

The
University
Of
Sheffield.

**BAYESIAN MODELLING FOR
SUPER-RESOLUTION LOCALISATION
MICROSCOPY**

By: Dimitrios Kiagias

A thesis submitted in partial fulfilment
of the requirements for the degree of
Doctor of Philosophy

University of Sheffield
Faculty of Science
School of Mathematics and Statistics

September, 2019

Copyright

All rights reserved. No part of this publication may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the publisher, except in the cases where correct citation is provided.

Student: **Mr Dimitrios
Kiagias**

Supervisor: **Dr Miguel
Juarez**

Supervisor: **Dr Kostas
Triantafyllopoulos**

Signature: _____

Signature: _____

Signature: _____

Date: _____

Date: _____

Date: _____

Abstract

Super resolution microscopy techniques, which allow live organisms imaging beyond the diffraction limit, are in the forefront of understanding living systems. Single Molecule Localisation Microscopy (SMLM) randomly turns on a sparse subset of photo-switchable molecules at different times, enabling temporal localisations of molecules otherwise limited by the diffraction barrier of light. Stacks of images, represented by lattices of intensities, are produced over time making both quantitative analysis (such as counting) and localisation of single molecules, an essential key to understand biological structures and interactions inside living cells.

In this thesis we introduce a novel set of structural functions named as Single Molecule Pattern, SiMPa, functions, which describe a molecule's diffraction of intensity across a predetermined neighbourhood based on its position within a pixel. The diffraction is obtained as a function of the single event intensity, a quantity directly related to the number of photons emitted by a molecule. An individual frame inference is presented by considering the SiMPa functions within a probabilistic scheme via a Bayesian approach, to count and localise molecules. This is achieved with our novel localisation scheme based on structures formed within a moving region that scans the frame.

We then propose a state-space model to model stacks of frames, that are time related, where the SiMPa functions are embedded within a Markov switching model. The latter accounts for the switching behaviour of pixels over time, consisting of active and de-active states, and their corresponding spatial structure. Additionally, we present a data generation mechanism to obtain sequences of SMLM frames using the SiMPa functions and a transition matrix for the states of the molecules. A thorough sensitivity and evaluation analysis is considered on a synthetic data set, and a comparison with one of most popular state-of-the-art alternatives, both on synthetic and realistic data.

Acknowledgements

I would like to thank my supervisor Miguel Juarez for his guidance, support and mostly for our excellent collaboration the past four years. I could not ask for a better person to work with. Also, I would like to thank my co-supervisor Kostas Triantafyllopoulos for his assistance and valuable advice during my PhD.

The biggest thanks goes to my parents, Katerina and Mpampis, and my brother, Konstantinos, for their continuous and unconditional love and support, not only during my PhD but throughout all these years of studies. This thesis is dedicated to them.

Contents

	Page
1 Introduction	1
2 Literature Review	5
2.1 Principles of fluorescence microscopy	5
2.2 Super-resolution localisation microscopy	7
2.2.1 Image formation	8
2.2.2 Principles of single molecule localisation microscopy	10
2.2.3 Current localisation methodologies and algorithms	15
2.3 Motivation of this research	26
3 Frame-wise image analysis	29
3.1 Introduction	29
3.2 Single Molecule Pattern (SiMPa) functions	30
3.2.1 Special cases of the SiMPa functions	39
3.2.2 Extended Single Molecule Pattern (SiMPa _X) functions	46
3.2.3 Function of relative side length of pixel l and molecule's radius $R, g(R, l)$	52
3.2.4 Power of diffraction parameter c	54
3.2.5 Single frame generation using the SiMPa functions	59
3.3 Individual frame inference based on SiMPa	61
3.3.1 Markov random fields (MRF)	62
3.3.2 Bayesian modelling of individual frames using the SiMPa func- tions	64
3.3.3 Pattern-Configuration-Realisation concept	72
3.3.4 Individual frame inference	84
3.3.5 Bayesian update Pseudo-algorithms	94
3.4 Implementation	97

4	Stack of frames analysis	104
4.1	Introduction	104
4.2	General form of state-space model	106
4.2.1	The Markov Switching model	106
4.2.2	Markov Switching model based on SiMPa functions	109
4.2.3	Bayesian inference for the Markov Switching model based on SiMPa functions	120
4.3	Sampling scheme for sequence of images using the SiMPa functions	150
4.4	Implementation	153
5	Application and sensitivity analysis	164
5.1	Introduction	164
5.2	Synthetic data generation	165
5.3	Application and comparisons on synthetic data	173
5.3.1	Application using individual frame inference based on SiMPa functions	173
5.3.2	Comparison with ThunderSTORM	194
5.4	Sensitivity analysis	201
5.4.1	Sensitivity analysis for counting and localisation	201
5.4.2	Sensitivity analysis on parameters	230
5.5	Application on realistic data	239
5.6	Application on SuReSim data and comparisons	259
6	Conclusion and Discussion	276
A	Appendix	282
A.1	Figures from implementation of MSM SiMPa model (Sec 4.4)	283
A.2	Figures - Application on synthetic data (Sec 5.3)	285
A.3	Figures - Sensitivity analysis (Sec 5.4)	296
A.4	Figures - Application on realistic data (Sec 5.5)	307
A.5	Figures - Application on SuReSim data (Sec 5.6)	314
	Appendix	320
	Bibliography	320

List of Figures

2.1.1	Jablonski diagram	6
2.2.1	Principle of single molecule localisation microscopy (SMLM) and photoswitching features	9
2.2.2	Spatial resolution against molecular density in SMLM.	11
3.2.1	Graphical representation of the Single Molecule Pattern (SiMPa) functions.	32
3.2.2	Minimum distances of a molecule placed r units with direction θ away from the origin from the neighbouring pixels.	35
3.2.3	Proposed use of SiMPa functions	38
3.2.4	Evaluation of the SiMPa functions for different combinations of r and θ	39
3.2.5	Special case when a molecule lies on one of the edges of the central pixel of the SiMPa functions	41
3.2.6	Minimum distances on the special case where the molecule lies in on of the edges of the central pixel	43
3.2.7	Implementation of special case when the molecule is placed on an edge of a pixel	44
3.2.8	Special case when a molecule lies on the intersection of two edges of the central pixel	46
3.2.9	Graphical representation of the Extended Single Molecule Pattern (SiMPa _X) functions	47
3.2.10	Minimum distances of the neighboring pixels of a molecule placed r units on direction θ away from the origin.	49
3.2.11	Function of relative side length of pixel l and molecule's radius R , $g(R, l)$	53
3.2.12	Maximum distance of a molecule on the diagonal line from the origin of the SiMPa _X functions	57
3.3.1	Neighbourhood systems of 1st and 2nd order.	62

3.3.2	Example of two Point Spread Functions (PSF) contributing in a moving region (MR)	70
3.3.3	Example of a <i>pattern</i> along with every possible <i>configuration</i> that can be formed from, with respect to the assumption of up to two PSFs allowed to overlap	74
3.3.4	Example of the periodicity of the SiMPa functions when the system is rotated	78
3.3.5	Implementation of case where a single pixel of a PSF lands on an MR.	81
3.3.6	Implementation of case where a single pixel of a PSF lands on an MR.	82
3.3.7	Simple example of a <i>pattern</i> along with every possible <i>configuration</i> it can arise from, with respect to the assumption of up to two PSFs allowed to overlap	90
3.3.8	Different examples of identification issues regarding localisation of molecules using the SiMPa functions.	92
3.4.1	Simulated individual frame using the SiMPa functions	98
3.4.2	Prior and posterior distributions for parameters in the simulated individual frame using the SiMPa functions	99
3.4.3	Posterior probability maps for the simulated individual frame based on SiMPa functions	102
4.2.1	Graphical representation of the state space model based on SiMPa functions	114
4.2.2	Transition of a state via a transition matrix and the corresponding neighbourhood region of the MRF.	118
4.2.3	Trajectories of raw data of two different pixels over all frames	140
4.4.1	Subset of consecutive frames from the entire stack generated for application on 'Circle within circle' synthetic data.	154
4.4.2	Prior and posterior distributions for parameters in the simulated individual frame using the SiMPa functions	155
4.4.3	Posterior probability heatmaps filtered by the average field used on the implementation of the Markov switching model based on SiMPa functions	157
4.4.4	Double event posterior probability heatmaps filtered by the average field used on the implementation of the Markov switching model based on SiMPa functions	158
4.4.5	Average field per frame used on the implementation of the Markov switching model based on SiMPa functions	159

4.4.6	Individual frame reconstructions used on the implementation of the Markov switching model based on SiMPa functions	162
5.2.1	Circle within circle - synthetic data using the SiMPa functions	168
5.2.2	Quantities used for prior elicitation	170
5.2.3	Priors elicited from the 'Circle within circle' synthetic data.	172
5.3.1	Selected frames from the on 'Circle within circle' synthetic data	174
5.3.2	Posterior probability maps filtered by the average field for a number of frames participating in the application on synthetic data.	180
5.3.3	Double event posterior probability maps filtered by the average field for a number of frames included in the application on synthetic data.	182
5.3.4	Precision-Recall curves for 'Circle within circle' frames used in application	183
5.3.5	Individual frame reconstructions for a number of frames included in the application on synthetic data, using two different thresholds for the posterior probability maps filtered by the average field.	188
5.3.6	Individual frame reconstructions for a number of frames included in the application on synthetic data, using the conventional threshold scheme based on the average field.	189
5.3.7	Final reconstructions of the super resolution image for the chosen frames in application on 'Circle within circle' synthetic data.	192
5.3.8	Final reconstructions of the super resolution image for the application on synthetic data.	193
5.3.9	Individual frame reconstructions for selected frames included in the application on synthetic data, using the conventional threshold scheme based on the average field and the localisations obtained by ThunderSTORM on single molecule fitting.	198
5.3.10	Individual frame reconstructions for selected frames included in the application on synthetic data, using the conventional threshold scheme based on the average field and the localisations obtained by ThunderSTORM on multiple-emitters fitting.	199
5.3.11	Final reconstructions of the super resolution image for the application on synthetic data.	200
5.4.1	Different noise levels when generating the 'Circle within circle' synthetic data.	203
5.4.2	Priors elicited for single event intensity I and precision τ_b when small and large noise levels are present	204

