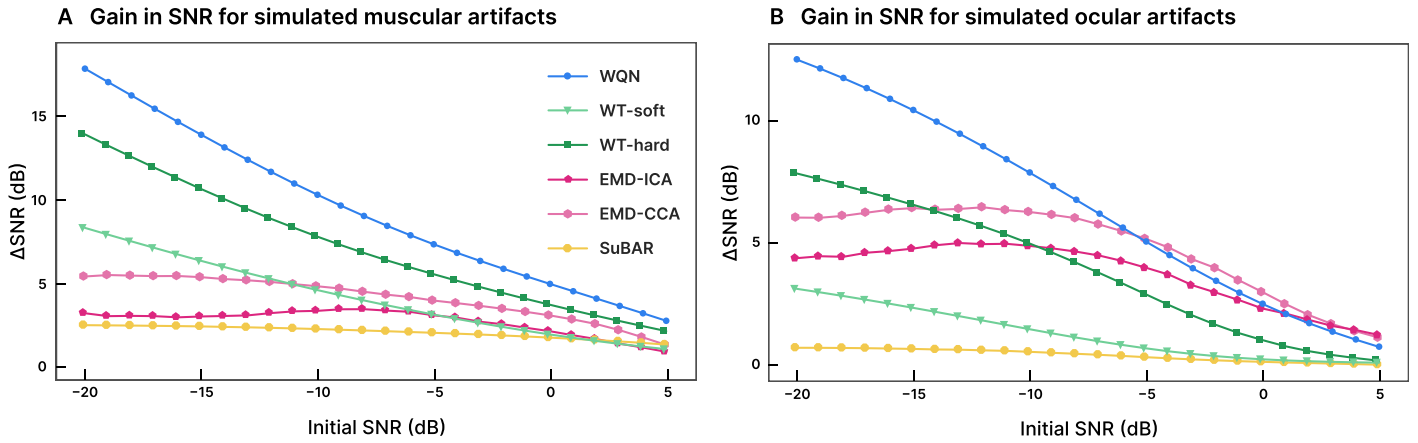


Figure 6.3 Noise reduction (Δ SNR) for varying level of simulated signal to noise ratio (SNR). **A** Simulated EMG dataset. **B** Simulated EOG dataset.

6.3.3 Methods parameters

We compare (Table 6.1) our proposed wavelet quantile normalization method against the classical wavelet thresholding and the SUBAR method [Chavez et al. 2018]. We implement the wavelet thresholding as described in section 6.2.2, using the universal threshold defined in equation 6.5 separately for each decomposition level. Estimation of the standard deviation σ_m values is obtained by the MAD estimator $\sigma_m = 1.4826 \cdot \text{median}(|w_m|)$. We report the results for both hard and soft thresholding functions. The SUBAR method is applied on 30 s windows with no overlapping, using 1000 surrogates. For all tested methods, we used the symlet wavelet family with 5 vanishing moments [Daubechies 1992], as it was found to be suitable for EEG artefact removal [Chavez et al. 2018] and allows us to perform a direct comparison with the SUBAR method. In all cases, the correction was only applied on regions which were labelled as artefactual. EMD-CCA and EMD-ICA were implemented according to [Sweeney, McLoone, and Ward 2012]. Up to 10 intrinsic mode functions were obtained from EMD and used as input for the CCA or ICA source separation. Artefactual sources were removed according to the methodology described in [Sweeney, McLoone, and Ward 2012]: sources which, upon removal, increased the correlation to the reference signal were considered artefactual and thus suppressed. While this method cannot be employed in practice, since the reference signal is unknown, it ensures that the effectiveness of artefact removal is independent of the source classification strategy.

6.4 Results and Discussion

Dataset	Method	Δ SNR	NMSE	Δ R	I_{coh}
Motion Artefacts [Sweeney et al. 2012]	WQN (ours)	16.57	0.66	0.51	0.27
	WT-hard (equation 6.6)	13.25	4.05	0.43	0.13
	WT-soft (equation 6.7)	7.60	9.69	0.22	0.10
	SUBAR [Chavez et al. 2018]	0.95	16.28	0.02	0.04
	EMD-ICA[Sweeney, McLoone, and Ward 2012]	7.45	9.80	0.25	0.04
	EMD-CCA [Sweeney, McLoone, and Ward 2012]	5.56	11.69	0.18	0.05
Semi-simulated EOG [Klados and Bamidis 2016]	WQN	6.13	-5.80	0.06	0.18
	WT-hard	3.32	-2.97	0.03	0.11
	WT-soft	1.55	-1.22	0.01	0.15
	SUBAR	1.72	-1.39	-0.01	-1.06
	EMD-ICA	2.08	-1.76	0.03	-0.02
	EMD-CCA	2.35	-2.03	0.03	0.02
Denoise-Net EMG [Zhang et al. 2020]	WQN	17.84	2.67	0.46	0.26
	WT-hard	13.90	6.48	0.34	0.19
	WT-soft	8.28	11.85	0.18	0.05
	SUBAR	2.43	17.62	0.04	0.03
	EMD-ICA	3.17	16.90	0.10	0.00
	EMD-CCA	5.35	14.96	0.14	0.01
Denoise-Net EOG [Zhang et al. 2020]	WQN	12.48	10.41	0.24	0.24
	WT-hard	7.81	13.95	0.17	0.30
	WT-soft	3.08	18.39	0.05	0.17
	SUBAR	0.65	20.83	0.00	0.21
	EMD-ICA	4.31	17.36	0.15	-0.10
	EMD-CCA	5.83	16.10	0.18	-0.03
Denoise-Net EOG+EMG [Zhang et al. 2020]	WQN	13.47	9.20	0.25	0.23
	WT-hard	8.46	13.13	0.16	0.17
	WT-soft	4.05	17.28	0.06	0.05
	SUBAR	1.24	19.90	0.01	0.04
	EMD-ICA	2.25	18.55	0.08	0.00
	EMD-CCA	2.93	18.02	0.10	0.01

Table 6.1 Result of validation on different datasets. For the proposed wqn method, wavelet hard thresholding (WT-hard), wavelet soft thresholding (WT-soft), SUBAR, EMD-ICA, and EMD-CCA methods. The values are obtained by averaging over all recordings in a dataset. Δ SNR and NMSE values are in dB. The best performance for each metric is highlighted in black.

6.4.1 Artefact removal effectiveness

To validate the effectiveness of the WQN method we benchmarked it against alternative artefact removal algorithms which are suitable for single-channel EEG recordings: wavelet thresholding (with both soft and hard thresholding functions), EMD-ICA, EMD-CCA [Sweeney, McLoone, and Ward 2012], and surrogate-based artefact removal (SUBAR) [Chavez et al. 2018]. We included artefacts of different nature (movement, ocular, muscular) using datasets of semi-simulated and real EEGs (section 6.3.1), where we used a ground truth signal to evaluate the performances of artefact removal algorithms. We quantified the reduction of artefactual components using the normalized mean squared error (NMSE), the gain in SNR (ΔSNR), the change in correlation (ΔR), and the improvement in spectral coherence (I_{coh}), obtained by comparing the output of the artefact removal algorithms to the ground truth signal (section 6.3.2). We note that in some cases, an algorithm can show an improvement in terms of ΔSNR and NMSE but at the same time destroy the structure of the signal, losing the information content of the data. This is particularly the case for high-power artefacts. To avoid this effect, we consider the change in correlation (ΔR) and coherence (I_{coh}) as well as NMSE and ΔSNR . Since a degradation of the signal will coincide with a lower correlation to the ground truth, we consider that a simultaneous improvement in ΔSNR , ΔR , and I_{coh} is a solid indication of reliable artefact removal. In particular, improvement of coherence allows detecting possible spectral distortions introduced by the artefact removal.

We found that WQN performed better among the tested methods in terms of SNR and NMSE, producing the highest increase to signal to noise ratio (ΔSNR) and the lowest normalized mean squared error on all datasets. WQN also outperformed the other methods in terms of improvement of the correlation and coherence, confirming absence of information loss or spectral distortion. For the tested data, the performances of the SUBAR method were lower to those of wavelet thresholding. This result can be understood because the SUBAR was designed to remove short, isolated artefacts and was reported to be ill-suited for large EOG artefacts [Chavez et al. 2018]. EMD-ICA and EMD-CCA showed better performances in removing ocular and motion artefacts, but were significantly outperformed by wavelet-based methods in the case of high-frequency muscular artefacts. This result confirms that ICA and similar BSS methods might not be optimal in removing EMG activity [Muthukumaraswamy 2013; Urigüen and Garcia-Zapirain 2015]. Our tests also revealed problems in the convergence of the FastICA algorithm when applied to the 2-second EEG samples from the Denoise-Net data, highlighting the limitations of this method for short time windows. The benchmark results are shown in Table 6.1.

We then investigated how the performances depend on the initial signal-to-noise level. To test this, we rescaled the amplitude of the artefacts in the Denoise-Net datasets (see section 6.3.1) to obtain SNRs between -20 dB (large, powerful artefacts) and 5 dB (small artefacts). In Figure 6.3 we show the ΔSNR for the various initial SNR values on muscular and ocular artefacts. WQN consistently outperformed the other

Method	Computational Complexity	Execution Time
Wavelet thresholding	$\mathcal{O}(MN)$	3 ms
WQN (ours)	$\mathcal{O}(MN + K_{\text{art}}N_{\text{art}} \log N_{\text{art}})$	2 ms
EMD-CCA	$\mathcal{O}(MN + M^2N)$ [Flandrin et al. 2013]	314 ms
EMD-ICA (FastICA)	$\mathcal{O}(MN + K_{\text{it}}M^2N)$ [Flandrin et al. 2013; Zarzoso, Comon, and Kallel 2006]	419 ms
SUBAR	$\mathcal{O}(K_{\text{su}} K_{\text{it}} N \log N)$ [Chavez et al. 2018]	31.1 s

Table 6.2 Computational costs for a signal of length N using M levels of decompositions. The execution time is the empirical CPU time required to analyse 30 s EEG signal at 256 Hz.

methods in the case of muscular artefacts, showing higher ΔSNR for all initial SNR levels. In this case, SUBAR outperformed soft wavelet thresholding for weak muscular artefacts (figure 6.3 A). For ocular artefacts, EMD-based methods and WQN showed the best performances, with WQN significantly outperforming other methods for low SNRS (< -10 dB).

In figure 6.4 we show three examples of corrected artefacts obtained with the WQN method: motion artefact (figure 6.4 A), eye blink (figure 6.4 B), and muscular activity (figure 6.4 C). The results show a significant reduction of the artefacts without distortion of the EEG signal, highlighting the flexibility and robustness of the proposed method.

6.4.2 Computational complexity

We consider a single-channel EEG signal of length N samples. The computational complexity of an M -level discrete wavelet transform is $\mathcal{O}(MN)$ [Mallat 1989; Safieddine et al. 2012]. In wavelet thresholding algorithms, the threshold value is typically fixed beforehand or can be determined in $\mathcal{O}(N)$ time based on the signal statistics (such as the universal thresholding presented in equation 6.5). Thus, the leading term in computational complexity is $\mathcal{O}(MN)$ of the DWT decomposition.

A similar result holds for the proposed WQN algorithm. The M -level DWT can be performed in $\mathcal{O}(MN)$ operations. For an artefact of length N_{art} , the empirical cumulative distribution is obtained by sorting the coefficients with a cost of $\mathcal{O}(N_{\text{art}} \log N_{\text{art}})$. The total computational complexity can thus be given as $\mathcal{O}(MN + K_{\text{art}}N_{\text{art}} \log N_{\text{art}})$ for K_{art} artefacts of maximum length N_{art} . Since the number and length of artefacts is typically limited, the computational complexity will be dominated by the $\mathcal{O}(MN)$ term. In practice, the WQN algorithm can be faster than WT since the DWT has to be computed only on short fragments of the signal around the artefact.

In the SUBAR method [Chavez et al. 2018] the leading term of computational complexity comes from the generation of Fourier surrogates via iterative amplitude adjusted Fourier transform (IAAFT) [Schreiber and Schmitz 1996]. Each iteration of

the IAAFT algorithm requires the computation of Fourier transform and sorting of the coefficients, with both operations requiring $\mathcal{O}(N \log N)$ time. The generation of K_{su} surrogates has thus a computational complexity of $\mathcal{O}(K_{\text{su}} K_{\text{it}} N \log N)$ where K_{it} is the number of iterations of IAAFT.

EMD-based methods such as EMD-ICA and EMD-CCA require the decomposition of the signal as a sum of intrinsic mode functions (IMFs). Although EMD is commonly reported to be computationally demanding [Wang et al. 2014], its order of complexity is $\mathcal{O}(N \log N)$ [Flandrin, Rilling, and Goncalves 2004; Wang et al. 2014]. Moreover, if only the first M IMFs are considered, the complexity becomes $\mathcal{O}(MN)$ [Safieddine et al. 2012; Wang et al. 2014], making it equivalent to the M -level DWT. In our empirical measures of the CPU time (table 6.2), the EMD-based algorithms remain significantly slower than wavelet decomposition (~ 300 ms versus 3 ms respectively). ICA is commonly implemented via the FastICA algorithm which has computational cost $\mathcal{O}(K_{\text{it}} M^2 N)$ [Safieddine et al. 2012; Zarzoso, Comon, and Kallel 2006] where K_{it} is the number of algorithm iterations. CCA can be efficiently implemented by singular value decomposition with complexity $\mathcal{O}(M^2 N)$ and typically outperforms ICA [Borga and Knutsson 2001; Sweeney, McLoone, and Ward 2012].

We then compared the performances of practical implementations of the algorithms by measuring the CPU time required to analyse a 30-second EEG sample taken from the Physiobank dataset [Sweeney et al. 2012]. WQN and wavelet thresholding showed similar performances (2 ms and 3 ms respectively) and were superior to EMD-based and SUBAR algorithms (314 ms for EMD-CCA, 419 ms for EMD-ICA, and 31 s for SUBAR).