5.4.3	'Circle within circle' data - Frame 57 under different noise levels along with posterior probability maps filtered by average field, using both counting and localisation algorithms.	206
5.4.4	Selected region A in Frame 57 along with posterior probability maps filtered by average field, using both counting and localisation algorithms - Well isolated PSF	209
5.4.5	Selected region B in Frame 57 along with posterior probability maps filtered by average field, using both counting and localisation algorithms - Two PSFs overlapping on four pixels	210
5.4.6	Selected region C in Frame 57 along with posterior probability maps filtered by average field, using both counting and localisation algorithms - Two cases of two PSFs overlapping on six and one pixel	211
5.4.7	'Circle within circle' data - Frame 6 under different noise levels along with posterior probability maps filtered by average field, using both counting and localisation algorithms.	213
5.4.8	Selected region A in Frame 6 along with posterior probability maps filtered by average field, using both counting and localisation algorithms - Three PSFs forming a small high intensity region without violation	214
5.4.9	'Circle within circle' data - Frame 62 under different noise levels along with posterior probability maps filtered by average field, using both counting and localisation algorithms.	216
5.4.10	Selected region A in Frame 62 along with posterior probability maps filtered by average field, using both counting and localisation algorithms - Three PSFs forming a large high intensity region without violation	217
5.4.11	'Circle within circle' data - Frame 23 under different noise levels along with posterior probability maps filtered by average field, using both counting and localisation algorithms.	219
5.4.12	Selected region B in Frame 23 along with posterior probability maps filtered by average field, using both counting and localisation algorithms - Three PSFs forming a small high intensity region under violation	220
5.4.13	Selected region A in Frame 23 along with posterior probability maps filtered by average field, using both counting and localisation algorithms - Five PSFs forming a large high intensity region under violation	222
5.4.14	Precision-Recall curves under small and free of noise levels for the frames used in sensitivity analysis, including both the counting and localisation procedures	223

5.4.15	Precision-Recall curves under regular and large noise levels for the frames used in sensitivity analysis, including both the counting and localisation procedures	224
5.4.16	Position inference using the counting algorithm on Frames 1 and 10 from 'Circles within circle' data used in application	226
5.4.17	Prior and Posterior distributions of I , τ_b and c for Frames 1 and 10 from 'Circles within circle' data used in application	228
5.4.18	Prior distributions for the parameters I , τ_b , c and β_0 & β_f used for sensitivity analysis	232
5.4.19	Prior/Posterior distributions obtained using different priors for the parameters I , τ_b , c and β_0 & β_f , part of the sensitivity analysis on the parameters	234
5.4.20	Prior and Posterior distributions of I , τ_b and c for Frame 338 from 'Circle within circle' data using different proportions of the background intensity.	236
5.4.21	Posterior probability maps filtered by the average field for Frame 338 from 'Circle within circle', using different proportions of the background intensity	238
5.4.22	Prior distributions for the parameters ξ_{11} and ξ_{00} of the transition matrix ξ used for sensitivity analysis	238
5.4.23	Prior and Posterior distributions for the parameters ξ_{11} and ξ_{00} of the transition matrix ξ used for sensitivity analysis.	239
5.5.1	Selected frames from the realistic LDLS Tubulins data.	241
5.5.2	Quantities used for prior elicitation - realistic LDLS Tubulins data.	242
5.5.3	Priors elicited or chosen from the realistic LDLS Tubulins data.	243
5.5.4	Posterior probability maps filtered by the average field for a number of frames included in the application on realistic LDLS Tubulins data.	247
5.5.5	Double event posterior probability maps filtered by the average field for a number of frames included in the application on realistic LDLS Tubulins data.	248
5.5.6	Individual frame reconstructions for a number of frames included in the application on realistic LDLS Tubulins data, using two different thresholds for the posterior probability maps filtered by the average field.	250
5.5.7	Individual frame reconstructions for a number of frames included in the application on realistic LDLS Tubulins data, using the conventional threshold scheme based on the average field.	253

5.5.8	Final reconstructions of the super resolution image for the chosen frames in the application on realistic LDLS Tubulins data.	254
5.5.9	Final reconstructions of the super resolution image for the application on realistic LDLS Tubulins data.	255
5.5.10	Final reconstructions of the super resolution image for the application on realistic LDLS Tubulins data using SiMPa and ThunderSTORM.	256
5.5.11	Individual frame reconstructions for a number of frames included in the application on realistic LDLS Tubulins data, using the conventional threshold scheme based on the average field.	258
5.6.1	Selected frames from SuReSim data	260
5.6.2	Quantities used for prior elicitation - SuReSim data.	262
5.6.3	Priors elicited or chosen from SuReSim data.	263
5.6.4	Posterior probability maps filtered by the average field for a number of frames participating in the application on SuReSim data.	265
5.6.5	Double event posterior probability maps filtered by the average field for a number of frames included in the application on SuReSim data.	266
5.6.6	Individual frame reconstructions for a number of frames included in the application on SuReSim data, using the conventional threshold scheme based on the average field.	270
5.6.7	Final reconstructions of the super resolution image for the application on SuReSim data.	271
5.6.8	Precision-Recall curves for a number of chosen frames used in the application on SuReSim data.	272
5.6.9	Final reconstructions of the super resolution image on SuReSim data, using SiMPa and ThunderSTORM.	275
A.1.1	Prior and posterior distributions for the frame wise parameters β_0 's used on the implementation of the Markov switching model based on SiMPa functions.	283
A.1.2	Prior and posterior distributions for the frame wise parameters β_f 's used on the implementation of the Markov switching model based on SiMPa functions.	284
A.2.1	Average fields for 'Circle within circle' frames used in the application	285
A.2.2	Sensitivity-Specificity (ROC) curves for 'Circle within circle' frames used in the application	286
A.2.3	Posterior distributions of single event intensity I for frames presented in the application on 'Circle within circle' synthetic data	287

A.2.4	Trace plots of the drawn posterior sample of the single event intensity I for the frames used in the application on 'Circle within circle' synthetic data	288
A.2.5	Posterior distributions of background error τ_b for frames presented in the application on 'Circle within circle' synthetic data	289
A.2.6	Trace plots of the drawn posterior sample of the background error τ_b for the frames used in the application on 'Circle within circle' synthetic data	290
A.2.7	Posterior distributions of power of spread parameter c for frames presented in the application on 'Circle within circle' synthetic data	291
A.2.8	Trace plots of the drawn posterior sample of the power of spread parameter c for the frames used in the application on 'Circle within circle' synthetic data	292
A.2.9	Posterior distributions of field control parameters β_0 and β_F for frames presented in the application on 'Circle within circle' synthetic data	293
A.2.10	Trace plots of the drawn posterior sample of the field control parameter β_0 for the frames used in the application on 'Circle within circle' synthetic data	294
A.2.11	Trace plots of the drawn posterior sample of the field control parameter β_f for the frames used in the application on 'Circle within circle' synthetic data	295
A.3.1	Prior/Posterior distributions obtained for the parameters τ_b , c and β_0 & β_f for Frame 338, using different prior distributions for I	296
A.3.2	Prior/Posterior distributions obtained for the parameters I , c and β_0 & β_f for Frame 338, using different prior distributions for τ_b	297
A.3.3	Prior/Posterior distributions obtained for the parameters I , τ_b and β_0 & β_f for Frame 338, using different prior distributions for c	298
A.3.4	Prior/Posterior distributions obtained for the parameters I , τ_b and c for Frame 338, using different prior distributions for β_0 and β_f	299
A.3.5	Posterior probability maps filtered by the average field for Frame 338 used on parameter sensitivity.	300
A.3.6	Precision-Recall curves under regular and large noise levels for the frames used in sensitivity analysis, including both the counting and localisation procedures	301
A.3.7	Sensitivity-Specificity (ROC) curves under small and free of noise levels for the frames used in sensitivity analysis, including both the counting and localisation procedures	302

A.3.8	Frame 57 - Posterior probability map filtered by the average field for specified locations within the pixels, according to our localisation algorithm based on SiMPa functions	303
A.3.9	Frame 6 - Posterior probability map filtered by the average field for specified locations within the pixels, according to our localisation algorithm based on SiMPa functions	304
A.3.10	Frame 62 - Posterior probability map filtered by the average field for specified locations within the pixels, according to our localisation algorithm based on SiMPa functions	305
A.3.11	Frame 23 - Posterior probability map filtered by the average field for specified locations within the pixels, according to our localisation algorithm based on SiMPa functions	306
A.4.1	Posterior distributions of single event intensity I for frames presented in the application on realistic LDLS Tubulins data.	307
A.4.2	Posterior distributions of background error precision τ_b for frames presented in the application on realistic LDLS Tubulins data.	308
A.4.3	Posterior distributions of power of spread parameter c for frames presented in the application on realistic LDLS Tubulins data.	309
A.4.4	Posterior distributions of field control parameters β_0 and β_F for frames presented in the application on realistic LDLS Tubulins data.	310
A.4.5	Average fields for a number of frames included in the application on realistic LDLS Tubulins data.	311
A.4.6	Posterior probability maps filtered by the average field for a number of frames included in the application on realistic LDLS Tubulins data, when the elicited prior distribution for τ_b is used.	312
A.4.7	Average fields for a number of frames included in the application on realistic LDLS Tubulins data. when the elicited prior distribution for τ_b is used.	313
A.5.1	Posterior distributions of single event intensity I for frames presented in the application on SuReSim data.	314
A.5.2	Posterior distributions of background error precision τ_b for frames presented in the application on SuReSim data.	315
A.5.3	Posterior distributions of power of spread parameter c for frames presented in the application on SuReSim data.	316
A.5.4	Posterior distributions of field control parameters β_0 and β_F for frames presented in the application on SuReSim data.	317
A.5.5	Average fields for a number of chosen frames used in the application on SuReSim data.	318

A.5.6	Sensitivity-Specificity (ROC) curves for a number of chosen frames used in the application on SuReSim data.	319
-------	---	-----