In conclusion, we found that the computational cost of the WQN method is comparable to that of wavelet thresholding and that both are significantly faster than EMD and surrogate-based methods. The computational complexity and the CPU execution time for each algorithm are summarized in Table 6.2.

6.4.3 Application to anaesthesia monitoring

We present an application of WQN to the monitoring of general anaesthesia. In this context, spectral measures of the EEG such as the spectral edge frequency (SEF) or the alpha-to-delta ratio (ADR) are routinely used by anaesthesiologists to control the depth of anaesthesia. Artefacts caused by inadvertent motion of the electrodes can alter the signal spectrum and make these measures unreliable (figure 6.5 A–B). Depending on the duration of the time window on which the spectrum is calculated, a single artefact can affect the output of the monitoring for up to one minute. In figure 6.5 we present an example of the EEG spectrogram of a patient during the maintenance phase of general anaesthesia, showing how artefacts (purple arrows) can negatively affect the calculation of the SEF₉₅ (figure 6.5 B). WQN can be used to recover a corrected spectrum (figure 6.5 C) and to obtain a more robust estimation of spectral measures such as the SEF, improving the reliability of anaesthesia monitoring.

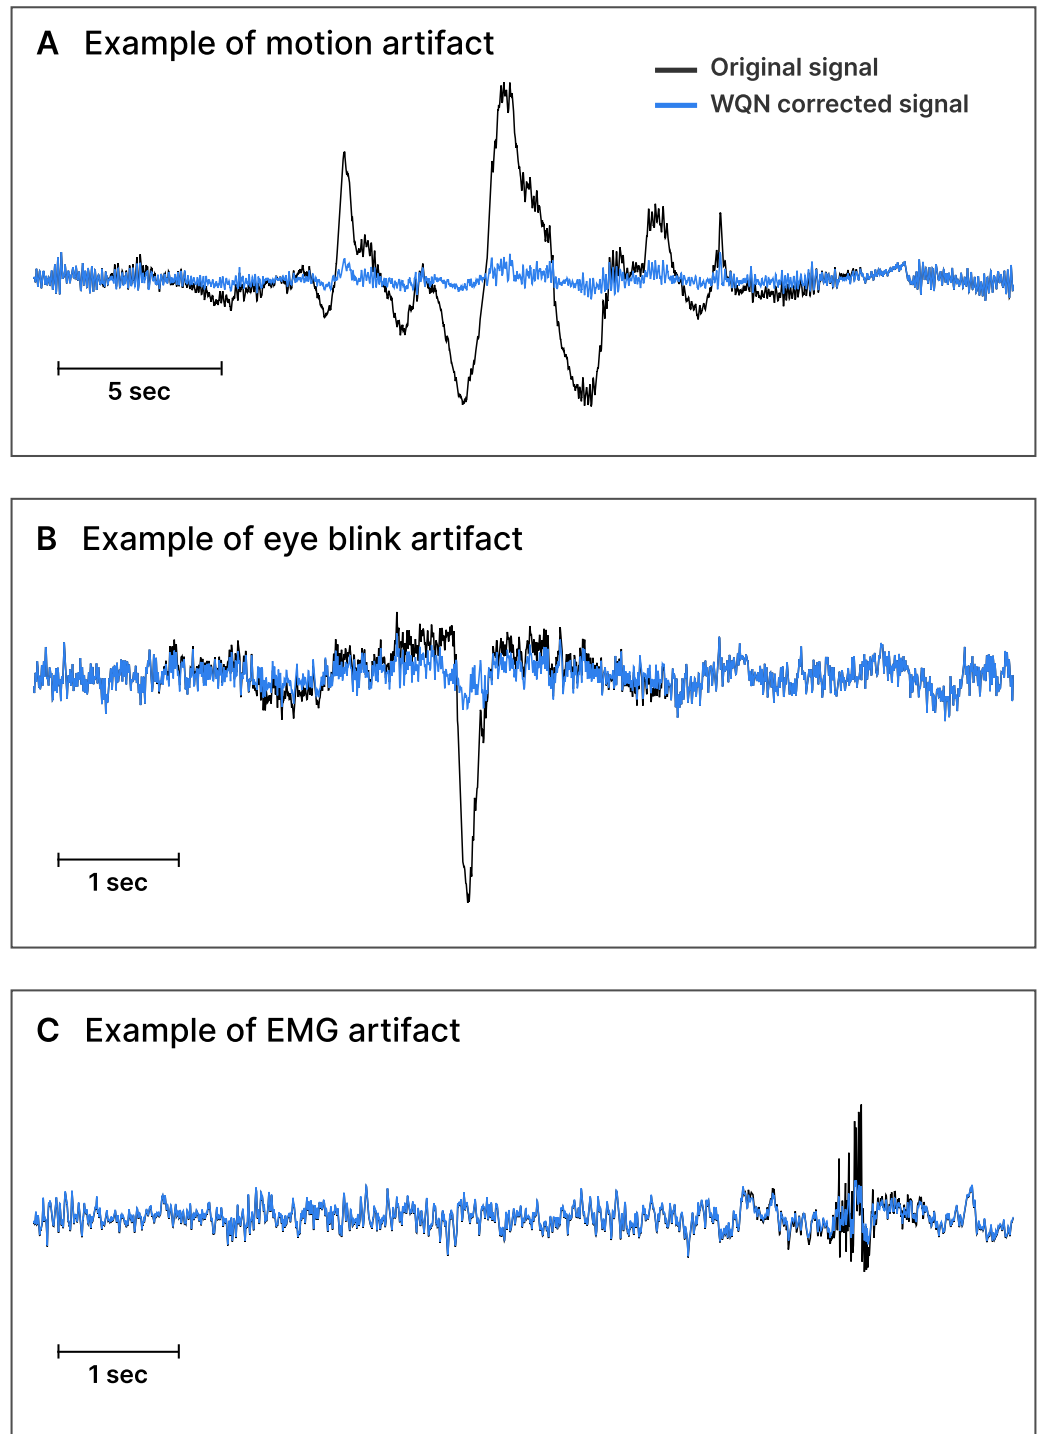


Figure 6.4 Example of artefact removal via WQN (original signal in black, corrected signal in blue). **A** Motion artefact. **B** Eye blink artefact. **C** Muscular artefact (EMG). (The samples were obtained from datasets [Agarwal and Sivakumar 2019; Cattan, Rodrigues, and Congedo 2018; Sweeney et al. 2012].)

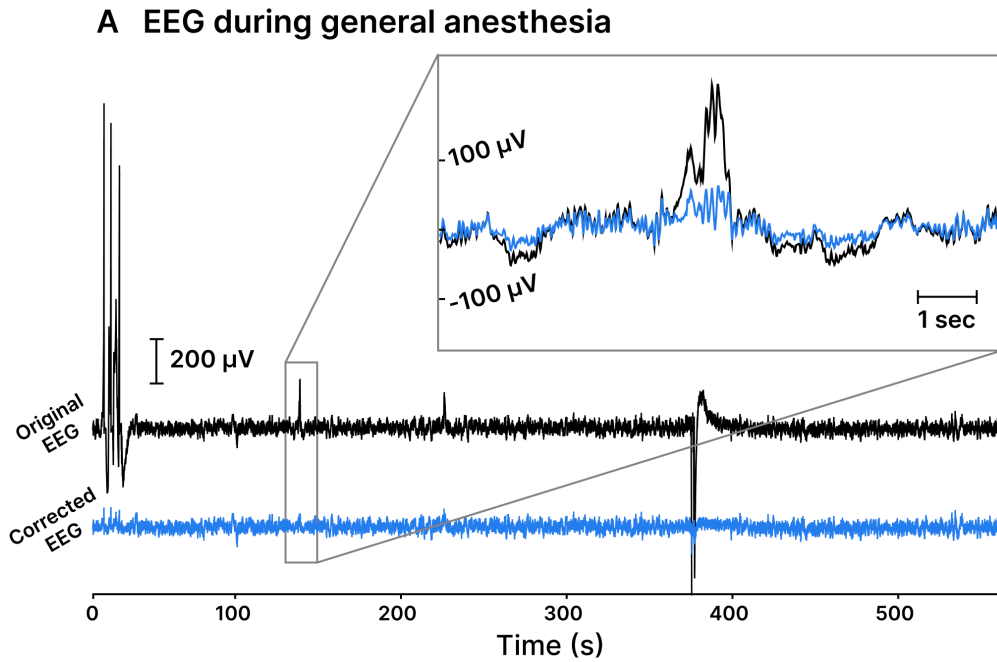
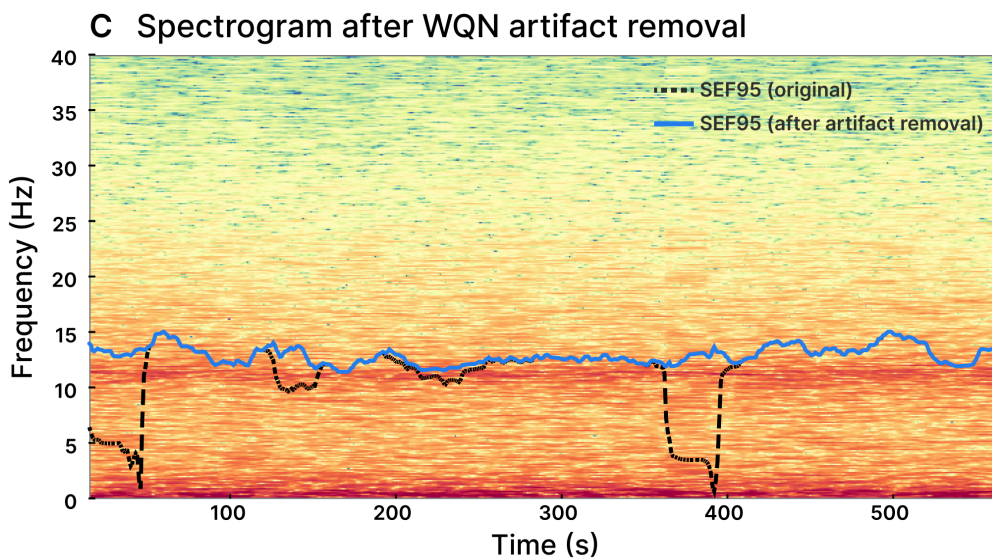



Figure 6.5 Example of artefact removal for EEG during anaesthesia. **A** Original EEG (black line) and WQN-corrected EEG (blue line) of a patient under general anaesthesia. **B** Spectrogram and spectral edge frequency at 95% (SEF95, black line). Artefacts are visibly contaminating the spectrum (purple arrows) and corrupting the SEF95 calculation. **C** Spectrogram and SEF95 after artefacts were removed using wavelet quantile normalization. The EEG signal was extracted from the VitalDB dataset [Lee and Jung 2018].



6.4.4 Limitations of the proposed method

The WQN method assumes gradual change of the brain activity, suppressing patterns which are associated to abrupt power changes. While this proved to be effective for several types of artefacts (table 6.1), the assumption cannot be extended to contexts where artefacts are associated to other transient activity. An example is represented by epileptic seizures which provoke sudden changes in the EEG power due to abnormal brain activity and are simultaneously associated to muscular or movement artefacts. Application of WQN in these cases would attenuate the epileptic patterns together with the artefactual components, potentially causing loss of important clinical information. Overcoming this limitation would require a spatiotemporal model to separate abnormal patterns of epileptic and artefactual origin [Hartmann et al. 2014], an issue closely related to the detection and classification of epileptic seizures.

6.5 Conclusions

In this work we presented and validated an artefact removal algorithm which can operate on single-channel EEG. The proposed WQN method is data-driven and requires no auxiliary input, parameter tuning or any human intervention. In our tests, WQN consistently outperformed comparable algorithms when applied on artefacts of different nature (table 6.1, figures 6.3 and 6.4) and for different signal-to-noise ratios. Moreover, we showed that the proposed algorithm has a similar computational cost to wavelet thresholding algorithms and that its performances are compatible with real-time applications. The characteristics of the method make it particularly suited for anaesthesia monitoring and other special care unit applications where only few electrodes are available and analysis has to be performed in real-time.

We presented an example of application to general anaesthesia (figure 6.5), showing how the WQN method could be used to improve the reliability of EEG monitoring. Future work includes the development and integration of an artefact detection algorithm, making WQN a fully automated technique to reduce artefacts in unsupervised EEG devices.

6.A Probability transformation

Consider a random variable X , with probability density function f_X , i.e. $X \sim f_X$. We would like to generate random variables distributed as f_Y , i.e. we want to find a monotonic transformation T such that $Y = T(X) \sim f_Y$. We obtain T in two steps. First, using the probability integral transform, we map X to a uniform random variable $U = F_X(X) \sim \text{Uniform}(0, 1)$, where F_X is the cumulative density function of X . Then, with the inverse integral transform we map U to Y , obtaining $Y = F_Y^{-1}(U) \sim f_Y$. Combining the two steps, we finally define

$$T = F_Y^{-1} \circ F_X. \quad (6.20)$$

The validity of this result can be easily verified by checking that the cumulative density function of $T(X)$ is indeed F_Y :

$$\begin{aligned} \mathbb{P}\{T(X) \leq y\} &= \mathbb{P}\{X \leq T^{-1}(y)\} = \mathbb{P}\{X \leq F_X^{-1}(F_Y(y))\} \\ &= F_X(F_X^{-1}(F_Y(y))) = F_Y(y). \end{aligned}$$

CHAPTER 7

Properties of the WQN algorithm

This last chapter supplements the previous one by providing some hints about the future direction of our work the on wavelet quantile normalization (WQN) algorithm. The content has been readapted from a conference proceedings proposal I have recently submitted in collaboration with Stéphane Jaffard (LAMA, Université Paris-Est Créteil) and David Holcman.¹ For this reason, there is a certain amount of repetition (e.g. in the definition of the WQN algorithm). The reader can directly skip to section 7.3, but I maintained the introductory section for completeness and coherence of notation. While the WQN method was developed empirically, here we try to develop a more precise intuition of why the algorithm works, and what differentiates it from traditional wavelet thresholding.

This chapter was adapted from “WQN-transport algorithm to adaptively correct artifacts in the EEG signal”, by Matteo Dora, Stéphane Jaffard, and David Holcman, submitted to the 30th European Signal Processing Conference (EUSIPCO 2022).

7.1 Introduction

Context. Real-time artefact removal in brain time series such as electroencephalogram (electroencephalography (EEG)) poses a general challenge to signal processing: how to suppress transient large amplitude artefacts without irreparably altering the statistical properties of the underlying signal? This situation is in contrast with traditional denoising approaches, which are designed to remove small additive noise which contaminates the signal. For this latter class of problems, wavelet-based methods [Averbuch, Lazar, and Israeli 1996; Buades, Coll, and Morel 2010; Donoho et al. 1995; Jaffard, Meyer, and Ryan 2001; Meyer 1992] have been shown to be very effective at identifying multiscale time-frequency patterns, allowing to separate signal from noise. The wavelet transform (WT) is now commonly used in many applications ranging from images [Averbuch, Lazar, and Israeli 1996], audio [Omer and Torrèsani 2017], to physiological signals [Ciuciu et al. 2008; Le Van Quyen and Bragin 2007; Lina et al. 2014; Unser and Aldroubi 1996; Worrell et al. 2012].

¹I would like to thank Stéphane Jaffard for supporting us with his profound expertise, and for revealing some of the mathematical details of the WQN algorithm.

During general anaesthesia, real-time monitoring of the EEG signal provides continuous feedback about the depth of anaesthesia, allowing to control the anaesthetic dose required to keep the patient in a safe sleep state. Anaesthesia monitoring machines are generally based on spectral features computed on time windows of 20–40 s [Fahy and Chau 2018]. When artefacts are detected, the output is suspended, preventing possible incorrect feedback but at the same time removing fundamental information about the state of the patient. Being able to continuously suppress transient artefacts, preserving the statistical properties of the signal, would thus allow a smoother monitoring.

Related works. Wavelet methods based on thresholding of small wavelet coefficients have been classically used to remove small amplitude noise with low computational cost. However, this approach requires to define the statistical properties of the signals in order to obtain a clear separation of the signal from noise [Urigüen and Garcia-Zapirain 2015]. Various methods have been proposed to optimally define threshold values based on these statistical properties [Donoho and Johnstone 1998; Donoho et al. 1995], focusing on signals contaminated by small additive white noise. The wavelet thresholding approach has been adapted to remove artefacts from EEG [Chavez et al. 2018; Inuso et al. 2007; Krishnaveni et al. 2006] but with opposite assumptions: the physiological signal is represented by a small-amplitude, stationary process which is corrupted by large and isolated artefacts. Contrary to traditional wavelet thresholding, artefact removal is obtained by suppressing a small number of *large* wavelet coefficients. Comparing the statistical properties of wavelet thresholding and WQN will allow to better understand how this latter algorithm works.

Goal, contribution, and outline. In this chapter, we analyse the WQN algorithm that we recently introduced [Dora and Holcman 2022] to remove artefacts from single-channel EEGs for real-time clinical monitoring. After having briefly recalled the WQN algorithm (in section 7.2), in section 7.3 we present wavelet coefficients statistics for two classes of EEG artefacts and compare them to those of physiological EEG. Then, in section 7.4, we compare the WQN to wavelet thresholding methods, showing how the latter can cause significant alteration of the coefficient statistics when they are employed backwards by suppressing large coefficients. In section 7.5 we study the regularity properties of the WQN algorithm, showing that it cannot introduce unwanted singularities in the signal. Finally, in section 7.6, we discuss the application of WQN to remove artefact from EEG in the operating room. Uninterrupted monitoring of these statistics can provide a continuous guidance for anaesthetists in taking the relevant decisions.

7.2 WQN algorithm

The wavelet quantile normalization (wQN) transport algorithm [Dora and Holcman 2022] allows eliminating transient artefacts [Schomer and Da Silva 2012; Tatum et al. 2011] from single-channel EEGs. The method consists in normalizing the wavelet coefficients of artefacted intervals using temporally adjacent uncontaminated signal as a reference. The artefacted and reference signals are decomposed with an M -level discrete wavelet transform (DWT) defined as [Daubechies 1992]:

$$x(t) = \sum_{m=1}^M \sum_n c_{m,n} \psi_{m,n}(t) + \sum_n c_{M+1,n} \varphi_{M,n}(t), \quad (7.1)$$

where $\psi_{m,n}(t) = 2^{-m/2} \psi(2^{-m}t - n)$, $\varphi_{m,n}(t) = 2^{-m/2} \varphi(2^{-m}t - n)$, represent wavelet and scaling functions, and $c_{M+1,n} = \langle x, \varphi_{M,n} \rangle$, $c_{m \leq M,n} = \langle x, \psi_{m,n} \rangle$ are the approximation and detail coefficients.

For each level m , the coefficients $c_m^{(\text{art})}$, $c_m^{(\text{ref})}$ are used to compute the empirical cumulative density functions (CDF) $F_m^{(\text{ref})}$, $F_m^{(\text{art})}$ of the coefficients amplitude for artefacted and reference signal respectively:

$$F_m(x) = \frac{1}{N_m} \sum_{n=1}^{N_m} 1_{|c_{m,n}| < x}, \quad (7.2)$$

where N_m indicates the number of coefficients and $1_{|c| < x}$ is the indicator function which takes value 1 if $|c| < x$ and 0 otherwise. The wavelet coefficients c_m^{art} are transformed such that the distribution of their amplitude matches that of the reference segment, via the mapping T_m defined as

$$T_m(x) = F_m^{(\text{ref})^{-1}} \left(F_m^{(\text{art})}(x) \right), \quad (7.3)$$

where $F_m^{(\text{ref})^{-1}}$ indicates the generalized inverse of $F_m^{(\text{ref})}$ (in the sense of completed graphs for discontinuous increasing functions, see figure 7.1). This mapping transports the amplitude distribution of $c_m^{(\text{art})}$ to the distribution of the reference signal. Finally, the normalization function

$$\lambda_m(c) = \text{sgn}(c) \cdot \min \{ |c|, T_m(|c|) \}, \quad (7.4)$$

maps a coefficient c from $c_m^{(\text{art})}$ to its possibly attenuated value. Equation 7.4 ensures that the norm of wavelet coefficients is never increased, a key requirement to guarantee the regularity of the algorithm as we will see in Sec. 7.5. The corrected coefficients are defined by

$$c_{m,n}^{(\text{corr})} = \lambda_m \left(c_{m,n}^{(\text{art})} \right), \quad m = 1, \dots, M + 1. \quad (7.5)$$

The corrected version of the artefacted segment is obtained by inverting the DWT using the corrected coefficients

$$x^{(\text{corr})}(t) = \sum_{m=1}^M \sum_n c_{m,n}^{(\text{corr})} \psi_{m,n}(t) + \sum_n c_{M+1,n}^{(\text{corr})} \varphi_{M,n}(t). \quad (7.6)$$

To illustrate how the wQN-algorithm performs, we added an electrooculogram (EOG) artefact to an EEG (figure 7.1 A) showing how wavelet coefficients are transported from the artefacted distribution to the reference (figure 7.1 B). The single-level coefficient mapping (figure 7.1 B) illustrates how the algorithm adapts to different scales.

7.3 Statistics of EEG and artefact signals


We further studied the distributions of the wavelet coefficients of the EEG signal and compared them to artefacts generated by eye movement (EOG) and muscular activity (EMG). We decomposed 2-second signals from the Denoise-Net dataset [Zhang et al. 2020] using a symlet wavelet with 5 vanishing moments to obtain the coefficient distributions for each scale (Fig. 7.2). Interestingly, the wavelet coefficients of the unperturbed EEG can be well described by generalized Gaussian distributions ($f(x) = \frac{\beta}{2\alpha\Gamma(1/\beta)} \exp[-(|x|/\alpha)^\beta]$) with shape parameter β between 1.5 and 2 (i.e. a standard Gaussian). This is in contrast with the coefficient distribution of the artefacted epochs. In EOG, non-zero coefficients are concentrated in the low frequency scale (level c_5); their distribution is peaked around small coefficients and shows a clear deviation from a generalized Gaussian due to skewness (Fig. 7.2, second column). In EMG, non-zero coefficients span several scales (Fig. 7.2, third column), similarly to the EEG, but they have a peaked distribution (with β close to 1, i.e. Laplace distribution). In summary, when the wQN algorithm is applied on EEG signals perturbed by additive artefacts, it will transport the coefficients from the EEG + EOG/EMG into an EEG-like distribution, allowing to restore the right statistics for the wavelet coefficients.

We conclude that the wQN algorithm can map artefacted signals to the smooth EEG distribution and thus remove possible jumps in the energy introduced by the artefacts, which are the cause of the monitoring issues we discuss below.

7.4 WQN versus wavelet thresholding

We compared wQN algorithm with the wavelet thresholding methods that were initially introduced to denoise signals without smoothing sharp patterns. In this context, the noisy component is assumed to have small wavelet coefficients, sharing the same statistical properties at each scale. The algorithm is efficient if the statistical properties of the signal to be recovered strongly differ from the noise e.g. if the wavelet coefficients of the signal form a *sparse sequence* (most of them almost vanish), and the non-zero coefficients are large. This situation is actually the opposite to the one considered in EEG artefact removal, where the artefacts to be eliminated have a sparse signature while the signal to be recovered presents the homogeneous statistical properties of such a noise. However, wavelet thresholding and wavelet shrinkage also have been used in such contexts [Chavez et al. 2018; Inuso et al. 2007; Krishnaveni et al. 2006; Sweeney, Ward, and McLoone 2012] (but,

A. WQN-transport correction of EOG artifact



B. Transport of wavelet coefficients




Figure 7.1 WQN correction of an EOG artefact. **A** Reconstruction of the original signal based on WQN (blue) from the artefacted signal (orange), compared to the original signal (black). **B** Effect of the coefficient normalization at scales c_4 and c_5 , showing the empirical CDF transport (left) and the mapping of the wavelet coefficients (right). The algorithm adaptively attenuates coefficients at the relevant scale c_5 , leaving c_4 mostly unmodified, thus demonstrating the relevance of a scale by scale procedure.

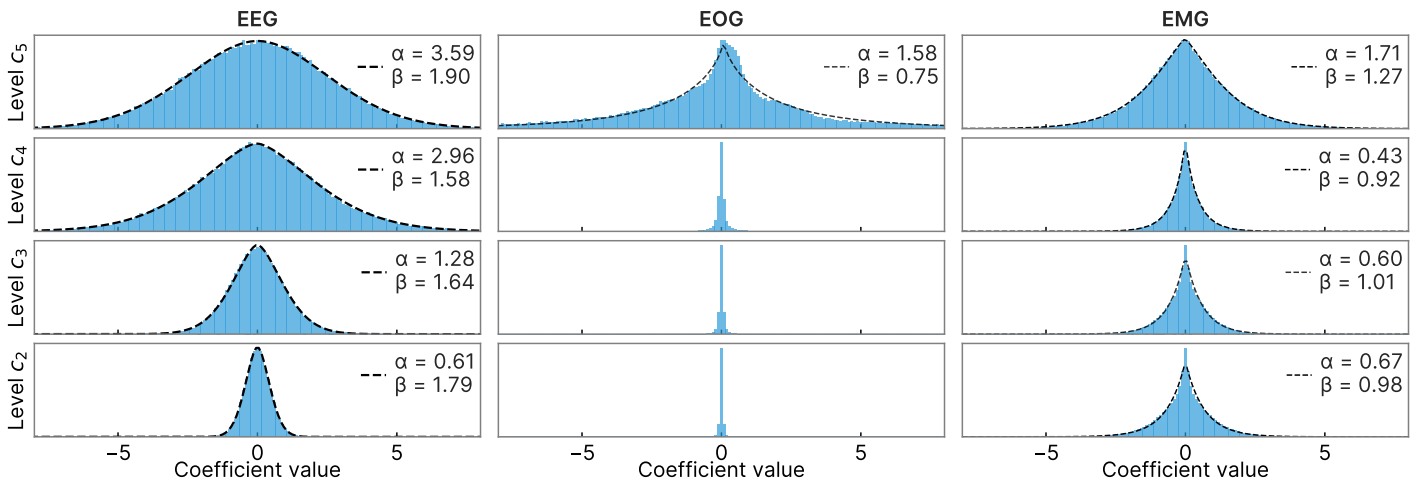


Figure 7.2 Distribution of the wavelet coefficients for EEG, EOG, and electromyogram (EMG). The wavelet distribution for the EEG signal can be well approximated by a centred generalized Gaussian distribution (dashed line), fitted by maximum-likelihood [Varanasi and Aazhang 1989]. The statistics were computed on EEG, EOG, and EMG recordings taken from the Denoise-Net dataset [Zhang et al. 2020]. Before computing the wavelet transform, each signal was normalized to have standard deviation equal to 1.

once the splitting has been performed, the noisy part is kept instead of the sparse one); it is therefore enlightening to compare their performance with the wqn algorithm. One drawback of both of these algorithms is that they are *local in the wavelet domain*, i.e. each wavelet coefficient is modified independently of the other ones, and therefore, they do not preserve the statistics of wavelet coefficients.