List of Tables

3.3.1	Number of different patterns that can be formed when the MR size is 3×3 pixels (same as PSF's size) and total number of configurations that can exist (with respect to the assumption of up to two PSF allowed to coincide).	75
3.3.2	Number of possible combinations a PSF's parts can fall within a MR, $N_{\{\text{MR}, n_{\text{PSF}}\}}$ for a single realisation.	79
3.3.3	Combinations of the SiMPa functions when a PSF contributes one or four pixels on a MR, along with the corresponding quadrants the molecule lies in, with respect to rotation and PSF's relative position with the MR.	83
3.3.4	Combinations of the SiMPa functions when a PSF contributes three or six pixels on a MR, along with the corresponding quadrants the molecule lies in, with respect to rotation and PSF's relative position with the MR.	83
4.4.1	Counting of molecules using different expected duration $E(D)$ of 'On' and 'Off' states.	163
5.3.1	Conventional threshold scheme based on the average field	187
5.3.2	Performance of different threshold approaches in terms of counting for a number of chosen frames on application	190
5.3.3	Performance of different threshold approaches in terms of counting for frames on application on synthetic data	194
5.3.4	Performance of ThunderSTORM and individual frame inference based on SiMPa functions for 'Circle within circle' synthetic data.	201
5.5.1	Conventional threshold scheme based on the average field for application on realistic LDLS Tubulins	252

5.6.1	Conventional threshold scheme based on the average field for SuReSim data	269
5.6.2	Performance of ThunderSTORM and individual frame inference based on SiMPa functions for SuReSim data.	274
6.0.1	Computational times for individual frame inference based on SiMPa functions	278

Chapter 1

Introduction

Single Molecule Localisation Microscopy (SMLM) enables overcoming the diffraction barrier of light in fluorescence microscopy, where light emitting molecules, called fluorophores, are focused within a specimen. A fluorophore's diffraction, commonly described by the Point Spread function (PSF), can prevent fluorophores in close proximity from being resolved as a result of overlapping PSFs. STORM involves the stochastic activation of a sparse subset of fluorophores on distinct time points, achieving temporal isolation of molecules that were originally spatially indistinguishable. Cycles of this procedure generate a stack of frames each one consisting of a few molecules, mainly different or the same ones due to fluorophore blinking properties, which corresponds to their capability of cycling between excitation and dark states. Processing these stack of frames is key in order to determine and retrieve information on molecules' positions, behaviour and interaction.

In this thesis, we present a novel mathematical model within the Bayesian perspective in order to process the spatial association on individual frames, also embedded in a state space model to allow for inference on stacks of frames. The main contribution of this research is providing a method capable of dealing with overlapping PSFs, thus avoiding the loss of potentially valuable information in the samples, while

propagating uncertainty in a coherent way. This loss corresponds to a significant amount of information regarding molecular interaction and evolution within cells to be left unexplored and has a substantial effect on both visualisation and quantitative analysis of target structures. Based on a novel probabilistic scheme capable of accommodating overlapping PSFs, our approach can perform molecule localisation to provide reconstructions of a specimen, and used as a quantitative mechanism to provide metrics for the number of molecules.

Chapter 2 offers a literature review on current methodologies for super resolution localisation microscopy, while also focusing on the principles and challenges one should consider during processing. In the first section, we provide a brief introduction on fluorescence microscopy highlighting the main features, as well as the major limitations when imaging goes down to the molecular level. In order to overcome this barrier, super resolution fluorescence microscopy techniques were developed, from which we focus on single molecule localisation microscopy (SMLM) and discuss in the second section. We present and analyse SMLM main principles and properties, with the following section presenting some localisation procedures developed in the past decade. For each technique we provide brief descriptions and we outline the most commonly used processing mechanisms, with the last section addressing obstacles we tackle and overcome with our proposed methodology.

Chapter 3 consists of two main parts. The first describes the novel Single Molecule Pattern (SiMPa) functions, which account for a molecule diffusion of the intensity across a predetermined region, based on its relative position within a pixel. A detailed description on their derivation is presented, along with a proposed way of their application, followed by the individual frame inference via a Bayesian perspective. This second part relies on a unobserved mechanism, modelled by a Markov random field (MRF), to identify regions of interest (ROI) that potentially contain active molecules. Molecule counting and localisation from overlapping, or not, PSFs is performed within a probabilistic scheme based on the novel pattern-configuration concept for a moving region which scans the frame. The chapter concludes with an

implementation on a toy example describing the output of our procedure.

In Chapter 4, our proposed methodology based on the SiMPa functions and the pattern-configuration scheme, is embedded within a Markov switching model which allows processing of entire stack of frames. The first section consists of an introduction to the Markov switching model, followed by the appropriate arrangements to incorporate the SiMPa functions and the probabilistic localisation scheme. We present the complete Bayesian inference split in two parts, one for parameter updating and localisation and the second for pixels state (active - de-active) estimation based on a Forward Filtering Backwards Sampling (FFBS) algorithm. The latter relies on a transition matrix along with a neighbourhood, with the states being directly related to the existence of an active molecule, since they identify pixels that are potentially parts of a PSF. A data generation mechanism for stack of frames using the SiMPa functions follows, with the chapter concluding on an implementation of the model on a toy example.

Chapter 5 contains the application and evaluation of our proposed methodology, as well as comparison against an off-the-shelf alternative, using a synthetic data set. In the first section we describe the data generation and discuss on the parameter selection, along with a simple elicitation procedure to obtain prior distributions. Next, we apply the individual frame inference based on SiMPa functions on a stack of frames, where we both perform reconstructions of the underlying structure and provide performance metrics to quantitatively evaluate our inference. We obtain similar output using ThunderSTORM, a popular choice in SMLM, with which we compare our results. In the following section we conduct a thorough sensitivity analysis on both the individual and stack of frames models. We investigate alterations in precision and accuracy our localisation scheme faces under various noise levels, and examine the effect these have on challenging overlapping scenarios. The chapter continues with an analysis on parameter sensitivities and concludes with applications on realistic data used for the Challenge 2013 [Sage et al., 2015] to evaluate current super resolution localisation techniques and simulated ground truth data using SuReSim

[[Venkataramani et al., 2016](#)].

In Chapter 6 we conclude and summarise the results, while also addressing the main advantages and issues of our proposed methodology. For the latter, we provide a detailed discussion on potential considerations for improvement and future work, regarding numerous parameters and parts of our model.

Chapter 2

Literature Review

2.1 Principles of fluorescence microscopy

Fluorescence is the emission of light of a specific wavelength, which takes place in a brief period of time, after light of generally shorter wavelength has been absorbed. Distinguishing between the absorbed and emitted light, in terms of considering the difference in their wavelengths [Stokes, 1852], allows visualisation only of molecules that fluoresce [Valeur, 2003].

The use of fluorescent molecules, or fluorophores, was introduced in biological experiments around the 1930s, initiating the development of fluorescence microscopy, nowadays one of the most important techniques in cell and molecular biology achieving insight into the nano-dimension. Fundamentally, fluorophores are chemical compounds or molecules, capable of absorbing light at a specific wavelength (excitation light) and emitting light at a different, usually longer, wavelength (emission light) [Croney et al., 2001]. More specifically, the outermost electron orbitals in the fluorophore regulates the wavelengths of absorption and emission, and when absorbing a photon that increases the energy levels from its low-energy ground state, define the potential alteration in any of the electronic, vibrational or rotational states of the

fluorophore [Lichtman and Conchello, 2005]. Fluorescence emission, along with vibrational relaxation, are responsible for the subsequent energy decrease which takes back the fluorophore to the low-energy ground state.

In Figure 2.1.1, the Jablonski diagram [Jablonski, 1933] is presented, displaying the fluorescence process in terms of the fluorophore's alterations in energy levels between different states [Sauer et al., 2010]. The ground energy state of the fluorophore before absorption is S_0 , with S_1 , S_2 representing the singlet excited states after absorption and T_1 the triplet excited state. For the singlet excited states S_1 and S_2 , the outer electron jumps into another orbital, whereas for the triplet state T_1 there is an additional reversal of the electron's spin, that can possibly, however unlikely, occur via inter-system crossing [Turro, 1991]. The bold horizontal lines on each state correspond to the lowest energy levels of each state respectively, with the grey horizontal lines the numerous vibrational levels (rotational levels are not displayed here).




Figure 2.1.1: Jablonski diagram showing the different energy levels of a fluorophore and their association with absorption and fluorescence emission. Reproduced by open access permission, Copyright © Fereja et al. [2013]

During excitation, if the amount of energy absorbed by a fluorophore is higher than needed for exact movement from ground state S_0 to lowest energy in singlet excited state S_1 , the vibration, rotation or orbital states of the fluorophore will be additionally altered (differences in pointing up blue arrows in Figure 2.1.1), with the latter corresponding to the higher energy singlet excited state S_2 (compared to S_1). Consequently, there are different wavelengths to excite a fluorophore, each one associated with alternative orbital, vibrational and/or rotational levels. However, regardless of the energy reached during excitation for the singlet states, the fluorophore's energy will eventually decrease to the lowest vibrational state of S_1 , from which fluorescence emission will occur (pointing down green arrows in Figure 2.1.1). In the case of the unlikely inter-system crossing, the fluorophore in the triplet excited state T_1 can either return to S_0 without emission or by fluorescent emission named phosphorescence (pointing down red dashed arrows in Figure 2.1.1). A more thorough analysis of the fluorescence process can be found in [Lakowicz \[2013\]](#).

2.2 Super-resolution localisation microscopy

Applications of fluorescence microscopy in biology can provide scientists with valuable information about structural details inside cell, tissues and whole organisms [[Meyer and Dworkin, 2007](#)], however, when visible light is used, the spatial resolution is limited due to the diffraction limit of light. Commonly, resolution refers to the level of detail that can be distinguished in an image [[Cox, 2015](#)]. [Abbe \[1874\]](#) was the first to describe the resolution of an optical system as a function of the wavelength of light and the angular objective aperture of the microscope. Abbe's lateral and axial resolution limits are described by the equations,

$$\text{Resolution}_{x,y} = \frac{\lambda}{2n \sin \theta} = \frac{\lambda}{2\text{NA}}, \quad \text{Resolution}_z = \frac{2\lambda}{(n \sin \theta)^2} = \frac{2\lambda}{(\text{NA})^2},$$

where λ is the wavelength of light, corresponding to excitation light wavelength in fluorescence, and NA the objective numerical aperture, with θ and n denoting the

angle and refractive index respectively. These limits correspond to the minimum resolvable distances where closely located light emitting point sources (here fluorophores) can be distinguished. Under ideal conditions when visible light is used, the microscope objective's NA can at most be varied between values of 1.3-1.6, leading to a maximum spatial resolution 200 nm and 500 nm in the lateral and axial planes respectively [Huang et al., 2009].