Non-local methods have been proposed (see e.g. block thresholding methods [Antoniadis 2007]), but they have the drawback of not restoring the correct coefficient statistics. This phenomenon has no negative impact when wavelet thresholding and wavelet shrinkage are used for their initial purpose, i.e. to restore the sparse part of the signal, but it becomes a major drawback when it is used for the opposite purpose of restoring the noisy component; for instance, on the specific position where the eliminated artefact was localized in time, the wavelet coefficients are set to 0, leading to inhomogeneities in the restored signal; one of the purposes of the wqn algorithm is to circumvent this drawback by restoring everywhere the correct anticipated statistics of the wavelet coefficients. As a consequence, it is not local in the wavelet domain: the value attributed to a coefficient depends on the whole statistic of coefficients at a given scale, and therefore the analysis of the regularity properties of the algorithm is more involved than for wavelet thresholding and wavelet shrinkage.

At this stage, we decided to compare the wqn algorithm with the wavelet thresholding methods. To this aim, we consider hard and soft thresholding on the wavelet

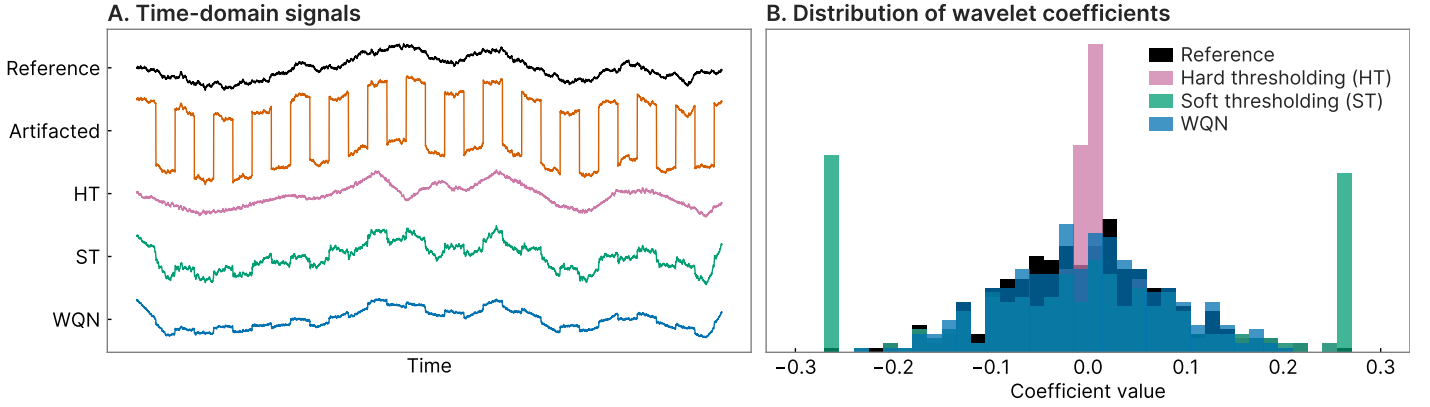


Figure 7.3 Example of artefact removal. **A** Comparison of wavelet hard-thresholding (HT), soft-thresholding (ST), and wqn on a signal that has been corrupted by adding a square wave. The reference signal is a realization of a Brownian motion process. **B** Distribution of the wavelet coefficients c_5 . The wavelet transport algorithm restores the right statistics of the EEG data, in contradistinction with the hard and soft thresholding which generate spurious peaks in the histogram respectively at the origin and at the threshold values.

coefficients w , defined by

$$\lambda_{\text{hard}}(w) = w \mathbf{1}_{|w| < \theta} \quad (7.7)$$

$$\lambda_{\text{soft}}(w) = \text{sgn}(w) \min(|w|, \theta) \quad (7.8)$$

where θ is the threshold and $\mathbf{1}$ is the indicator function. To emulate a stationary signal we use a Brownian motion process. The distribution of the wavelet coefficients is a centred Gaussian of variance 2^{5m} and with short-range correlations (see [Flandrin 1992] for explicit values). We added a square wave to simulate an artefact. We then applied the hard, soft thresholding and the wqn algorithm (figure 7.3A): while the reference and the signal restored by wqn have similar distributions, hard thresholding projects most of the large coefficients to zero, creating a peak at 0 in the distribution (Fig. 7.3B, pink), while the soft thresholding replaced them with the threshold value, creating two peaks at values $\pm\theta$ (figure 7.3 B, green). We conclude that the presence of various jump discontinuities leads to a large distribution of wavelet coefficients with high values. HT and ST algorithms with a threshold separating well the distribution of the artefact from the Brownian motion do not preserve the coefficient distribution, contrary to the wqn algorithm.

The empirical difference between the artefacted and the reference wavelet coefficient distribution seems sufficient to guarantee a good performance of the wqn algorithm: indeed the artefact can usually be decomposed into a sparse set of wavelet coefficients, some with large values, associated with fast jumps or/and large fast oscillation, leading to a long tail in the appropriate wavelet level distribution. The large coefficients are then mapped back into the distribution of the reference, possibly increasing all coefficients by a small value.

7.5 Regularity properties of WQN

A key property of wavelet bases is that, unlike the trigonometric system or other classical bases, most classical functional spaces used in mathematical analysis, such as Sobolev spaces $H^{s,p}$ for $1 < p < \infty$ or Besov spaces $B_p^{s,q}$ for $0 < p, q \leq \infty$, can be characterized by conditions on their coefficients [Meyer 1992]; this is referred to as the *unconditional basis property* which plays a key role in statistics, see [Donoho et al. 1995]. These spaces have a wavelet characterization which bears on the moduli of the wavelet coefficients, and which is an increasing function of each of these moduli. As an example, $f \in B_p^{s,q}$ if its wavelet coefficients satisfy the condition

$$\left(\sum_n (2^{(1/p-1/2-s)m} |c_{m,n}|)^p \right)^{1/p} \in l^q. \quad (7.9)$$

Since, by construction, the wavelet-transport algorithm is *wavelet decreasing* (i.e. it does not increase the size of the wavelet coefficients, see equation 7.4), the following regularity property, also shared by the wavelet thresholding algorithms is satisfied: For $0 < p, q \leq \infty$, and for any $s \in \mathbb{R}$, it maps functions of a Besov or Sobolev space into the same space. Note that this boundedness property also holds in the S^v spaces, which were introduced in [Jaffard 2004], and are defined by conditions on the distributions of wavelet coefficients at each scale, and therefore supply a functional setting particularly well-fitted to the study of the WQN-wavelet transport algorithm. The same remark also applies to pointwise singularities, which can also be characterized by conditions bearing on the moduli of the wavelet coefficients, see e.g. [Jaffard 2006]. To conclude, the regularity property above guarantees that the WQN algorithm does not introduce unwanted singularities, but on the contrary smooths a signal contaminated by a sparse, highly non-Gaussian artefact.

7.6 WQN to remove artefacts in real-time monitoring

We discuss now a direct application of the WQN algorithm to remove artefacts located in an EEG signal. An artefact can last few seconds, but can contaminate computations in a longer time interval of tens of seconds. This is usually the case when the spectrogram is computed over a sliding window with a typical width of 20 to 40 seconds. The artefact is visible on the spectrogram, with a large spread of the frequencies (Fig. 7.4A), compromising the quality of estimations. In some cases, algorithms to evaluate the depth of anaesthesia such as the bispectral index (BIS) simply stop processing the data, leading to blank epochs. To illustrate the performance of WQN algorithm, we present artefacts that are corrupting the spectrogram of an EEG recorded during general anaesthesia and consequently the computations of parameters measuring the depth of sedation (Fig. 7.4A). After applying the WQN algorithm, the artefacted regions present in the EEG and the spectrogram can be restored (Fig. 7.4B). Indeed, the spectrogram now clearly shows a continuous alpha band around 10 Hz.




Figure 7.4 Example of restoring an EEG signal during general anaesthesia. **A** Spectrogram and EEG with artefact (shaded intervals). Artefacts are corrupting the spectrogram calculation and consequently the computations of parameters measuring the depth of anaesthesia. **B** Spectrogram and EEG after applying the wqn algorithm that has restored the artefacted regions. The spectrogram now clearly shows a continuous alpha band around 10 Hz.

7.7 Conclusions

The WQN algorithm allows replacing an artefacted epoch by a corrected signal, where the wavelet coefficient distribution is similar to the reference EEG. In contrast with hard and soft thresholding methods, this algorithm does not require a threshold to be chosen. In addition, it performs well on EEG signal recorded during anaesthesia as there is usually an energy gap between the artefacts (characterized by large jumps and sudden fast oscillations) and the unperturbed EEG. As shown in the present work, the corrected signal is smoother than the artefact. This property is relevant because it allows direct automated computations from on the corrected signal that can be used to measure the depth of anaesthesia. Clinical brain monitors are often programmed to stop operating following the detection of an artefact, as the artefact can undermine the computation of relevant parameters in the neighbouring signal. We conclude that the WQN procedure could allow to improve the continuity of monitoring and real-time computation of the parameters associated to depth of sedation for general anaesthesia.

Concluding remarks

My colleague, Fréd, likes to say that the conclusion section is just the same as the introduction—but in the past tense. I will thus avoid the pointless exercise of repeating what has already been said (an overview of the main results can be found at page x) and limit myself to a few very general remarks. What we have seen in this thesis is how theoretical tools can complement experimental research. In some cases, these tools improve our capacities by allowing us to see a little further ahead, in the fog where we could not see before. In others, they assume a more fundamental and unique role, in bringing new perspectives that could not have been revealed by any other mean. One surprising—albeit well known—aspect of all this is how diverse and apparently unrelated fields of applied research could be traced back to a very limited family of theoretical methods. As we have seen, a few mathematical descriptions recurrently emerged when describing different biological systems operating at very different scales. Obviously, not much can be said about the origins of this effectiveness and apparent universality. My idea, however, is that this has to do with the ease of representation that the mathematical language provides. The advantage of the theoretical world is that one does not have to account for hard reality, and can instead work with malleable representations of the phenomena under examination. Moving from one representation to the other takes very little effort, usually in the form of few computations to get, for example, to a Fourier or wavelet basis. This flexibility goes along with the fact that there always seems to be a representation that makes our problems simpler. Moreover, some of these representations can apparently seize profound structural patterns which are not limited to the single phenomenon, but to a more universal class of systems. The desire to understand representations at a more fundamental level is what I think will guide me in my future route.

*Why is my verse so barren of new pride,
So far from variation or quick change?
Why with the time do I not glance aside
To new-found methods, and to compounds strange?
Why write I still all one, ever the same,
And keep invention in a noted weed,
That every word doth almost tell my name,
Showing their birth, and where they did proceed?
O know, sweet love, I always write of you,
And you and love are still my argument,
So all my best is dressing old words new,
Spending again what is already spent:
 For as the sun is daily new and old,
 So is my love still telling what is told.*

— William Shakespeare, *Sonnet 76*

Bibliography

- Agarwal, Mohit and Raghupathy Sivakumar (2019). “Blink: A Fully Automated Un-supervised Algorithm for Eye-Blink Detection in Eeg Signals”. In: *2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, pp. 1113–1121.
- Alberts, Bruce et al. (2013). *Essential Cell Biology*. Garland Science.
- (2015). *Essential Cell Biology*. Garland Science.
- Alim, Karen et al. (2013). *Random Network Peristalsis in Physarum Polycephalum Organizes Fluid Flows across an Individual*. In: *Proceedings of the National Academy of Sciences* 110.33, pp. 13306–13311.
- André-Obadia, Nathalie (2018). *Recommendations for the Use of Electroencephalography and Evoked Potentials in Comatose Patients*. In: *Clinical Neurophysiology*, p. 27.
- Angermann, Bastian R et al. (2012). *Computational Modeling of Cellular Signaling Processes Embedded into Dynamic Spatial Contexts*. In: *Nature methods* 9.3, pp. 283–289.
- Antoniadis, Anestis (Jan. 2007). *Wavelet Methods in Statistics: Some Recent Developments and Their Applications*. In: *Statistics Surveys* 1 (none), pp. 16–55. ISSN: 1935-7516. DOI: 10.1214/07-SS014.
- Averbuch, Amir, Danny Lazar, and Moshe Israeli (1996). *Image Compression Using Wavelet Transform and Multiresolution Decomposition*. In: *IEEE Transactions on Image Processing* 5.1, pp. 4–15.
- Averbuch, Amir Z, Pekka Neittaanmäki, and Valery A Zheludev (2019). “Acoustic Detection of Moving Vehicles”. In: *Spline and Spline Wavelet Methods with Applications to Signal and Image Processing*. Springer, pp. 219–239.
- Azzerboni, Bruno et al. (2002). “A New Approach to Detection of Muscle Activation by Independent Component Analysis and Wavelet Transform”. In: *Italian Workshop on Neural Nets*. Springer, pp. 109–116.
- Bajaj, Nikesh et al. (Jan. 2020). *Automatic and Tunable Algorithm for EEG Artifact Removal Using Wavelet Decomposition with Applications in Predictive Modeling during Auditory Tasks*. In: *Biomedical Signal Processing and Control* 55, p. 101624. ISSN: 17468094. DOI: 10.1016/j.bspc.2019.101624.
- Bando, Yuki et al. (Sept. 12, 2019). *Genetic Voltage Indicators*. In: *BMC biology* 17.1, p. 71. ISSN: 1741-7007. DOI: 10.1186/s12915-019-0682-0.