A number of techniques were developed to increase resolution, achieving the broadly known super resolution. In simple words, this describes imaging better or beyond the diffraction barrier of light. Without much detail, the most popular techniques include structured illumination microscopy (SIM) [Gustafsson, 2005] [Gustafsson et al., 2008], stimulated emission depletion (STED) [Hell and Wichmann, 1994] and saturated pattern excitation microscopy (SPEM) [Heintzmann et al., 2002]. Alternatively and while being of main focus throughout this thesis, a set of techniques referred to as single molecule localisation microscopy (SMLM) super-resolution imaging, rely on single molecules along with a key feature which allows them to alternate their state between 'On' (excited) and 'Off' (ground).

2.2.1 Image formation

Generally, a light emitting point source appears as a diffraction limited spot, or commonly known Airy disc, when imaged by an optical system. In physics, such a spot is characterised by the Point Spread Function (PSF) [Rottenfusser, Wilson, and Davidson], which is the description of the shape of the blur on imaging systems. In our case, such point sources can be represented by fluorophores, since they are only of a few nanometers in size with the property of absorbing light to move into the excited state. This is useful because an image that is formed from the emission of an active fluorophore can be established in a way by the true PSF of the imaging system, initially identified by Ambrose et al. [1991]. In order to obtain spatial information about a fluorophore's location that overcomes the diffraction barrier, the observed PSF is numerically fitted to a hypothetical model PSF, most often of a Gaussian or

Airy shape [Wolter et al., 2011], to estimate the position of its center. We specifically focus on a variety of localisation procedures in Section 2.2.3.

In SMLM, obtaining super resolution is based on a key feature allowing only a sparse subset of fluorophores to stochastically reach the excited state at a specific point in time, with the majority of them remaining at the ground state. This sparsity corresponds to a subset of temporally isolated fluorophores, originally closely distanced and spatially indistinguishable, which are imaged and then localised. Cycles of this process are repeated until numerous fluorophores are sampled, before the majority reach the inevitable photo-bleached state [Egging et al., 1998] which makes them no longer able to fluoresce. A super resolution image of the specimen can be obtained by combining the individual localisations from the acquisition cycles. In Figure 2.2.1 part (a), the SMLM procedure is displayed for a toy ring structure, whereas in part (b) the different modes to achieve this key property of state switching are presented [Dempsey et al., 2011].




Figure 2.2.1: Principle of single molecule localisation microscopy (SMLM) and photoswitching features. (a) A ring diffraction limited structure is resolved using SMLM, by sparsely activating and localising individual fluorophores with high precision. PSFs for active fluorophores are indicated with the red filled in circle with yellow crosses indicating their centred positions. (b) Photoswitching and non-photoswitching features of fluorescent labels. Reprinted by permission from Springer Nature, Dempsey et al. [2011], © 2011 Nature Methods

Probably the most popular SMLM techniques, introduced simultaneously but independently, each one adapting a different labelling approach to obtain photo-switchable fluorophores, are Photo-activated Localisation Microscopy (PALM) [Betzig et al., 2006], Fluorescent Photo-activation Localisation Microscopy (fPALM) [Hess et al., 2006] and Stochastic Optical Reconstruction Microscopy (STORM) [Rust et al., 2006]. As mentioned, all of these methods rely on photoswitching fluorophores (Figure 2.2.1 part (b)), with the choice of appropriate fluorescent labels playing a crucial role. Reviews and performances for a large number of different fluorophores are available in Gould et al. [2009] and [Thompson et al., 2010].

2.2.2 Principles of single molecule localisation microscopy

The gain in spatial resolution is one of the main advancements in single molecule localisation microscopy (SMLM). The main factors contributing to the increase in resolution include fluorophores localisation precision and the density of fluorophores that have been localised on the reconstructed image [Shroff et al., 2008], a term commonly referred to as molecular density. Regarding molecular density, in order to account for increased resolution, the Nyquist-Shannon sampling theorem [Nyquist, 1928] [Shannon, 1949] adequately describes the in-between connection, requiring at least two fluorophores localisation within each resolution unit. In case where labelling is insufficient, super resolution reconstruction of structures may be incoherent presenting unrefined details. In Figure 2.2.2, the effect of molecular density on resolution is presented based on a toy pattern. An insufficient sampling interval can create artifacts on the pattern consisting of discontinuities, thus affecting the spatial resolution. Increasing the number of pixels/lines (left to right) while reducing the number of pixels measured (bottom to top) can lead to an unresolved pattern.




Figure 2.2.2: Spatial resolution against molecular density in SMLM. The figure shows the relationship between the fractions of pixels measured and number of pixels/lines, with increase displayed from top to bottom and left to right respectively. When increasing the number of pixels/lines while also having a small fraction of pixels measured, results in an unresolved structure due to an insufficient sampling interval (Nyquist sampling theorem). Reprinted by permission from Springer Nature, [Shroff et al., 2008], © 2008 Nature Methods

Regarding localisation precision of any individual localisation estimate, corresponding to the position of a fluorophore associated with an uncertainty level, is mainly dependent on the number of photons arriving on the detector's pixels during the fluorophores' emission. We note here that the photon count capabilities can vary according to the fluorescent probe's choice [Dempsey et al., 2011]. Apart from photon detection, which has a stochastic nature therefore considered to form a Poisson process, the size of the detector's pixels and background noise also have a smaller impact however affect the accuracy of localisation, as thoroughly discussed by Ober et al. [2004] and Thompson et al. [2002]. All the concepts that must be taken into consideration for fluorophore localisation are discussed next.

We stress again here that the localisation precision of a localisation procedure determines the level of consistency and reliability of a position estimate, which associated

with molecular density account for the improvement in resolution. Theoretical resolution measures have been already studied [Nieuwenhuizen et al., 2013] [Mukamel and Schnitzer, 2012], however this is an area still open for improvement. Next, we discuss numerous features that contribute to the concept of fluorophore localisation, and make it in some of the cases quite challenging. The main characteristics that need to be taken into consideration regardless of the localisation procedure, are addressed to be (i) *pixelation* (ii) *shot noise* (iii) *background emission* (iv) *single and/or multiple fluorophore(s) acquisition* (v) *PSF misspecification* and (vi) *camera/readout noise*, as discussed in Small and Stahlheber [2014]. All, or most, of them should be taken into consideration when performing a localisation algorithm in order to achieve better accuracy and precision on the fluorophore’s position estimates.

Pixelation Pixelation accounts for any asymmetric allocation of photons if the fluorophore is shifted even by a small fraction from the ideally focused center. Such asymmetry contains any hidden spatial information which plays a crucial role when estimating a fluorophore’s position. More specifically, as analysed by [Small and Stahlheber, 2014], adding the photons arriving anywhere alongside a pixel form the observed signal, which is proportional to the integral of the PSF over the pixel region. If the PSF’s width is noticeably larger than the pixel size, then the PSF value at the center of the pixel can arguably be a good approximation of the observed signal. In contrast, for a larger pixel size, closer to the PSF’s width, the signal must be considered by integrating the PSF over the pixel region. The latter also applies in the case where high precision is of main interest.

Shot noise Probably the most important factor affecting any localisation approach has to be the noise coming from the recorded signal during imaging, referred to as shot noise. A light emitter, usually isotropic, corresponds to the diffraction of light from a focused point source within the specimen, with the associated spatial information being essential for estimating the point’s position. The recorded signal is highly dependent on the amount of photons collected around the area of a point source, a process potentially affected by a number of factors related to the imaging

experiment. As the photon emission has a stochastic nature, translating into their collection inside pixels to be following a random process, a captured image of a light emitter will differ from any chosen or preferred PSF. For instance, one can consider two separate light emitters where the underlying fluorophores have exactly the same position within the specimen. Using the same PSF in both cases, could result in different position estimates due to the stochastic arrival of photons on the detector's pixels. Typically, the photon arrival is described by a Poisson distribution.

Model misspecification Another factor is considered to be the choice of a suitable PSF model basically in terms of structure and shape. A particular choice of a PSF model could provide different position estimates of fluorophores compared to a differently shaped one. In general, misspecification of the PSF model, either in terms of width, orientation, shape or any structural parameters, can potentially lead to inaccurate and/or biased localisation estimates. A comparison of different position estimators, even in misspecified scenarios, is analysed by [Abraham et al. \[2009\]](#).

Background effect An issue highly affecting the distribution of the photons in a captured image, as well adding an additional layer of uncertainty regarding estimation of fluorophore's positions, is scattered light or background effect from fluorophores that are out of focus. Inevitably, due to the process of fluorophores that interchange between states, adds a supplementary level of noise coming from fluctuations in the background. Additionally, this scattered light can have an impact and cause drifts on the average photon position leading to biased position estimates during localisation. A uniform background, translating into the same average of background photons on every pixel, tends to drag the average photon position to the center of the image, whereas a non-uniform background closer to the pixel area with the highest photon count [[Small and Stahlheber, 2014](#)].

Individual and/or multiple fluorophores Arguably one of the most challenging aspects in a localisation scheme is dealing with fluorophores in close proximity, translating into overlapping PSFs on different levels. Identifying positions for over-

lapping PSFs is feasible, however it increases the complexity of the estimation process. Achieving desirable separation of fluorophores is often quite challenging depending on the nature of the experiment, and sometimes can result into photo-toxication of living cells [Cox, 2015]. The capability of a system to capture well isolated fluorophores and not coinciding ones, as well entire fluorophore's activation cycle, is a concept directly associated with the frame rate of the camera used for imaging and the blinking rates of the chosen fluorescent probes, a procedure referred to as acquisition scheme. The most commonly used cameras for localisation microscopy are the Electron Multiplying Charge Coupled Device (EMCCD) and scientific Composite Metal-Oxide Semiconductor (sCMOS) [Huang et al., 2011b], with a comparison of their capabilities on localisation addressed in Quan et al. [2010]. It should be noted here that as the fluorophore activation is a stochastic process, there will always be a probability of spotting light emitters close to each other, resulting in overlapping patches, even under the most sophisticated set ups.

As mentioned, the blinking rate interaction with the camera frame rate have a substantial effect during imaging. The frame rate refers to the capacity and efficiency of the camera in use to capture images of the specimen under investigation. In a sense, the faster the rate the more reliable the images are in terms of collecting every piece of information in an excitation cycle. However, this can arguably not always be ideal for the localisation procedure. The duty cycle is the period of time a fluorophore remains on the excited state and varies according to the particular fluorescent probe. As thoroughly discussed by Dempsey et al. [2011], the most accurate determination of a fluorophore's position would arise from a low duty cycle under a high photon emission, with low duty cycle being a desirable property in general. Regarding single molecule localisation, the total amount of fluorophores able to be localised is inversely proportional to the fluorophore's duty cycle. High photon emission translates into higher density of the focused fluorophore, hence allows for more accurate localisation estimates.