- Basnayake, K., Z. Schuss, and D. Holcman (Sept. 15, 2018). *Asymptotic Formulas for Extreme Statistics of Escape Times in 1, 2 and 3-Dimensions*. In: *Journal of Nonlinear Science*. ISSN: 1432-1467. DOI: 10.1007/s00332-018-9493-7.
- Basnayake, K et al. (2018). *Extreme Narrow Escape: Shortest Paths for the First Particles among n to Reach a Target Window*. In: *Physics Letters A*.
- Basnayake, Kanishka and David Holcman (2019). *Fastest among Equals: A Novel Paradigm in Biology. Reply to Comments: Redundancy Principle and the Role of Extreme Statistics in Molecular and Cellular Biology*. In: *Physics of life reviews* 28, pp. 96–99.
- Batson, Joshua and Loic Royer (2019). “Noise2Self: Blind Denoising by Self-Supervision”. In: *International Conference on Machine Learning*, pp. 524–533.
- Baum, Buzz and David A Baum (2020). *The Merger That Made Us*. In: *BMC biology* 18.1, pp. 1–4.
- Bernasconi, Riccardo et al. (2010). *Stringent Requirement for HRD1, SEL1L, and OS-9/XTP3-B for Disposal of ERAD-LS Substrates*. In: *Journal of Cell Biology* 188.2, pp. 223–235.
- Berridge, Michael J (2002). *The Endoplasmic Reticulum: A Multifunctional Signaling Organelle*. In: *Cell calcium* 32.5-6, pp. 235–249.
- Bloxham, William et al. (Dec. 7, 2021). *Linearly Polarized Excitation Enhances Signals from Fluorescent Voltage Indicators*. In: *Biophysical Journal* 120.23, pp. 5333–5342. ISSN: 1542-0086. DOI: 10.1016/j.bpj.2021.10.028.
- Bollobas, Bela, Bela Bollobás, and Oliver Riordan (Sept. 21, 2006). *Percolation*. Cambridge University Press. 334 pp. ISBN: 978-0-521-87232-4.
- Bonifacino, Juan S et al. (1989). *Pre-Golgi Degradation of Newly Synthesized T-cell Antigen Receptor Chains: Intrinsic Sensitivity and the Role of Subunit Assembly*. In: *The Journal of cell biology* 109.1, pp. 73–83.
- Borga, Magnus and Hans Knutsson (June 5, 2001). *A Canonical Correlation Approach to Blind Source Separation*. In: p. 12.
- Bray, Alan J, Satya N Majumdar, and Grégory Schehr (2013). *Persistence and First-Passage Properties in Nonequilibrium Systems*. In: *Advances in Physics* 62.3, pp. 225–361.
- Broadbent, Simon R and John M Hammersley (1957). “Percolation Processes: I. Crystals and Mazes”. In: *Mathematical Proceedings of the Cambridge Philosophical Society*. Vol. 53. 3. Cambridge University Press, pp. 629–641.
- Buades, A., B. Coll, and J. M. Morel (Jan. 2010). *Image Denoising Methods. A New Nonlocal Principle*. In: *SIAM Review* 52.1, pp. 113–147. ISSN: 0036-1445. DOI: 10.1137/090773908.
- Buchanan, E. Kelly et al. (July 17, 2018). *Penalized Matrix Decomposition for Denoising, Compression, and Improved Demixing of Functional Imaging Data*. arXiv: 1807.06203 [q-bio, stat].
- Budnik, Annika and David J Stephens (2009). *ER Exit Sites–Localization and Control of COPII Vesicle Formation*. In: *FEBS letters* 583.23, pp. 3796–3803.

- Bénichou, Olivier and R Voituriez (2014). *From First-Passage Times of Random Walks in Confinement to Geometry-Controlled Kinetics*. In: *Physics Reports* 539.4, pp. 225–284.
- Cai, Changjia et al. (Apr. 2021). *VolPy: Automated and Scalable Analysis Pipelines for Voltage Imaging Datasets*. In: *PLoS computational biology* 17.4, e1008806. ISSN: 1553-7358. DOI: 10.1371/journal.pcbi.1008806.
- Cattan, Grégoire, Pedro Luiz Coelho Rodrigues, and Marco Congedo (2018). *EEG Alpha Waves Dataset*. DOI: 10.5281/zenodo.2348891.
- Chavez, M. et al. (Mar. 2018). *Surrogate-Based Artifact Removal From Single-Channel EEG*. In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 26.3, pp. 540–550. ISSN: 1534-4320, 1558-0210. DOI: 10.1109/TNSRE.2018.2794184.
- Ciuciu, Philippe et al. (Dec. 2008). *Log Wavelet Leaders Cumulant Based Multifractal Analysis of EVI fMRI Time Series: Evidence of Scaling in Ongoing and Evoked Brain Activity*. In: *IEEE Journal of Selected Topics in Signal Processing* 2.6, pp. 929–943. ISSN: 1941-0484. DOI: 10.1109/JSTSP.2008.2006663.
- Coates, Steven et al. (2012). *Tele-EEG in the UK: A Report of over 1000 Patients*. In: *Journal of Telemedicine and Telecare* 18.5, pp. 243–246.
- Cole, Nelson B et al. (1996). *Diffusional Mobility of Golgi Proteins in Membranes of Living Cells*. In: *Science* 273.5276, pp. 797–801.
- Coombs, Daniel (2019). *First among Equals: Comment on “Redundancy Principle and the Role of Extreme Statistics in Molecular and Cellular Biology” by Z. Schuss, K. Basnayake and D. Holcman*. In: *Physics of life reviews*.
- Correa, A Garcés et al. (2007). “Artifact Removal from EEG Signals Using Adaptive Filters in Cascade”. In: *Journal of Physics: Conference Series*. Vol. 90. 1. IOP Publishing, p. 012081.
- Costantini, Lindsey M and Erik Lee Snapp (2013). *Fluorescent Proteins in Cellular Organelles: Serious Pitfalls and Some Solutions*. In: *DNA and cell biology* 32.11, pp. 622–627.
- Costantini, Lindsey M et al. (2015). *A Palette of Fluorescent Proteins Optimized for Diverse Cellular Environments*. In: *Nature communications* 6.1, pp. 1–13.
- Dabov, Kostadin, Alessandro Foi, and Karen Egiazarian (Sept. 2007). “Video Denoising by Sparse 3D Transform-Domain Collaborative Filtering”. In: *2007 15th European Signal Processing Conference*. 2007 15th European Signal Processing Conference, pp. 145–149.
- Dabov, Kostadin et al. (Aug. 2007). *Image Denoising by Sparse 3-D Transform-Domain Collaborative Filtering*. In: *IEEE Transactions on Image Processing* 16.8, pp. 2080–2095. ISSN: 1941-0042. DOI: 10.1109/TIP.2007.901238.
- Danielyan, Aram, Vladimir Katkovnik, and Karen Egiazarian (Apr. 2012). *BM3D Frames and Variational Image Deblurring*. In: *IEEE Transactions on Image Processing* 21.4, pp. 1715–1728. ISSN: 1057-7149, 1941-0042. DOI: 10.1109/TIP.2011.2176954. arXiv: 1106.6180.

- Daubechies, Ingrid (Jan. 1, 1992). *Ten Lectures on Wavelets*. CBMS-NSF Regional Conference Series in Applied Mathematics. Society for Industrial and Applied Mathematics. 369 pp. ISBN: 978-0-89871-274-2. DOI: 10.1137/1.9781611970104.
- Deneux, Thomas et al. (Nov. 2016). *Accurate Spike Estimation from Noisy Calcium Signals for Ultrafast Three-Dimensional Imaging of Large Neuronal Populations in Vivo*. In: *Nature Communications* 7.1, p. 12190. ISSN: 2041-1723. DOI: 10.1038/ncomms12190.
- Ding, Lingjun et al. (Jan. 26, 2022). *Juxtacellular Opto-Tagging of Hippocampal CA1 Neurons in Freely Moving Mice*. In: *eLife* 11, e71720. ISSN: 2050-084X. DOI: 10.7554/eLife.71720.
- Donoho, David L and Iain M Johnstone (1994a). “Threshold Selection for Wavelet Shrinkage of Noisy Data”. In: *Proceedings of 16th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. Vol. 1. IEEE, A24–A25.
- Donoho, David L. and Iain M. Johnstone (June 1, 1998). *Minimax Estimation via Wavelet Shrinkage*. In: *The Annals of Statistics* 26.3. ISSN: 0090-5364. DOI: 10.1214/aos/1024691081.
- Donoho, David L and Iain M Johnstone (1994b). *Ideal Spatial Adaptation by Wavelet Shrinkage*. In: *biometrika* 81.3, pp. 425–455.
- Donoho, David L et al. (1995). *Wavelet Shrinkage: Asymptopia?* In: *Journal of the Royal Statistical Society: Series B (Methodological)* 57.2, pp. 301–337.
- Dora, Matteo and David Holcman (2020). *Active Flow Network Generates Molecular Transport by Packets: Case of the Endoplasmic Reticulum*. In: *Proceedings of the Royal Society B* 287.1930, p. 20200493.
- (2022). *Adaptive Single-Channel EEG Artifact Removal With Applications to Clinical Monitoring*. In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 30, pp. 286–295. ISSN: 1558-0210. DOI: 10.1109/TNSRE.2022.3147072.
- Ehret, Thibaud and Pablo Arias (2020). *Implementation of the VBM3D Video Denoising Method and Some Variants*. In.
- Ellenberg, Jan, Jennifer Lippincott-Schwartz, and John F Presley (1999). *Dual-Colour Imaging with GFP Variants*. In: *Trends in cell biology* 9.2, pp. 52–56.
- Encell, Lance P et al. (2012). *Suppl 1: Development of a Dehalogenase-Based Protein Fusion Tag Capable of Rapid, Selective and Covalent Attachment to Customizable Ligands*. In: *Current chemical genomics* 6, p. 55.
- Esposito, Fabrizio et al. (2003). *Real-Time Independent Component Analysis of fMRI Time-Series*. In: *Neuroimage* 20.4, pp. 2209–2224.
- Fahy, Brenda G and Destiny F Chau (2018). *The Technology of Processed Electroencephalogram Monitoring Devices for Assessment of Depth of Anesthesia*. In: 126.1, p. 7.
- Farneböck, Gunnar (2003). “Two-Frame Motion Estimation Based on Polynomial Expansion”. In: *Image Analysis*. Ed. by Josef Bigun and Tomas Gustavsson. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, pp. 363–370. ISBN: 978-3-540-45103-7. DOI: 10.1007/3-540-45103-X_50.

- Feldt, Sarah, Paolo Bonifazi, and Rosa Cossart (May 2011). *Dissecting Functional Connectivity of Neuronal Microcircuits: Experimental and Theoretical Insights*. In: *Trends in Neurosciences* 34.5, pp. 225–236. ISSN: 1878-108X. DOI: 10.1016/j.tins.2011.02.007.
- Fiorillo, Luigi et al. (2019). *Automated Sleep Scoring: A Review of the Latest Approaches*. In: *Sleep medicine reviews* 48, p. 101204.
- Flandrin, P. (Mar. 1992). *Wavelet Analysis and Synthesis of Fractional Brownian Motion*. In: *IEEE Transactions on Information Theory* 38.2, pp. 910–917. ISSN: 1557-9654. DOI: 10.1109/18.119751.
- Flandrin, Patrick, Gabriel Rilling, and Paulo Goncalves (2004). *Empirical Mode Decomposition as a Filter Bank*. In: *IEEE signal processing letters* 11.2, pp. 112–114.
- Flandrin, Patrick et al. (2013). *Time-Frequency Analysis and Applications*. In: *IEEE signal processing magazine* 30.6, p. 19.
- Foi, Alessandro (Dec. 2009a). *Clipped Noisy Images: Heteroskedastic Modeling and Practical Denoising*. In: *Signal Processing* 89.12, pp. 2609–2629. DOI: 10.1016/j.sigpro.2009.04.035.
- (2009b). *Optimization of Variance-Stabilizing Transformations*. In: *Preprint*.
- Foi, Alessandro et al. (2008). *Practical Poissonian-Gaussian Noise Modeling and Fitting for Single-Image Raw-Data*. In: *IEEE Transactions on Image Processing* 17.10, pp. 1737–1754.
- Ford Jr, Lester Randolph and Delbert Ray Fulkerson (2015). *Flows in Networks*. Vol. 54. Princeton university press.
- Getreuer, Pascal (July 30, 2012). *Total Variation Deconvolution Using Split Bregman*. In: *Image Processing On Line* 2, pp. 158–174. ISSN: 2105-1232. DOI: 10.5201/ipol.2012.g-tvdc.
- Ghemrawi, Rose and Mostafa Khair (Aug. 25, 2020). *Endoplasmic Reticulum Stress and Unfolded Protein Response in Neurodegenerative Diseases*. In: *International Journal of Molecular Sciences* 21.17, p. 6127. ISSN: 1422-0067. DOI: 10.3390/ijms21176127. pmid: 32854418.
- Giovannucci, Andrea et al. (Jan. 17, 2019). *CaImAn an Open Source Tool for Scalable Calcium Imaging Data Analysis*. In: *eLife* 8, e38173. ISSN: 2050-084X. DOI: 10.7554/eLife.38173.
- Giovannucci, Andrea et al. (Oct. 4, 2021). *FIOLA: An Accelerated Pipeline for Fluorescence Imaging OnLine Analysis*. In Review.
- Goldberger, Ary L et al. (2000). *PhysioBank, PhysioToolkit, and PhysioNet: Components of a New Research Resource for Complex Physiologic Signals*. In: *circulation* 101.23, e215–e220.
- Goldstein, Tom and Stanley Osher (Jan. 2009). *The Split Bregman Method for L1-Regularized Problems*. In: *SIAM Journal on Imaging Sciences* 2.2, pp. 323–343. DOI: 10.1137/080725891.
- Goyal, Uma and Craig Blackstone (2013). *Untangling the Web: Mechanisms Underlying ER Network Formation*. In: *Biochimica et Biophysica Acta (BBA)-Molecular Cell Research* 1833.11, pp. 2492–2498.

- Grimm, Jonathan B et al. (2016). *Bright Photoactivatable Fluorophores for Single-Molecule Imaging*. In: *Nature methods* 13.12, pp. 985–988.
- Guarnieri, Roberto et al. (2018). *Online EEG Artifact Removal for BCI Applications by Adaptive Spatial Filtering*. In: *J. Neural Eng.*, p. 13.
- Hartmann, M.M. et al. (Nov. 2014). *PureEEG: Automatic EEG Artifact Removal for Epilepsy Monitoring*. In: *Neurophysiologie Clinique/Clinical Neurophysiology* 44.5, pp. 479–490. ISSN: 09877053. DOI: 10.1016/j.neucli.2014.09.001.
- Hashimoto, Shoko and Takaomi C. Saido (Apr. 25, 2018). *Critical Review: Involvement of Endoplasmic Reticulum Stress in the Aetiology of Alzheimer’s Disease*. In: *Open Biology* 8.4, p. 180024. ISSN: 2046-2441. DOI: 10.1098/rsob.180024. pmid: 29695619.
- “The Graphics Gems Series” (Jan. 1, 1994). In: *Graphics Gems*. Ed. by Paul S. Heckbert. Academic Press, p. ii. ISBN: 978-0-12-336156-1. DOI: 10.1016/B978-0-12-336156-1.50001-X.
- Heinrich, Larissa et al. (2021). *Whole-Cell Organelle Segmentation in Volume Electron Microscopy*. In: *Nature* 599.7883, pp. 141–146.
- Hirschberg, Koret et al. (1998). *Kinetic Analysis of Secretory Protein Traffic and Characterization of Golgi to Plasma Membrane Transport Intermediates in Living Cells*. In: *The Journal of cell biology* 143.6, pp. 1485–1503.
- Hoitzing, Hanne, Iain G Johnston, and Nick S Jones (2015). *What Is the Function of Mitochondrial Networks? A Theoretical Assessment of Hypotheses and Proposal for Future Research*. In: *BioEssays* 37.6, pp. 687–700.
- Holcman, David and Zeev Schuss (2015). *Stochastic Narrow Escape in Molecular and Cellular Biology: Analysis and Applications*. Springer.
- Holcman, David et al. (2018). *Single Particle Trajectories Reveal Active Endoplasmic Reticulum Luminal Flow*. In: *Nature cell biology* 20.10, pp. 1118–1125.
- Hozé, Nathanaël and David Holcman (2017). *Statistical Methods for Large Ensembles of Super-Resolution Stochastic Single Particle Trajectories in Cell Biology*. In: *Annual Review of Statistics and Its Application* 4.1, pp. 189–223. DOI: 10.1146/annurev-statistics-060116-054204.
- Inuso, Giuseppina et al. (Aug. 2007). “Wavelet-ICA Methodology for Efficient Artifact Removal from Electroencephalographic Recordings”. In: *2007 International Joint Conference on Neural Networks*. 2007 International Joint Conference on Neural Networks. Orlando, FL, USA: IEEE, pp. 1524–1529. ISBN: 978-1-4244-1379-9 978-1-4244-1380-5. DOI: 10.1109/IJCNN.2007.4371184.
- Islam, Md Kafiul, Parviz Ghorbanzadeh, and Amir Rastegarnia (2021). *Probability Mapping Based Artifact Detection and Removal from Single-Channel EEG Signals for Brain-Computer Interface Applications*. In: *Journal of Neuroscience Methods*, p. 109249.
- Islam, Md Kafiul, Amir Rastegarnia, and Zhi Yang (2016). *Methods for Artifact Detection and Removal from Scalp EEG: A Review*. In: *Neurophysiologie Clinique/Clinical Neurophysiology* 46.4-5, pp. 287–305.

- Jaffard, Stéphane (May 1, 2004). *Beyond Besov Spaces Part 1: Distributions of Wavelet Coefficients*. In: *Journal of Fourier Analysis and Applications* 10.3, pp. 221–246. ISSN: 1531-5851. DOI: 10.1007/s00041-004-0946-z.
- (2006). *Wavelet techniques for pointwise regularity*. In: *Annales de la Faculté des sciences de Toulouse : Mathématiques* 15.1, pp. 3–33. DOI: 10.5802/afst.1111.
- Jaffard, Stéphane, Yves Meyer, and Robert D Ryan (2001). *Wavelets: Tools for Science and Technology*. SIAM.
- Jakobs, Maximilian AH, Andrea Dimitracopoulos, and Kristian Franze (2019). *KymoButler, a Deep Learning Software for Automated Kymograph Analysis*. In: *Elife* 8, e42288.
- Jas, Mainak et al. (2017). *Autoreject: Automated Artifact Rejection for MEG and EEG Data*. In: *NeuroImage* 159, pp. 417–429.
- Johnstone, Iain M and Bernard W Silverman (1997). *Wavelet Threshold Estimators for Data with Correlated Noise*. In: *Journal of the royal statistical society: series B (statistical methodology)* 59.2, pp. 319–351.
- Jouhanneau, Jean-Sébastien and James F. A. Poulet (2019). *Multiple Two-Photon Targeted Whole-Cell Patch-Clamp Recordings from Monosynaptically Connected Neurons in Vivo*. In: *Frontiers in Synaptic Neuroscience* 11, p. 15. ISSN: 1663-3563. DOI: 10.3389/fnsyn.2019.00015.
- Kaberniuk, Andrii A et al. (2018). *moxMaple3: A Photoswitchable Fluorescent Protein for PALM and Protein Highlighting in Oxidizing Cellular Environments*. In: *Scientific reports* 8.1, pp. 1–10.
- Keemink, Sander W. et al. (Feb. 22, 2018). *FISSA: A Neuropil Decontamination Toolbox for Calcium Imaging Signals*. In: *Scientific Reports* 8.1, p. 3493. ISSN: 2045-2322. DOI: 10.1038/s41598-018-21640-2.
- Kingma, Diederik P. and Jimmy Ba (Jan. 29, 2017). *Adam: A Method for Stochastic Optimization*. arXiv: 1412.6980 [cs].
- Klados, Manousos A. and Panagiotis D. Bamidis (Sept. 2016). *A Semi-Simulated EEG/EOG Dataset for the Comparison of EOG Artifact Rejection Techniques*. In: *Data in Brief* 8, pp. 1004–1006. ISSN: 23523409. DOI: 10.1016/j.dib.2016.06.032.
- Klug, Marius and Klaus Gramann (2020). *Identifying Key Factors for Improving ICA-based Decomposition of EEG Data in Mobile and Stationary Experiments*. In: *European Journal of Neuroscience*.
- Krishnaveni, V et al. (Dec. 1, 2006). *Removal of Ocular Artifacts from EEG Using Adaptive Thresholding of Wavelet Coefficients*. In: *Journal of Neural Engineering* 3.4, pp. 338–346. ISSN: 1741-2560, 1741-2552. DOI: 10.1088/1741-2560/3/4/011.
- Kurokawa, Kazuo and Akihiko Nakano (2019). *The ER Exit Sites Are Specialized ER Zones for the Transport of Cargo Proteins from the ER to the Golgi Apparatus*. In: *The Journal of Biochemistry* 165.2, pp. 109–114.
- Kurokawa, Kazuo, Michiyo Okamoto, and Akihiko Nakano (2014). *Contact of Cis-Golgi with ER Exit Sites Executes Cargo Capture and Delivery from the ER*. In: *Nature communications* 5.1, pp. 1–7.

- Lai, Chunwei Walter et al. (2012). *ERdj4 Protein Is a Soluble Endoplasmic Reticulum (ER) DnaJ Family Protein That Interacts with ER-associated Degradation Machinery*. In: *Journal of Biological Chemistry* 287.11, pp. 7969–7978.
- Lasierra, N et al. (2009). “Experience of a Real-Time Tele-EEG Service”. In: *2009 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, pp. 5211–5214.
- Le Van Quyen, Michel and Anatol Bragin (2007). *Analysis of Dynamic Brain Oscillations: Methodological Advances*. In: *Trends in neurosciences* 30.7, pp. 365–373.
- Lebrun, Marc (Aug. 8, 2012). *An Analysis and Implementation of the BM3D Image Denoising Method*. In: *Image Processing On Line* 2, pp. 175–213. ISSN: 2105-1232. DOI: 10.5201/ipol.2012.l-bm3d.
- Lecoq, Jérôme, Natalia Orlova, and Benjamin F. Grewe (Oct. 2, 2019). *Wide. Fast. Deep. Recent Advances in Multi-Photon Microscopy of in vivo Neuronal Activity*. In: *The Journal of Neuroscience*, pp. 1527–18. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.1527-18.2019.
- Lecoq, Jérôme et al. (Nov. 2021). *Removing Independent Noise in Systems Neuroscience Data Using DeepInterpolation*. In: *Nature Methods* 18.11, pp. 1401–1408. ISSN: 1548-7091, 1548-7105. DOI: 10.1038/s41592-021-01285-2.
- Lee, Christopher and Lan Bo Chen (1988). *Dynamic Behavior of Endoplasmic Reticulum in Living Cells*. In: *Cell* 54.1, pp. 37–46.
- Lee, Daniel D. and H. Sebastian Seung (Oct. 1999). *Learning the Parts of Objects by Non-Negative Matrix Factorization*. In: *Nature* 401.6755, pp. 788–791. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/44565.
- Lee, Hyung-Chul and Chul-Woo Jung (2018). *Vital Recorder—a Free Research Tool for Automatic Recording of High-Resolution Time-Synchronised Physiological Data from Multiple Anaesthesia Devices*. In: *Scientific reports* 8.1, pp. 1–8.
- Leeb, William (Apr. 6, 2021). *Matrix Denoising for Weighted Loss Functions and Heterogeneous Signals*. arXiv: 1902.09474 [eess, math, stat].
- Lehtinen, Jaakko et al. (July 10, 2018). “Noise2Noise: Learning Image Restoration without Clean Data”. In: ed. by Jennifer Dy and Andreas Krause. Vol. 80. *Proceedings of Machine Learning Research*. Stockholmsmässan, Stockholm Sweden: PMLR, pp. 2965–2974.
- Lemesle, M et al. (2015). *Tele-Transmission of EEG Recordings*. In: *Neurophysiologie Clinique/Clinical Neurophysiology* 45.1, pp. 121–130.
- Lina, J. M. et al. (Aug. 2014). *Wavelet-Based Localization of Oscillatory Sources from Magnetoencephalography Data*. In: *IEEE transactions on bio-medical engineering* 61.8, pp. 2350–2364. ISSN: 1558-2531. DOI: 10.1109/TBME.2012.2189883. pmid: 22410322.
- Lipovsek, Marcela et al. (Feb. 3, 2021). *Patch-Seq: Past, Present, and Future*. In: *The Journal of Neuroscience* 41.5, pp. 937–946. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.1653-20.2020.

- Lippincott-Schwartz, Jennifer, Nihal Altan-Bonnet, and George H Patterson (2003). *Photobleaching and Photoactivation: Following Protein Dynamics in Living Cells*. In: *Nature cell biology*, S7–14.
- Lippincott-Schwartz, Jennifer, Erik Snapp, and Anne Kenworthy (2001). *Studying Protein Dynamics in Living Cells*. In: *Nature reviews Molecular cell biology* 2.6, pp. 444–456.
- Lippincott-Schwartz, Jennifer, Erik Lee Snapp, and Robert D. Phair (Oct. 2, 2018). *The Development and Enhancement of FRAP as a Key Tool for Investigating Protein Dynamics*. In: *Biophysical Journal* 115.7, pp. 1146–1155. ISSN: 0006-3495. DOI: 10.1016/j.bpj.2018.08.007. pmid: 30219286.
- Los, Georgyi V et al. (2008). *HaloTag: A Novel Protein Labeling Technology for Cell Imaging and Protein Analysis*. In: *ACS chemical biology* 3.6, pp. 373–382.
- Lovász, László et al. (1993). *Random Walks on Graphs: A Survey*. In: *Combinatorics, Paul Erdos is eighty* 2.1, pp. 1–46.
- Luisier, Florian, Thierry Blu, and Michael Unser (Mar. 2011). *Image Denoising in Mixed Poisson–Gaussian Noise*. In: *IEEE Transactions on Image Processing* 20.3, pp. 696–708. ISSN: 1941-0042. DOI: 10.1109/TIP.2010.2073477.
- Maas, Andrew IR et al. (2017). *Traumatic Brain Injury: Integrated Approaches to Improve Prevention, Clinical Care, and Research*. In: *The Lancet Neurology* 16.12, pp. 987–1048.
- Maier, Benjamin F. and Dirk Brockmann (Oct. 2017). *Cover Time for Random Walks on Arbitrary Complex Networks*. In: *Physical review E* 96.4, p. 042307.
- Majumdar, Satya N, Sanjib Sabhapandit, and Grégory Schehr (2016). *Exact Distributions of Cover Times for N Independent Random Walkers in One Dimension*. In: *Physical Review E* 94.6, p. 062131.
- Makitalo, M and A Foi (Jan. 2011). *Optimal Inversion of the Anscombe Transformation in Low-Count Poisson Image Denoising*. In: *IEEE Transactions on Image Processing* 20.1, pp. 99–109. ISSN: 1057-7149. DOI: 10.1109/TIP.2010.2056693.
- Makitalo, M. and A. Foi (Jan. 2013). *Optimal Inversion of the Generalized Anscombe Transformation for Poisson-Gaussian Noise*. In: *IEEE Transactions on Image Processing* 22.1, pp. 91–103. DOI: 10.1109/tip.2012.2202675.
- Mallat, Stephane G (1989). *A Theory for Multiresolution Signal Decomposition: The Wavelet Representation*. In: *IEEE transactions on pattern analysis and machine intelligence* 11.7, pp. 674–693. DOI: 10.1109/34.192463.
- Mammone, Nadia, Fabio La Foresta, and Francesco Carlo Morabito (Mar. 2012). *Automatic Artifact Rejection From Multichannel Scalp EEG by Wavelet ICA*. In: *IEEE Sensors Journal* 12.3, pp. 533–542. ISSN: 1530-437X, 1558-1748. DOI: 10.1109/JSEN.2011.2115236.
- Mangeol, Pierre, Bram Prevo, and Erwin JG Peterman (2016). *KymographClear and KymographDirect: Two Tools for the Automated Quantitative Analysis of Molecular and Cellular Dynamics Using Kymographs*. In: *Molecular biology of the cell* 27.12, pp. 1948–1957.