Camera/read noise Regardless of the detector alternatives, mainly consisting of

a camera system capturing blinking cycles of fluorophores, the uncertainty added during signal transformation needs to be addressed. The photon arrival from a focused light emitter on a pixelated region of the detector, forms the known optical signal of a source. Before storing the signal into the pixels, a conversion procedure onto electrical signal occurs, varying with respect to the associated characteristics of the camera. Inevitably, such a process would be subject to some kind of conversion errors, commonly known as read noise, distorting the original signal. In the case where photon emission of a focused fluorophore is high enough without complicated background structure, potentially sufficient for a potential accurate localisation, camera noise could be negligible. However, this ideal scenario is not usually the case in reality.

2.2.3 Current localisation methodologies and algorithms

Precision and accuracy on the position estimates of fluorophores during the stage of frame processing is key to visualise structures and interactions within a cell. Hence, localisation techniques have a very if not the most important role, constituting the basis for inference on super resolution localisation microscopy. Extensive literature exists on localisation methodologies and/or algorithms, constantly expanding at a high rate, which can be classified into five different groups (i) single (ii) multiple light emitters localisation methods via fitting PSF models (iii) single (iv) multiple light emitters localisation algorithms based on deconvolution or other approaches and (v) mechanistic or probabilistic localisation methodologies. Each localisation technique on either of the categories has advantages and drawbacks, with usage choices to be usually related to the nature of the imaging experiment as well as performance on inference.

The basis of the localisation method in terms of considering either single or multiple emitters, can be considered highly dependent on the type of data obtained during the imaging experiment. In principle, SMLM relies on the activation of a sparse subset of photo-switchable fluorophores, hence leading to images of low density. As a result,

the majority of the focused emitters correspond to single fluorophores allowing the use of single molecule localisation methods. However, long acquisition times are required accompanied with limitations on the spatial resolution of the image, not an ideal situation when investigating living cells [Cox, 2015]. On the contrary, increasing the density of active fluorophores (high density data) can allow shorter acquisition times and improvement on spatial resolution, however translating into more challenging localisation procedures due to the presence of mostly overlapping fluorophores.

2.2.3.1 Fitting related localisation approaches

Probably the most broadly used approach in super resolution localisation microscopy follows the concept of fitting a PSF model to the collected data from the focused light emitters. In a way, fitting a PSF requires minimising the distance of the observed signal and the proposed model, evaluated on all possible combinations of its associated parameters.

As described by Small and Parthasarathy [2014], a signal from a single isotropic light emitter can be described using an appropriate PSF model associated with a few parameters,

$$I(x, y) = \underbrace{I_0 h(x - x_0, y - y_0, \tau)}_{\text{expected signal}} + b \quad (2.2.1)$$

where I_0 is a parameter accounting for the detected signal proportional to the photon count (emission) from the focused emitter, b is a background baseline of each pixel, where h is the chosen PSF model with (x_0, y_0) to be denoting the fluorophore's position and τ is a parameter related to the width of the PSF. Arguably, a Gaussian model is the most popular PSF model for an isotropic light source, basically due to its symmetric bell shaped structure. As shown by Stallinga and Rieger [2010] and Wolter et al. [2011], the two-dimensional (2-D) Gaussian model is a good approximation of the true PSF of an imaging system, holding when extending to 3D [Zhang

et al., 2007] , but being noticeably more complicated [Kirshner et al., 2013] . The 2D Gaussian PSF has the general form

$$\text{PSF}_{\text{G-2D}}(x, y | \boldsymbol{\phi}) = \frac{1}{2\pi\sigma^2} \exp \left\{ -\frac{(x - x_0)^2 + (y - y_0)^2}{2\sigma^2} \right\},$$

with the expected signal from Equation (2.2.1) being

$$\begin{aligned} S_{\text{G-2D}}(x, y | \boldsymbol{\phi}) &= I_0 \text{PSF}_{\text{G-2D}}(x, y | \boldsymbol{\phi}) + b \\ &= \frac{I_0}{2\pi\sigma^2} \exp \left\{ -\frac{(x - x_0)^2 + (y - y_0)^2}{2\sigma^2} \right\} + b, \end{aligned} \quad (2.2.2)$$

and $\boldsymbol{\phi} = \{x_0, y_0, I_0, \sigma^2, b\}$.

Localisation of a fluorophore within a focused diffraction limited region can be performed by fitting the PSF model of the form in Equation (2.2.2). The best fit corresponds to the set of evaluated parameters in $\boldsymbol{\phi}$ which minimise the discrepancy between the model and the observed signal on the region. In order to obtain the best fit, the parameters are varied within an optimisation procedure, most frequently carried out through either the Least Squares method (LS) or the maximum likelihood estimation (MLE), measuring the compatibility of the candidate parameter set.

Ways of identifying the diffraction limited patches on a frame, usually described as regions of interest (ROI), vary across every localisation techniques in the literature. This can consist of a threshold, adaptive or not, accompanied with a determined size for each ROI, or a procedure of identifying local maxima based in a filtering approach.

Least-squares (LS) fit The least squares (LS) method basically relies on the minimisation of the so called error terms in order to obtain the best fit. For every pixel in a ROI, the error is defined as the distance of the observed signal and the predicted one according to the chosen PSF model. For a entire ROI, the sum of the

errors is calculated through cycles of varying the parameters until the best fit has been obtained. It should be mentioned that the LS method does not require any specific information about the noise of the imaging system.

Occasionally, the individual error terms are additionally weighted according to their expected errors coming from the fitted PSF model, before adding them all together. The use, or not, of weighting can arguably be a decision based on the chosen PSF model, as a large error does not necessarily have a negative contribution resulting in a poor fit, if the corresponding expected error is also large. On the contrary, if the expected error of the model is small enough even tiny fitting errors have a noticeable negative effect. Therefore, one can claim that the width of the PSF as well the behaviour in the tails of the distribution contribute in a potential weighting decision, particularly useful when the variance is unstable. For instance, when the photon count is considered to follow a Poisson distribution, the expected variance of the errors is the same as the expected signal, and can be approached with a Gaussian distribution for a large count, however this is not the case with lower photon counts [Small and Parthasarathy, 2014]. An early implementation of least squares Gaussian fitting on individual fluorophores was done by Thompson et al. [2002].

Using the 2D Gaussian PSF model, as in Equation (2.2.2), the best fit using the least squares method can be obtained by minimizing with respect to the parameters ϕ ,

$$\hat{\phi} = \arg \min_{\phi} \left\{ \sum_{\forall \{x,y\} \in \text{ROI}} w_{xy} \left(\tilde{S}(x,y) - S_{\text{G-2D}}(x,y | \phi) \right)^2 \right\}, \quad (2.2.3)$$

where $\tilde{S}(x,y)$ denotes the observed signal of pixel (x,y) within the focused region (or the region of interest (ROI)) and w_{xy} the weights of the expected signal, with $w_{xy} = 1$ the un-weighted least squares estimates.

Single fluorophore fitting methods using LS [Tang et al., 2015] comprised a number of steps, involving search for local maxima, smoothing and background esti-

mation on the photo-electrons counts, a transformation of the observed intensities, to identify potential single emitters, followed by a 2D Gaussian PSF fitting using the LS method. A similar approach, named Auto-Bayes, presented by [Tang et al. \[2016\]](#), where an automatic threshold from a constructed histogram of photo-electron count accounts for classification of emitters. In 2D, [Kechkar et al. \[2013\]](#) introduced a real-time localisation algorithm based on a watershed algorithm to extract positions from a wavelet-based segmented reconstruction, using non-linear LS for 3D PSF fitting. In a weighted LS fitting scheme, 3D reconstructions on PSF models extracted using some interpolation techniques was presented by [Baddeley et al. \[2011\]](#).

Multiple fluorophore fitting methods using LS A procedure, known as DAOS-TORM presented by [Holden et al. \[2011\]](#), is a localisation algorithm for multiple fluorophore fitting on high-density super resolution data. Multiple PSF fitting is carried out with the use of weighted least squares, where each pixel is weighted according to the expected photon count of the fitting PSF model. Initial localisation followed by a repetition of multiple fitting on residual images accounts for the final localisation reconstruction. The PSF model is measured and derived from a procedure using a set of individual molecules within an image.

Maximum Likelihood estimates (MLE) In order to obtain a maximum likelihood estimate (MLE), one has to retain specific details about the system under investigation. In the case of localisation microscopy, this information includes a chosen PSF model which describes the signal coming from focused emitter(s) under the presence of noise. The likelihood of the parameters in the PSF model given the observed signal, is associated with the discrepancy between the corresponding model prediction and the observed signal with respect to the noise system. Conceptually, small distances will lead to a higher likelihood than any noticeable mismatches associated with large differences. Maximisation of the likelihood with respect to the model parameters determines the MLE, as in the best fit of the chosen model to the signal.

From the statistical and mathematical perspective, the MLE estimate has some effective properties. Different unbiased estimates can exist, however any one of them has a minimum achievable variance, or maximum precision similarly, which cannot be necessarily attained. This variance limit for unbiased estimates is known as the Cramer-Rao lower bound (CRLB) [Cramér, 2016] [Rao, 1992], with the MLE being the unbiased estimates achieving maximum precision for a large number of observations, when this limit is achievable [Kay, 1993]. Ober et al. [2004] introduced an approach for determination of MLE's localisation accuracy with respect to fluorophore's emission and the imaging system, whereas a comprehensive comparison between the LS and MLE fitting approaches was implemented by Abraham et al. [2009].

In the general case where the photon count captured by the camera is considered to arise from a Poisson distribution, accounting for shot noise as described earlier, with the expected photon count to be described by a 2D Gaussian PSF model as in Equation (2.2.2), the likelihood of the parameters ϕ for $\forall\{x, y\} \in \text{ROI}$ reads in

$$\mathcal{L}(\phi \mid \{x, y\} \in \text{ROI}) = \prod_{\forall\{x, y\} \in \text{ROI}} \frac{S_{\text{G-2D}}(x, y \mid \phi)^{\tilde{S}(x, y)} \exp\{-S_{\text{G-2D}}(x, y \mid \phi)\}}{\tilde{S}(x, y)},$$

with $\tilde{S}(x, y)$ being the observed signal. The MLE of ϕ is acquired by maximizing $\mathcal{L}(\phi \mid \{x, y\} \in \text{ROI})$ with respect to ϕ , or more often done using the log-likelihood, hence

$$\begin{aligned} \hat{\phi} &= \arg \max_{\phi} \ell(\phi \mid \forall\{x, y\} \in \text{ROI}) \quad \text{with} \\ \ell(\phi \mid \forall\{x, y\} \in \text{ROI}) &= \sum_{\forall\{x, y\} \in \text{ROI}} \tilde{S}(x, y) \ln \left(S_{\text{G-2D}}(x, y \mid \phi) \right) - S_{\text{G-2D}}(x, y \mid \phi). \end{aligned} \tag{2.2.4}$$

Single fluorophore fitting methods using MLE Numerous localisation methods exist for estimating fluorophore’s positions using the MLE on regions of interest which contain single fluorophores. [Smith et al. \[2010a\]](#) implemented a 2D (and 3D) Gaussian MLE localisation using the integrated expected signal of $S_{G-2D}(x, y | \phi)$, formulated with the Gaussian error functions, on pre-selected ROIs. An iterative mechanism using the Newton-Raphson method [[Seiler and Seiler, 1989](#)] accounted for maximising the likelihood. The separable property of the Gaussian PSF model allowed an increased localisation speed when estimation performed also with the integrated 2D Gaussian PSF [[Starr et al., 2012](#)]. [Brede and Lakadamyali \[2012\]](#) introduced an open-source software consisting of a variety of features for rendering and calibrating high resolution images. The main processing method includes identification of single molecule ROIs based on a threshold, before Gaussian PSF MLE fitting on a Poisson noise model is performed. Smoothing the frames before determination of the local maxima, hence the ROIs, followed by fitting Gaussian PSF in both 2D and 3D for single molecule localisation using either MLE or LS, was proposed by [Wolter et al. \[2012\]](#).

Multiple fluorophore fitting methods using MLE When high density frames are obtained from an super resolution imaging experiment, the probability of having multiple fluorophores on a ROI is considerably high, as a result of overlapping PSFs, hence the need of localisation procedures which can simultaneously deal with multiple emitters. [Huang et al. \[2011a\]](#) presented an extension of an MLE fitting procedure for multiple light emitters, using the integrated expected signal for a single fluorophore based on the Gaussian error functions [[Smith et al., 2010a](#)], implemented on a GPU hardware allowing for fast processing. [Quan et al. \[2011\]](#) introduced an algorithm in which the number of active molecules in a focal plane are initially roughly identified by the use of the Structured Sparse model, before using the BIC criterion in order to select an optimal model, potentially free of false positives, from multiple fitted MLE positions. Using GPU parallel computing, the algorithm was extended and streamlined into a high density localisation method, known as PALMER [[Wang](#)

[et al., 2012](#)]. An algorithm for processing high density super resolution images in 3D, allowing fit of overlapping PSFs of different shapes, modelled as elliptical Gaussian PSFs in which flattening of the ellipsoids depends on the z position of the focused emitters, was introduced by [Babcock et al. \[2012\]](#). [Huang et al. \[2013\]](#) presented a localisation algorithm based on multi-emitter fitting using MLE, where the super resolution images were obtained with a sCMOS camera [[Huang et al., 2011b](#)].

ThunderSTORM [[Ovesný et al., 2014](#)], an increasingly popular open source software, consists of a large variety of features and different single molecule localisation methods, similar to the ones described throughout the section. It offers a number of choices for processing SMLM frames both in 2D and 3D, along with a collection of options for visualisation and analysis in a post-processing manner.

2.2.3.2 Localisation algorithms based on deconvolution or other techniques

Apart from localisation methods based on a PSF model fitting approach, either with the use of MLE or LS, a variety of algorithmic approaches have been developed in order to process both low and high density super resolution images. Following a similar structure as for the fitting approaches, we classify the algorithms into single and multiple fluorophores approaches, accounting for low and high density super resolution data respectively, while also having a group for mechanistic approaches. The latter involves methods attempting to model the fluorophore’s dynamics and properties in a probabilistic manner.

Single fluorophore localisation algorithms A fast algorithmic method adopting a modified center of mass algorithm for parameter estimation on focused ROIs and position identification, capable of implemented both in 2D and 3D, presented by [Henriques et al. \[2010\]](#). [Parthasarathy \[2012\]](#) used the fact that a line of any point obtained parallel to that point’s intensity gradient will cross the molecule’s true position, if the intensity distribution is radially symmetric. For noisy measurements,

the true molecule's position can be derived by minimising the distance of true position and every corresponding line. In a similar manner, by using the radial symmetry, [Ma et al. \[2012\]](#) calculated the gradient using a function formed by the image gradients on the x and y directions, obtained as the convolution of the image and gradient operators, followed by a minimisation step to obtain the best fit for the center of symmetry. An non-numerical fitting algorithm [[Andersson, 2008](#)] uses the location of a measurement, which is the center of a pixel for a EMCCD camera, and the potential true position of the molecule to obtain a range equation of the true source. To determine the estimate of the true position of the source, an approximate least squares solution is derived from the range equation of multiple measurements. [Yu et al. \[2011\]](#) used the intensity spectrum derived with a zero padded Fourier transform to model the intensity of a source, and estimated the position of the molecule using the approximate phase shift of the intensity spectrum. A reconstruction algorithm, obtaining molecule's coordinates of local maxima using cubic spline interpolation on a self calibration set up, followed by a statistical test to get fluorescence patches introduced by [Köthe et al. \[2014\]](#).

Multiple fluorophore localisation algorithms CSSTORM [[Zhu et al., 2012](#)], an abbreviation for compressed sensing for STORM images, relies on linear programming to perform the localisation procedure. Instead of coordinates, a discrete sub-pixel grid accounts for the molecule's position from a camera's frame signal with each sub-pixel representing the brightness of a molecule positioned within. A convex optimisation problem is considered by favouring sparsity priors using a weighted L1 norm under a Poisson noise constraint. Several patches are analysed and added together to produce the super resolution image. A similar first step was adapted by [Min et al. \[2014\]](#) to obtain a sparse image, followed by deconvolution on a fixed spatial support to correct pixel values. Their algorithm, known as FALCON, uses a Taylor approximation for the PSF to refine the initial positions over a continuum, obtained using local maxima from the deconvolution. [Mukamel et al. \[2012\]](#) created an algorithm named deconSTORM, which uses a stack of frames on

a adjusted iterative Lucy-Richardson deconvolution algorithm based on a Markov process accounting for molecules state transition throughout the stack. Using the Expectation-Maximization (EM) algorithm, the maximum likelihood sample estimate without explicit molecule localisation is obtained, using the sparsity property of activated fluorophores within a frame. Takeshima et al. [2018] developed a multi-emitter fitting technique based on a wedge-shaped template matching algorithm, hence the name wedge template matching WTM, in order to localise molecules with overlapping PSFs. Template matching is the procedure of matching sectors of an observed image to a template image, whereas WTM depends on partial feature-based templates. Accounting for both the spatial and temporal aspects of fluorophores, Cox et al. [2012] developed a Bayesian approach for high density data by taking into consideration the natural fluorophore processes of blinking and bleaching. Modelling the entire stack simultaneously, a density map is obtained according to conditional probabilities, denoting the likelihood of molecules existing on the spatial domain of an image. The super resolution image is constructed by averaging the density maps opposed to the most popular super resolution reconstructions using localisations of point estimates.

2.2.3.3 Quantitative imaging

The importance of localisation microscopy is indisputable since it allows visualisation and analysis of both structures and molecular behaviour within living cells. Besides the qualitative evaluation of cells, knowledge on quantifiable features is equally essential and can provide substantial information about intermolecular evolution and interaction. A valuable component in quantitative imaging relates to molecular counting, which accounts for accurate identification of the number of molecules on a specimen under investigation. In a way, super resolution localisation microscopy provides such a measurement, by relying on mostly well isolated fluorophores imaged during a stochastic activation process over a period of time, however a number of factors can affect and complicate counting. Some of the main challenges are associated

with multiple localisations of fluorophores that undergo blinking (over-counting), missed active fluorophores due to low signal (under-counting) or generally inaccuracy on the localisation outcome, thoroughly discussed by [Shivanandan et al. \[2014\]](#) and [Deschout et al. \[2014\]](#).

In the past years, quantitative imaging is at the forefront of interest along with the development of localisation procedures, and here we outline a few of the approaches. [Rollins et al. \[2015\]](#) introduced a stochastic approach by adapting an aggregated Markov model to study the blinking dynamics. Instead of a temporal threshold to distinguish between re-activated fluorophores the kinetic rates are an output of the process, estimated using the maximum likelihood. [Hummer et al. \[2016\]](#) obtained a functional form for the number of blinking events independent of the photophysics, called as model-independent counting procedure. Their approach does not require any specific knowledge about the photophysical properties such as states or kinetic rates, with the corresponding effects described by at most three parameters in the model, depending on the overlapping (or not) of fluorophores. [Nino et al. \[2017\]](#) considered the distribution of blinking events from a single fluorophore, and used Bayesian analysis to obtain an estimate for the number of fluorophores conditional on the observed blinks. An analytical approach providing quantitative evaluation of super resolution images introduced by [Culley et al. \[2018\]](#). Their algorithm, named as SQUIRREL, requires a reference diffraction limited image associated with a corresponding super resolution one and a representative resolution scaling function image. Under the assumption of a spatially invariant PSF on the observed image, convolution of the super resolution image with the representative resolution function produces an image similar to the observed diffraction limited image. These diffraction limited images are used to construct a map of errors which can indicate and identify defects. Although these approaches provide alternatives to quantitative assessment in super resolution imaging, this is an area still open to improvement, with molecular counting being of particular importance.

2.3 Motivation of this research

Super resolution localisation microscopy is capable of allowing inference up to the molecular scale, a case not feasible with common fluorescence imaging alternatives, thus processing methods are in the forefront of interest. Knowledge on the number and positions of differently labelled molecules within cells can provide substantial information and give valuable insight not only about cellular structure, but also about functionality, clustering and molecular behaviour and interaction. Therefore, processing schemes should be considered within the area of quantitative imaging, where localisation procedures can provide both quantifiable features and measurements apart from visual super resolution reconstructions. From a mathematical point of view, one can argue that processing super resolution images involves analysing the spatial association of molecules within a frame as well as the temporal aspects arising from the nature of the experiment, including the fluorophore's blinking and bleaching attributes.