- Manley, Suliana et al. (2008). *High-Density Mapping of Single-Molecule Trajectories with Photoactivated Localization Microscopy*. In: *Nature methods* 5.2, pp. 155–157.
- Maruyama, Ryuichi et al. (July 2014). *Detecting Cells Using Non-Negative Matrix Factorization on Calcium Imaging Data*. In: *Neural Networks: The Official Journal of the International Neural Network Society* 55, pp. 11–19. ISSN: 1879-2782. DOI: 10.1016/j.neunet.2014.03.007.
- Mayeli, Ahmad et al. (2016). *Real-Time EEG Artifact Correction during fMRI Using ICA*. In: *Journal of neuroscience methods* 274, pp. 27–37.
- Mehnert, Martin, Thomas Sommer, and Ernst Jarosch (2010). *ERAD Ubiquitin Ligases: Multifunctional Tools for Protein Quality Control and Waste Disposal in the Endoplasmic Reticulum*. In: *Bioessays* 32.10, pp. 905–913.
- Meusser, Birgit et al. (Aug. 2005). *ERAD: The Long Road to Destruction*. In: *Nature Cell Biology* 7.8 (8), pp. 766–772. ISSN: 1476-4679. DOI: 10.1038/ncb0805-766.
- Meyer, Yves (1992). *Wavelets and Operators: Volume 1*. Cambridge University Press. 248 pp. ISBN: 978-0-521-45869-6. Google Books: y5L5HVlh3ngC.
- Mijović, Bogdan et al. (2010). *Source Separation from Single-Channel Recordings by Combining Empirical-Mode Decomposition and Independent Component Analysis*. In: *IEEE transactions on biomedical engineering* 57.9, pp. 2188–2196.
- Mitchell, Benjamin R. (Mar. 17, 2017). “The Spatial Inductive Bias of Deep Learning”. Thesis. Johns Hopkins University.
- Mukamel, Eran A., Axel Nimmerjahn, and Mark J. Schnitzer (Sept. 24, 2009). *Automated Analysis of Cellular Signals from Large-Scale Calcium Imaging Data*. In: *Neuron* 63.6, pp. 747–760. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2009.08.009.
- Muthukumaraswamy, Suresh (2013). *High-Frequency Brain Activity and Muscle Artifacts in MEG/EEG: A Review and Recommendations*. In: *Frontiers in human neuroscience* 7, p. 138.
- Nehls, Sarah et al. (2000). *Dynamics and Retention of Misfolded Proteins in Native ER Membranes*. In: *Nature cell biology* 2.5, p. 288.
- Niblack, Wayne (1985). *An Introduction to Digital Image Processing*. DNK: Strandberg Publishing Company. 215 pp. ISBN: 978-87-87200-55-4.
- Nichols, Benjamin J et al. (2001). *Rapid Cycling of Lipid Raft Markers between the Cell Surface and Golgi Complex*. In: *The Journal of cell biology* 153.3, pp. 529–542.
- Nixon-Abell, Jonathon et al. (2016). *Increased Spatiotemporal Resolution Reveals Highly Dynamic Dense Tubular Matrices in the Peripheral ER*. In: *Science* 354.6311, aaf3928.
- Nolan, H., R. Whelan, and R.B. Reilly (Sept. 2010). *FASTER: Fully Automated Statistical Thresholding for EEG Artifact Rejection*. In: *Journal of Neuroscience Methods* 192.1, pp. 152–162. ISSN: 01650270. DOI: 10.1016/j.jneumeth.2010.07.015.
- Oltmanns, Sebastian et al. (2020). *NOSA, an Analytical Toolbox for Multicellular Optical Electrophysiology*. In: *Frontiers in Neuroscience* 14, p. 712. ISSN: 1662-4548. DOI: 10.3389/fnins.2020.00712.

- Omer, H. and B. Torr sani (July 1, 2017). *Time-Frequency and Time-Scale Analysis of Deformed Stationary Processes, with Application to Non-Stationary Sound Modeling*. In: *Applied and Computational Harmonic Analysis* 43.1, pp. 1–22. ISSN: 1063-5203. DOI: 10.1016/j.acha.2015.10.002.
- Onton, Julie et al. (Jan. 2006). *Imaging Human EEG Dynamics Using Independent Component Analysis*. In: *Neuroscience & Biobehavioral Reviews* 30.6, pp. 808–822. ISSN: 01497634. DOI: 10.1016/j.neubiorev.2006.06.007.
- Pain, Charlotte et al. (Feb. 28, 2019). *Quantitative Analysis of Plant ER Architecture and Dynamics*. In: *Nature Communications* 10.1 (1), p. 984. ISSN: 2041-1723. DOI: 10.1038/s41467-019-08893-9.
- Patterson, George H and Jennifer Lippincott-Schwartz (2002). *A Photoactivatable GFP for Selective Photolabeling of Proteins and Cells*. In: *Science* 297.5588, pp. 1873–1877.
- Pignat, Jean Michel et al. (2013). *The Impact of Denoising on Independent Component Analysis of Functional Magnetic Resonance Imaging Data*. In: *Journal of neuroscience methods* 213.1, pp. 105–122.
- Pnevmatikakis, Eftychios A. (Apr. 2019). *Analysis Pipelines for Calcium Imaging Data*. In: *Current Opinion in Neurobiology* 55, pp. 15–21. ISSN: 1873-6882. DOI: 10.1016/j.conb.2018.11.004.
- Pnevmatikakis, Eftychios A. et al. (Jan. 20, 2016). *Simultaneous Denoising, Deconvolution, and Demixing of Calcium Imaging Data*. In: *Neuron* 89.2, pp. 285–299. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2015.11.037.
- Prinz, William A, Alexandre Toulmay, and Tamas Balla (2020). *The Functional Universe of Membrane Contact Sites*. In: *Nature Reviews Molecular Cell Biology* 21.1, pp. 7–24.
- Purdon, Patrick L et al. (2015). *Clinical Electroencephalography for Anesthesiologists part I: Background and Basic Signatures*. In: *Anesthesiology: The Journal of the American Society of Anesthesiologists* 123.4, pp. 937–960.
- Rapoport, Tom A, Long Li, and Eunyong Park (2017). *Structural and Mechanistic Insights into Protein Translocation*. In: *Annual review of cell and developmental biology* 33, pp. 369–390.
- Redner, S (1981). *Percolation and Conduction in a Random Resistor-Diode Network*. In: *Journal of Physics A: Mathematical and General* 14.9, p. L349.
- Remondelli, Paolo and Maurizio Renna (2017). *The Endoplasmic Reticulum Unfolded Protein Response in Neurodegenerative Disorders and Its Potential Therapeutic Significance*. In: *Frontiers in Molecular Neuroscience* 10. ISSN: 1662-5099.
- Ron, David and Peter Walter (2007). *Signal Integration in the Endoplasmic Reticulum Unfolded Protein Response*. In: *Nature reviews Molecular cell biology* 8.7, pp. 519–529.
- Ronneberger, Olaf, Philipp Fischer, and Thomas Brox (2015). “U-Net: Convolutional Networks for Biomedical Image Segmentation”. In: *Lecture Notes in Computer Science*. Springer International Publishing, pp. 234–241. DOI: 10.1007/978-3-319-24574-4_28.

- Safieddine, Doha et al. (2012). *Removal of Muscle Artifact from EEG Data: Comparison between Stochastic (ICA and CCA) and Deterministic (EMD and Wavelet-Based) Approaches*. In: *EURASIP Journal on Advances in Signal Processing* 2012.1, pp. 1–15.
- Sato, Yoshinobu et al. (June 1998). *Three-Dimensional Multi-Scale Line Filter for Segmentation and Visualization of Curvilinear Structures in Medical Images*. In: *Medical Image Analysis* 2.2, pp. 143–168. ISSN: 13618415. DOI: 10.1016/S1361-8415(98)80009-1.
- Schehr, Grégory and Satya N Majumdar (2014). “Exact Record and Order Statistics of Random Walks via First-Passage Ideas”. In: *First-Passage Phenomena and Their Applications*. World Scientific, pp. 226–251.
- Schomer, Donald L and Fernando Lopes Da Silva (2012). *Niedermeyer’s Electroencephalography: Basic Principles, Clinical Applications, and Related Fields*. Lippincott Williams & Wilkins.
- Schreiber, Thomas and Andreas Schmitz (1996). *Improved Surrogate Data for Nonlinearity Tests*. In: *Physical review letters* 77.4, p. 635.
- (2000). *Surrogate Time Series*. In: *Physica D: Nonlinear Phenomena* 142.3-4, pp. 346–382.
- Schröder, Martin and Randal J. Kaufman (Jan. 6, 2005). *ER Stress and the Unfolded Protein Response*. In: *Mutation Research* 569.1-2, pp. 29–63. ISSN: 0027-5107. DOI: 10.1016/j.mrfmmm.2004.06.056. pmid: 15603751.
- Schuss, Z. (1980). *Theory and Applications of Stochastic Differential Equations*. Wiley Series in Probability and Statistics - Applied Probability and Statistics Section.
- Schuss, Z, K Basnayake, and D Holcman (2019). *Redundancy Principle and the Role of Extreme Statistics in Molecular and Cellular Biology*. In: *Physics of life reviews* 28, pp. 52–79.
- Schuss, Zeev (2009). *Theory and Applications of Stochastic Processes: An Analytical Approach*. Springer Science & Business Media. 486 pp. ISBN: 978-1-4419-1605-1.
- Shen, Jingshi et al. (2005). *Stable Binding of ATF6 to BiP in the Endoplasmic Reticulum Stress Response*. In: *Molecular and cellular biology* 25.3, pp. 921–932.
- Shibata, Yoko, Gia K. Voeltz, and Tom A. Rapoport (Aug. 11, 2006). *Rough Sheets and Smooth Tubules*. In: *Cell* 126.3, pp. 435–439. ISSN: 0092-8674. DOI: 10.1016/j.cell.2006.07.019. pmid: 16901774.
- Sigl, Jeffrey C and Nassib G Chamoun (1994). *An Introduction to Bispectral Analysis for the Electroencephalogram*. In: *Journal of clinical monitoring* 10.6, pp. 392–404.
- Snapp, Erik L et al. (2006). *Monitoring Chaperone Engagement of Substrates in the Endoplasmic Reticulum of Live Cells*. In: *Proceedings of the National Academy of Sciences* 103.17, pp. 6536–6541.
- Sokolov, Igor M. (2019). *Extreme Fluctuation Dominance in Biology: On the Usefulness of Wastefulness: Comment on “Redundancy Principle and the Role of Ex-*

- treme Statistics in Molecular and Cellular Biology*” by Z. Schuss, K. Basnayake and D. Holcman. In: *Physics of life reviews*.
- Soltanian-Zadeh, Somayyeh et al. (Apr. 23, 2019). *Fast and Robust Active Neuron Segmentation in Two-Photon Calcium Imaging Using Spatiotemporal Deep Learning*. In: *Proceedings of the National Academy of Sciences of the United States of America* 116.17, pp. 8554–8563. ISSN: 1091-6490. DOI: 10.1073/pnas.1812995116.
- Somers, Ben, Tom Francart, and Alexander Bertrand (June 1, 2018). *A Generic EEG Artifact Removal Algorithm Based on the Multi-Channel Wiener Filter*. In: *Journal of Neural Engineering* 15.3, p. 036007. ISSN: 1741-2560, 1741-2552. DOI: 10.1088/1741-2552/aaac92.
- Sood, Vishal, Sidney Redner, and Dani Ben-Avraham (2004). *First-Passage Properties of the Erdős–Rényi Random Graph*. In: *Journal of Physics A: Mathematical and General* 38.1, p. 109.
- Sun, Yunhao et al. (2021). *Unraveling Single-Particle Trajectories Confined in Tubular Networks*. arXiv: 2112.05884.
- Sun, Yunhao et al. (Jan. 7, 2022). *Unraveling Single-Particle Trajectories Confined in Tubular Networks*. arXiv: 2112.05884 [physics, q-bio].
- Sweeney, K. T., T. E. Ward, and S. F. McLoone (May 2012). *Artifact Removal in Physiological Signals—Practices and Possibilities*. In: *IEEE Transactions on Information Technology in Biomedicine* 16.3, pp. 488–500. ISSN: 1089-7771, 1558-0032. DOI: 10.1109/TITB.2012.2188536.
- Sweeney, K. T. et al. (Sept. 2012). *A Methodology for Validating Artifact Removal Techniques for Physiological Signals*. In: *IEEE Transactions on Information Technology in Biomedicine* 16.5, pp. 918–926. ISSN: 1089-7771, 1558-0032. DOI: 10.1109/TITB.2012.2207400.
- Sweeney, Kevin T, Seán F McLoone, and Tomas E Ward (2012). *The Use of Ensemble Empirical Mode Decomposition with Canonical Correlation Analysis as a Novel Artifact Removal Technique*. In: *IEEE transactions on biomedical engineering* 60.1, pp. 97–105.
- Sweeney, Patrick et al. (Mar. 13, 2017). *Protein Misfolding in Neurodegenerative Diseases: Implications and Strategies*. In: *Translational Neurodegeneration* 6.1, p. 6. ISSN: 2047-9158. DOI: 10.1186/s40035-017-0077-5.
- Tatum, William O et al. (2011). *Artifact: Recording EEG in Special Care Units*. In: *Journal of Clinical Neurophysiology* 28.3, p. 14.
- Terasaki, Mark (2018). *Axonal Endoplasmic Reticulum Is Very Narrow*. In: *Journal of Cell Science* 131.4, jcs210450.
- Terasaki, Mark, Lan Bo Chen, and Keigi Fujiwara (1986). *Microtubules and the Endoplasmic Reticulum Are Highly Interdependent Structures*. In: *The Journal of cell biology* 103.4, pp. 1557–1568.
- Tremblay, Robin, Soohyun Lee, and Bernardo Rudy (July 20, 2016). *GABAergic Interneurons in the Neocortex: From Cellular Properties to Circuits*. In: *Neuron* 91.2, pp. 260–292. ISSN: 1097-4199. DOI: 10.1016/j.neuron.2016.06.033.