Development of localisation methods showed increasing attention since the evolution of super resolution localisation imaging. As analysed in Section 2.2.3, a vast variety of alternatives exist for either fitting PSF models to single and/or multiple light emitters or non-fitting algorithmic processing approaches. The majority of the fitting based techniques use the MLE approach, proved to be the most accurate tool for single molecule localisation in terms of precision. The main requirement is knowledge of the noise model which associated with a respective PSF accounts for modelling the light emission of a focused source. The noise model can be usually adjusted to the properties of the imaging system, however the large computational times needed for these type of approaches cannot be avoided. This can be a crucial issue when it comes to imaging living organelles. On the other hand, non-fitting approaches can substantially improve on computing times, however the fact that there is no explicit determination of a noise model justifies the decreased reliability on the molecule's localisation. In a sense, a combination of the positive attributes from these two

alternatives would be ideal, but unfortunately a potential not necessarily feasible.

One of the main challenges that arise with high density super resolution data is the higher probability of having overlapping fluorophores, that is closely distanced molecules having their PSFs coinciding on some level. On high density data, a larger amount of molecules is allowed to reach their emission state on a single time point, a matter that can be either controllable or enforced necessity. Depending on the nature of the experiment and the sample under investigation, the labelling strategy might differ, in some of which the ideal sparse fluorophore blinking can be limited, forming denser fluorescent frames. In such cases, the chosen model should be able to accommodate overlapping situations. A few methods have been introduced based on fitting techniques or alternative approaches as discussed on the previous section, however this is still an area of development and improvement based on the challenging nature of the process. In a way, we want a method which can allow localisation based on the number of photons, translated into registered intensities, of fluorophores from overlapping PSFs, while taking into consideration the relative noise levels of the system.

Most of the state-of-the-art alternatives rely their localisation on single molecules, hence they aim to identify regions of well separated fluorophores before a localisation scheme is applied. Sometimes, these algorithms focus on pixels with the higher photon count, as potential pixel containing the molecules, without taking into account the possibility of this pixel being formed from overlapping PSFs. In some cases, regions of potentially overlapping fluorophores are discarded, due to the incapability of the localisation scheme to be reliable and accurate. Therefore, apart from a procedure capable of dealing with multiple emitters, a sophisticated way of identifying the regions of interest (ROI) is of high importance. In order to isolate regions of interest, the majority of methods adopt a fixed mechanism without taking into consideration the uncertainty within the frame. This can either involve an adaptive, or not, threshold scheme or deterministic choices and formation of windows using center of mass or brightest pixels, sometimes associated with a filtering procedure.

Furthermore, only a few of the current techniques take into consideration the time dynamics of the fluorophores associated with a SMLM experiment. These dynamics correspond to the blinking cycles, where a fluorophore can interchange between excited and dark states before exhaustion, leading to the irreversible photo-bleached state. In a sense, the realisation of whether or not two closely located light emitters captured on separate time points, or different frames within the stack, represent the same molecule undergoing a blinking cycle, is key if molecule counting is needed. Quantitatively, counting relates to gaining knowledge on the distribution of identified molecules within cells as well the scale of the entire population, but also contributes to acquiring substantial information about intermolecular evolution and interactions. Hence, my approach aims at building a model that takes into account the diffraction of light around an active fluorophore and the possible interactions with those close by, with the capability to perform both qualitative and quantitative inference.

Chapter 3

Frame-wise image analysis

3.1 Introduction

According to SMLM, an individual image of a specimen under investigation consists of a sparse number of active molecules, i.e. molecules on the light emission state, as discussed in Section 2.2. These molecules cannot be directly identified and localised, since their emission is distributed around a neighbourhood of their true positions, where further complications may arise if these are in close proximity to each other having their neighbourhoods overlapping. Regardless of the camera system, every captured image of the specimen is stored to a pixel-based frame, consisting of measurements from corresponding transformations of the recorded signal.

In this chapter, we present a complete Bayesian model which takes into account the spatial association of an active molecule within a frame and is capable of resolving closely located, or overlapping, molecules. In Section 3.2, we introduce the novel Single Molecule Pattern (SiMPa) functions, a set of functions describing the detected emission pattern of an emitting molecule in a predetermined pixel region, based on the position of the molecule within a pixel. A detailed discussion on the derivation of the functions is provided, as well a proposed way of their application, with the section

concluding on a pseudo-algorithm of single frame generation using the SiMPa functions (Section 3.2.5). In Section 3.3, we present our individual frame inference based on the SiMPa functions, via a Bayesian approach to allow coherent propagation of uncertainty (Section 3.3.2). Inference on positions of the molecules within a frame is performed through our novel probabilistic pattern-configuration-realisation scheme (Section 3.3.3), with the entire Bayesian inference scheme for individual frames summarised on a pseudo-algorithm at the end of the section (Section 3.3.5). The chapter concludes with an implementation of our individual frame inference scheme on a single simulated frame, in Section 3.4, where the output of our procedure is displayed and analysed.

3.2 Single Molecule Pattern (SiMPa) functions

In this section, the novel Single Molecule Pattern, or SiMPa, functions are introduced. The need of a tool to describe the behaviour of the detected intensities pattern, or intensities diffraction, when a molecule is on the active state, i.e. when it is emitting light, is of high importance. These intensities represent the transformed signal of light as recorded by the corresponding imaging system. Naturally, an individual intensities diffraction is a blurry representation of an active molecule, or fluorophore, working as an emitting point source. The diffraction, commonly described by the Point Spread Function (PSF), does not immediately allow identification of the precise position of the point source and especially in the case in which two, or more, fluorophores are really close to each other having their PSFs overlapping, this becomes even more challenging. The aim of the SiMPa functions is to provide a formal representation for the behaviour of the PSF of a fluorophore, according to its relative position within the lattice. This representation is then used in order to extend this procedure to more difficult scenarios.

The SiMPa functions are a set of functions that describe the intensities diffraction of an active molecule across a predetermined pixel region, based on its position within

the central pixel of that region. The main, naturally derived, assumption of the SiMPa functions comes from the fact that a fluorophore will be positioned somewhere within a pixel, as the imaging procedure relies on a pixel grid to record emission. Let a specific active molecule exist somewhere within a pixel. We define an origin within that pixel to be the center of it. The molecule's position is determined by introducing two different parameters, an angle θ and a distance r from that origin. These two parameters define the exact position of the molecule within that pixel with respect to the predetermined origin and are two of the key components that will control the behaviour of the intensities diffraction. This simply means that different positions of a molecule within a pixel will result in different spreads across the same fixed region. These pixel regions or neighbourhoods, as commonly known in lattices, will contain all the necessary information needed in order to describe the intensity emission of a molecule on the spatial domain. In the simplest case of the SiMPa functions, these neighbourhoods obey the second order neighbourhood structure consisting of a 3×3 area of pixels. We define $\alpha_k(r, \theta, c) \equiv \alpha_k$, to be the SiMPa functions describing the intensities diffraction on a predetermined neighbourhood, with k being a neighbouring pixel's index, $k = 0, \dots, 8$. The SiMPa functions $\alpha_k(r, \theta, c)$ depend on the distance r on direction θ the molecule lies away from the origin, as well as on a parameter c accounting for the power of the diffraction on the entire neighbourhood.

A graphical representation of how the SiMPa functions behave within the neighbourhood is presented in Figure 3.2.1. On the left panel of the figure a lattice of pixels is drawn representing an image of a focused specimen. A number of molecules, represented by the black dots, have been placed within pixels under different combinations of r and θ in accordance with their corresponding origin. These molecules are wrapped with black circles which denote the spread of the intensities when they are active, or emitting light. A specific region has been marked, with the dashed black lines, showing the diffraction along the pixel-wise neighbourhood of a molecule based on the 3×3 pixel area. The central pixel of the neighbourhood, in which the molecule lies onto, is the pixel marked with the thick black lines. The position

of the molecule using the distance r and angle θ from the predetermined origin, is presented on a magnified version on the middle part of the figure. The blue solid lines d_1 and d_2 are the fixed distances from the origin and l is the side length of the pixel. These quantities will contribute to the construction of the SiMPa functions which are shown on the right panel of the figure along with the corresponding indexing for each pixel in the neighbourhood.




Figure 3.2.1: Graphical representation of the Single Molecule Pattern (SiMPa) functions. On the left had side of the figure a lattice of pixels is shown. The black dots within, denote different molecules with their diffraction denoted by the surrounding black circles. The dashed lines denote the neighbourhood of a molecule placed within the bold pixel at the center of it. The middle part of the figure contains the center pixel of the neighbourhood showing the position of the molecule with respect to the distance r and angle θ . The fixed parameters d_1 , d_2 and l are the two fixed distances from the origin and the length of the side of the pixel respectively. The right hand side of the figure illustrates the spatial configuration of the SiMPa functions of that molecule.

The SiMPa functions, as introduced, denote how the intensities of an active molecule are distributed across a predetermined neighbourhood (3×3), according to the molecule's position within the central pixel. They are derived in a structural way, by taking into consideration the molecule's distance from the nearest edge of each one of the neighbouring pixels, while a variety of quantities contribute in order to obtain the appropriate behaviour. We will require the SiMPa functions to be bounded between

0 and 1, a choice based on our proposed scheme to use them. The boundaries can be acquired by creating an upper limit for the power of spread parameter c for all possible combinations of r and θ . Every aspect will be explained in detail throughout this chapter.