- Tsien, Roger Y (1998). *The Green Fluorescent Protein*. In: *Annual review of biochemistry* 67.1, pp. 509–544.
- Ulyanov, Dmitry, Andrea Vedaldi, and Victor Lempitsky (2018). “Deep Image Prior”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 9446–9454.
- Unser, M. and A. Aldroubi (Apr. 1996). *A Review of Wavelets in Biomedical Applications*. In: *Proceedings of the IEEE* 84.4, pp. 626–638. ISSN: 1558-2256. DOI: 10.1109/5.488704.
- Urigüen, Jose Antonio and Begoña Garcia-Zapirain (June 1, 2015). *EEG Artifact Removal—State-of-the-Art and Guidelines*. In: *Journal of Neural Engineering* 12.3, p. 031001. ISSN: 1741-2560, 1741-2552. DOI: 10.1088/1741-2560/12/3/031001.
- Valm, Alex M et al. (2017). *Applying Systems-Level Spectral Imaging and Analysis to Reveal the Organelle Interactome*. In: *Nature* 546.7656, pp. 162–167.
- Van Meer, Gerrit, Dennis R Voelker, and Gerald W Feigenson (2008). *Membrane Lipids: Where They Are and How They Behave*. In: *Nature reviews Molecular cell biology* 9.2, pp. 112–124.
- Varanasi, Mahesh K. and Behnaam Aazhang (Oct. 1989). *Parametric Generalized Gaussian Density Estimation*. In: *The Journal of the Acoustical Society of America* 86.4, pp. 1404–1415. ISSN: 0001-4966. DOI: 10.1121/1.398700.
- Villette, Vincent et al. (Dec. 12, 2019). *Ultrafast Two-Photon Imaging of a High-Gain Voltage Indicator in Awake Behaving Mice*. In: *Cell* 179.7, 1590–1608.e23. ISSN: 0092-8674. DOI: 10.1016/j.cell.2019.11.004.
- Vu, Trung, Evgenia Chunikhina, and Raviv Raich (Oct. 15, 2021). *Perturbation Expansions and Error Bounds for the Truncated Singular Value Decomposition*. In: *Linear Algebra and its Applications* 627, pp. 94–139. ISSN: 0024-3795. DOI: 10.1016/j.laa.2021.05.020.
- Wang, Bowen et al. (2021). *The Endoplasmic Reticulum Adopts Two Distinct Tubule Forms*. In: *bioRxiv*.
- Wang, Yung-Hung et al. (2014). *On the Computational Complexity of the Empirical Mode Decomposition Algorithm*. In: *Physica A: Statistical Mechanics and its Applications* 400, pp. 159–167.
- Weng, Tongfeng et al. (2017). *Multiple Random Walks on Complex Networks: A Harmonic Law Predicts Search Time*. In: *Physical Review E* 95.5, p. 052103.
- White, Diane M and C Anne Van Cott (2010). *EEG Artifacts in the Intensive Care Unit Setting*. In: *American journal of electroneurodiagnostic technology* 50.1, pp. 8–25.
- Witten, Daniela M., Robert Tibshirani, and Trevor Hastie (July 2009). *A Penalized Matrix Decomposition, with Applications to Sparse Principal Components and Canonical Correlation Analysis*. In: *Biostatistics (Oxford, England)* 10.3, pp. 515–534. ISSN: 1468-4357. DOI: 10.1093/biostatistics/kxp008.
- Worrell, Gregory Alan et al. (2012). *Recording and Analysis Techniques for High-Frequency Oscillations*. In: *Progress in neurobiology* 98.3, pp. 265–278.

- Xiao, Sheng et al. (Nov. 19, 2021). *Large-Scale Voltage Imaging in Behaving Mice Using Targeted Illumination*. In: *iScience* 24.11, p. 103263. ISSN: 2589-0042. DOI: 10.1016/j.isci.2021.103263.
- Xie, Michael E. et al. (Apr. 6, 2021). *High-Fidelity Estimates of Spikes and Subthreshold Waveforms from 1-Photon Voltage Imaging in Vivo*. In: *Cell Reports* 35.1, p. 108954. ISSN: 2211-1247. DOI: 10.1016/j.celrep.2021.108954.
- Xu, C. (Oct. 1, 2005). *Endoplasmic Reticulum Stress: Cell Life and Death Decisions*. In: *Journal of Clinical Investigation* 115.10, pp. 2656–2664. ISSN: 0021-9738. DOI: 10.1172/JCI26373.
- Xu, C Shan et al. (2021). *An Open-Access Volume Electron Microscopy Atlas of Whole Cells and Tissues*. In: *Nature* 599.7883, pp. 147–151.
- Young, Laurence J, Florian Ströhl, and Clemens F Kaminski (2016). *A Guide to Structured Illumination TIRF Microscopy at High Speed with Multiple Colors*. In: *JoVE (Journal of Visualized Experiments)* 111, e53988.
- Zamponi, Nahuel et al. (2018). *Mitochondrial Network Complexity Emerges from Fission/Fusion Dynamics*. In: *Scientific reports* 8.1, pp. 1–10.
- Zarzoso, Vicente, Pierre Comon, and Mariem Kallel (2006). “How Fast Is FastICA?” In: *2006 14th European Signal Processing Conference*. IEEE, pp. 1–5.
- Zhang, Haoming et al. (2020). *EEGdenoiseNet: A Benchmark Dataset for Deep Learning Solutions of EEG Denoising*. arXiv: 2009.11662.
- Zhang, Kai et al. (July 2017). *Beyond a Gaussian Denoiser: Residual Learning of Deep CNN for Image Denoising*. In: *IEEE Transactions on Image Processing* 26.7, pp. 3142–3155. ISSN: 1941-0042. DOI: 10.1109/TIP.2017.2662206.
- Zhang, Tongjie Y. and Ching Y. Suen (1984). *A Fast Parallel Algorithm for Thinning Digital Patterns*. In: *Communications of the ACM* 27.3, pp. 236–239.
- Zhang, Yide et al. (2019). “A Poisson-Gaussian Denoising Dataset with Real Fluorescence Microscopy Images”. In: *CVPR*.
- Zhao, Weisong et al. (Nov. 15, 2021). *Sparse Deconvolution Improves the Resolution of Live-Cell Super-Resolution Fluorescence Microscopy*. In: *Nature Biotechnology*, pp. 1–12. ISSN: 1546-1696. DOI: 10.1038/s41587-021-01092-2.
- Zhou, H Mimi, Ingrid Brust-Mascher, and Jonathan M Scholey (2001). *Direct Visualization of the Movement of the Monomeric Axonal Transport Motor Unc-104 along Neuronal Processes in Living Caenorhabditis Elegans*. In: *Journal of Neuroscience* 21.11, pp. 3749–3755.
- Zhu, Ciyu et al. (Dec. 1, 1997). *Algorithm 778: L-BFGS-B: Fortran Subroutines for Large-Scale Bound-Constrained Optimization*. In: *ACM Transactions on Mathematical Software* 23.4, pp. 550–560. ISSN: 0098-3500. DOI: 10.1145/279232.279236.

ABSTRACT

This dissertation is an investigation of biological phenomena related to the brain by means of mathematical models and quantitative methods. The leitmotiv of this work is the analysis of spatiotemporal series which naturally arise in biological systems at different scales.

In the first part of the thesis, we study the finest of such scales. I analyse intracellular protein dynamics in the endoplasmic reticulum (ER), an organelle of eukaryotic cells formed by a network of tubular membrane structures. The ER plays a key role in protein transport, and its dysfunction has been associated with numerous diseases, including, in particular, neurodegenerative disorders. Previous experimental observations suggested a possible deviation of ER luminal transport compared to classical diffusion. Based on this hypothesis, I introduce a graph model to describe ER protein dynamics. I analyse the model and develop numerical simulations, revealing a possible mechanism of aggregated protein transport that deviates from purely diffusive motion.

Then, to further test the predictions of the model, we turn to the analysis of experimental data. While protein mobility has been traditionally characterized by fluorescence imaging, the morphological characteristics of the ER pose new challenges to a quantitative analysis of such small scale dynamics. To address these issues, I introduce a novel image processing method to analyse ER dynamics based on photoactivatable fluorescent proteins. By joining analysis and reduction of noise with automatic segmentation of the ER, the technique can provide a robust estimation of the timescale of transport. Moreover, it allows us to characterize the spatial heterogeneity of the protein mixing process. I present and compare results for luminal, membrane, and misfolded proteins in the ER.

In the second part of the dissertation, we study neuronal signals at coarser scales. First, at the scale of the single neuron, I present a denoising method suitable for optical recording of single-cell activity in awake, behaving mice via fluorescent voltage indicators. I show how it is possible to reduce instrumental and photon-counting noise in such time series, allowing us to extract spike patterns at lower acquisition frequency. Such results enable simultaneous recording of multiple cells, thus allowing to explore the correlation of spikes and voltage oscillations within ensembles of neurons.

Finally, in the last chapters, we reach the coarsest scale with the study of electroencephalograms (EEG) which record the activity of the entire brain. Motivated by the applications of EEG in clinical monitoring, I introduce a new wavelet-based method that can attenuate undesired artefacts which contaminate the recording of the physiological EEG signal. The method is based on the remapping of the wavelet coefficients according to a reference distribution extracted from clean portions of the EEG signal. This technique can provide a flexible alternative to traditional approaches such as wavelet thresholding in the context of real-time clinical monitoring.

In conclusion, this thesis illustrates how an interdisciplinary approach combining experimental data with mathematical modelling and signal processing can provide new tools for the understanding of a wide variety of biological mechanisms, ranging from protein transport to EEG monitoring.

KEYWORDS

Data analysis, signal processing, mathematical modelling, endoplasmic reticulum, neural time series, EEG

RÉSUMÉ

Cette thèse est une étude de phénomènes biologiques liés au cerveau à travers des modèles mathématiques et méthodes quantitatives. Le leitmotiv de ce travail est l'analyse des séries spatio-temporelles qui apparaissent naturellement dans les systèmes biologiques à différentes échelles.

La première partie de la thèse se penche sur l'échelle la plus fine. J'étudie la dynamique des protéines au sein du réticulum endoplasmique (RE), un organe de la cellule eucaryote formé d'un réseau de structures membranaires tubulaires. Le RE joue un rôle clé dans le transport des protéines et son dysfonctionnement a été associé à de nombreuses maladies, dont notamment les troubles neurodégénératifs. Des observations expérimentales antérieures ont suggéré une possible déviation du transport luminal du RE par rapport à la diffusion classique. Pour tester cette hypothèse, j'introduis un modèle de réseau pour décrire la dynamique des protéines dans le RE. J'analyse le modèle et développe des simulations numériques, en révélant un possible mécanisme de transport de protéines agrégées qui s'écarte du mouvement purement diffusif.

Pour tester davantage les prédictions du modèle, nous nous tournons ensuite vers l'analyse des données expérimentales. Alors que la mobilité des protéines a été traditionnellement caractérisée par l'imagerie par fluorescence, les caractéristiques morphologiques du RE posent de nouveaux défis à une analyse quantitative d'une telle dynamique à petite échelle. Pour résoudre ces problèmes, j'introduis une nouvelle méthode de traitement d'image pour analyser la dynamique du RE basée sur des protéines fluorescentes photoactivables. Cette technique joint l'analyse et la réduction du bruit à la segmentation automatique du RE, et peut fournir une estimation robuste de l'échelle de temps de transport. Cela nous permet également de caractériser l'hétérogénéité spatiale du processus de mélange des protéines. Je présente et compare les résultats pour les protéines luminales, membranaires et mal repliées dans le RE. Dans la deuxième partie de la thèse, nous étudions les signaux neuronaux à des échelles supérieures. Tout d'abord, à l'échelle du neurone, je présente une méthode de débruitage adaptée à l'enregistrement optique de l'activité unicellulaire chez les souris éveillées via des indicateurs de voltage fluorescents. Je montre comment il est possible de réduire le bruit instrumental et quantique dans de telles séries temporelles, ce qui permet d'extraire les potentiels d'action à une fréquence d'acquisition plus faible. Ce résultat rend possible l'enregistrement simultané de plusieurs cellules, permettant ainsi d'explorer la corrélation des pointes et des oscillations de voltage au sein d'un ensemble de neurones.

Enfin, dans les derniers chapitres, nous atteignons le niveau le plus global avec l'étude des électroencéphalogrammes (EEG), qui enregistrent l'activité de l'ensemble du cerveau. Motivés par les applications de l'EEG dans la surveillance clinique, j'introduis une nouvelle méthode basée sur les ondelettes pour atténuer les artefacts indésirables qui contaminent l'enregistrement du signal EEG physiologique. La méthode est basée sur le remappage des coefficients d'ondelettes selon une distribution de référence extraite de portions propres du signal EEG. Notre technique peut fournir une alternative flexible aux approches traditionnelles, comme le seuillage d'ondelettes, dans le contexte de la surveillance clinique en temps réel.

En conclusion, cette thèse illustre la manière dont une approche interdisciplinaire combinant données expérimentales, modélisation mathématique et traitement du signal peut fournir de nouveaux outils pour la compréhension d'une large variété de mécanismes biologiques, allant du transport des protéines au suivi EEG.

MOTS CLÉS

Analyse des données, traitement du signal, modélisation mathématique, réticulum endoplasmique, séries temporelles neuronales, EEG