Firstly, we define the SiMPa functions $\alpha_k(r, \theta, c) \equiv \alpha_k$, $k = 0, \dots, 8$ as

$$\alpha_k = \begin{cases} c \left(d_1 - r \cos \left[\frac{(k-1)\pi}{4} + \theta \right] g(R, l) \right), & k = 1, 3, 5, 7 \\ c \left(\left((rg(R, l))^2 + d_2^2 + 2d_2 s_k r \cos \left(\frac{(k+1)\pi}{4} - \theta \right) g(R, l) \right)^{1/2} \right), & k = 2, 4, 6, 8, \\ 0, & k = 0 \end{cases} \quad (3.2.1)$$

where l is the side length of the pixel, where without loss of generality $l = 1$ units, and the constants d_1 and d_2 are the fixed distances of the neighbouring pixels from the origin with $d_1 = \frac{l}{2}$ and $d_2 = \frac{l}{\sqrt{2}}$. These quantities can be graphically seen in the middle part of Figure 3.2.1 along with distance r and angle θ . The choice of the indices is based on the fixed distances d_1 and d_2 from the predetermined origin. The even numbers $k = 2, 4, 6, 8$ represent the corner pixels of the neighbourhood based on the central pixel, which is indexed with $k = 0$, while the odd numbers $k = 1, 3, 5, 7$ denote the adjacent pixels to the central one (Right panel of Figure 3.2.1). The polar coordinate system is used to measure the molecule's position under the appropriate restrictions. As r is a distance it should be positive, $r \geq 0$, and due to the fact that we allow the molecule to lie anywhere within the pixel, every combination of r and θ should satisfy the inequalities $r|\cos(\theta)| < \frac{l}{2} - R$ and $r|\sin(\theta)| < \frac{l}{2} - R$ where R is the radius of the molecule. These restrictions imply that the distance r a molecule can be placed on direction θ is bounded by the size of the molecule, measured by its radius R , with respect to the pixel's side length.

Regarding the quantity R , we make the assumption that the projection of the

molecule on the 2-dimensional space is a circular object, with the parameter R denoting its radius. We define a function $g(R, l)$ which accounts for identical performance of the SiMPa functions for different choices of the size of the molecule. It is multiplied with the chosen distance r in order to achieve this similarity. This function, $g(R, l)$, takes into account the relative side length of the pixel l and the radius of the molecule R and defined as $g(R, l) = \frac{1}{1 - 2R/l}$. When the size of the molecule increases (bigger radius) the distance r it can lie within the pixel decreases and $g(R, l)$ works as a correction mechanism to obtain the same diffraction along the neighbouring pixels. A more detailed discussion can be found in Section 3.2.3.

Based on the position of the molecule within the central pixel its minimum distance from the neighboring pixels will be obtained. For our calculation without loss of generality, we assume that the molecule is a single point, i.e. $R = 0$, therefore it is straightforward to get $g(R, l) = 1$. Let the quantities \textcircled{A} and \textcircled{B} from the SiMPa functions in Equation (3.2.1) to be:

$$\begin{aligned} \textcircled{A} &= d_1 - r \cos \left[\frac{(k-1)\pi}{4} + \theta \right] \\ \textcircled{B} &= \left(r^2 + d_2^2 + 2d_2 s_k r \cos \left(\frac{(k+1)\pi}{4} - \theta \right) \right)^{1/2}. \end{aligned} \quad (3.2.2)$$

In order to construct the SiMPa functions the minimum distances of the molecule from the neighboring pixels need to be obtained. These distances, presented in Figure 3.2.2, correspond to the closer point of a neighbouring pixel to the molecule. On the left hand side of the figure the minimum distances of the molecule from the adjacent pixels to the central one are shown, while on the right panel of the figure the minimum distances of the molecule from the corner pixels the central one. For the adjacent pixels with $k = 1, 3, 5, 7$, we calculate the distances l_1, l_3, l_5 and l_7 , as shown in the Figure 3.2.2 (a), by using polar coordinates and basic trigonometric identities and shifts. The minimum distances on that case are the ones perpendicular to the

pixel, therefore:

$$\begin{aligned}
 l_1 &= \frac{l}{2} - r \cos(\theta) \\
 l_3 &= \frac{l}{2} + r \cos(\theta) = \frac{l}{2} - r \sin(-\theta) = \frac{l}{2} - r \cos\left(\frac{\pi}{2} + \theta\right) \\
 l_5 &= \frac{l}{2} + r \cos(\theta) = \frac{l}{2} - r \cos(\theta + \pi) \\
 l_7 &= \frac{l}{2} - r \sin(\theta) = \frac{l}{2} - r \cos\left(\frac{\pi}{2} - \theta\right) = \frac{l}{2} - r \cos\left(\frac{3\pi}{2} + \theta\right).
 \end{aligned}$$




Figure 3.2.2: Minimum distances of a molecule placed r units with direction θ away from the origin from the neighbouring pixels. The distances are denoted with the blue solid lines. (a) Minimum distances of the molecule from the adjacent pixels to the central pixel, l_1 , l_3 , l_5 and l_7 . (b) Minimum distances of the molecule from the corner pixels of the neighbourhood, l_2 , l_4 , l_6 and l_8 .

Putting all of them together by allowing the index of each distance to be k , with $k = 1, 3, 5, 7$, we obtain $l_k = d_1 - r \cos\left[\frac{(k-1)\pi}{4} + \theta\right] \equiv \textcircled{A}$ in Equation (3.2.2), with $k = 1, 3, 5, 7$ and $d_1 = l/2$.

In a similar manner we derive the distances l_2, l_4, l_6 and l_8 for the corner pixels with $k = 2, 4, 6, 8$, as shown in Figure 3.2.2 (b) using polar coordinates. We apply Pythagoras theorem along with the sum to product identities for the cosine functions and some basic co-functions. Only the calculation of the distance l_2 is shown while the calculations of l_4, l_6 and l_8 follow similar steps.

$$\begin{aligned}
l_2^2 &= \left(\frac{l}{2} - r \cos \theta\right)^2 + \left(\frac{l}{2} + r \sin \theta\right)^2 \\
&= \frac{l^2}{2} + r^2 - rl(\cos \theta - \sin \theta) \\
&= \frac{l^2}{2} + r^2 - rl \left(\cos \theta - \cos \left(\frac{\pi}{2} - \theta\right)\right) \\
&= \frac{l^2}{2} + r^2 + 2rl \sin \left(\frac{\pi}{4}\right) \sin \left(\theta - \frac{\pi}{4}\right) \\
&= \frac{l^2}{2} + r^2 + \sqrt{2}rl \cos \left(\frac{3\pi}{4} - \theta\right).
\end{aligned}$$

Therefore, $l_2 = \left(r^2 + d_2^2 + 2d_2r \cos \left(\frac{3\pi}{4} - \theta\right)\right)^{1/2}$, with $d_2 = l/\sqrt{2}$. Similarly,

$$\begin{aligned}
l_4^2 &= \left(\frac{l}{2} + r \sin \theta\right)^2 + \left(\frac{l}{2} + r \cos \theta\right)^2 \Rightarrow l_4 = \left(r^2 + d_2^2 + 2d_2r \cos \left(\frac{\pi}{4} - \theta\right)\right)^{1/2} \\
l_6^2 &= \left(\frac{l}{2} + r \cos \theta\right)^2 + \left(\frac{l}{2} - r \sin \theta\right)^2 \Rightarrow l_6 = \left(r^2 + d_2^2 - 2d_2r \cos \left(\frac{3\pi}{4} - \theta\right)\right)^{1/2} \\
l_8^2 &= \left(\frac{l}{2} - r \cos \theta\right)^2 + \left(\frac{l}{2} - r \sin \theta\right)^2 \Rightarrow l_8 = \left(r^2 + d_2^2 - 2d_2r \cos \left(\frac{\pi}{4} - \theta\right)\right)^{1/2}.
\end{aligned}$$

Similarly, we adopt a common index k with $k = 2, 4, 6, 8$, which sums up these functions into $l_k = \left(r^2 + d_2^2 + 2d_2s_kr \cos \left(\frac{(k+1)\pi}{4} - \theta\right)\right)^{1/2} \equiv \textcircled{\text{B}}$ in Equation (3.2.2), with $k = 2, 4, 6, 8$ and $d_2 = l/\sqrt{2}$. The quantity s_k appropriately alters the sign of the cosine function with $s_k = +1, -1, +1, -1$ for $k = 2, 4, 6, 8$ respectively.

The parameter c controls the power of the molecule's diffraction. This practically means that it describes the level of brightness or darkness the entire neighbourhood will have. High values of c imply that the neighbouring pixels will have higher values, since $\alpha_k(r, \theta)$ is boosted when multiplied with c , hence the SiMPa functions have larger influence than when c is small. In terms of intensities as we discuss next, large values of c correspond to a faster intensity decay within the neighbourhood compared to smaller ones. The SiMPa functions are bounded between 0 and 1, achievable by bounding c with $0 < c \leq 1/2d_2$, when $R = 0$ without loss of generality. The derivation of the boundaries can be found in Section 3.2.4.

Next, we discuss our proposed way to apply the SiMPa functions. This involves the introduction of a parameter I , which we call the single event intensity. This parameter has been considered based on the natural implementation of the imaging technique. During the set up of the imaging procedure a baseline value is chosen to be the background intensity. When a fluorophore is active, its emission intensity can be related to the background one in the sense of separating the purely background values from the higher ones. Practically, we consider the emission intensity to be proportional to the background or vice versa. The emission intensity is what we call the single event intensity which we denote with I .

The SiMPa functions, α_k , $k = 0, \dots, 8$, are used with respect to the single event intensity I in a way that they specify the appropriate proportions of I each pixel in the neighbourhood should have. This will be altered when different positions of a molecule are applied, according to distances r and angles θ from the predetermined origin. By construction the SiMPa functions are bounded between 0 and 1, $0 \leq \alpha_k(r, \theta, c) \leq 1$, where values close to 0 imply a high influence from the choice of r and θ and values close to 1 a small one. Translating this influence into intensities we allow pixels highly affected by (or closer to) the position of the molecule to have a higher intensity than the pixels less influenced (or further away). Therefore, the appropriate proportions of the single event intensity I for each one of the pixels in the neighbourhood according to the SiMPa functions are summarised by the quantities

$(1 - \alpha_k(r, \theta, c))I$, $k = 0, \dots, 8$. An illustration for a single molecule's diffraction is presented in Figure 3.2.3. On the left hand side of that figure a molecule is placed on a distance r_1 on direction θ_1 from the origin. On the right panel of the figure, the SiMPa functions are evaluated based on its position, creating the intensities diffraction for the entire neighbourhood with respect to the single event intensity I .




Figure 3.2.3: Proposed use of SiMPa functions. On the left hand side of the figure a molecule has been placed on a distance r_1 with angle θ_1 from the origin. On the right hand side an implementation is presented, with the SiMPa functions constructing the intensities diffraction with respect to r_1 and θ_1 and the single event intensity I .

A numeric illustration of the proposed use of the SiMPa functions according to the single event intensity I can be seen in Figure 3.2.4. A number of r and θ combinations show the effect of altering the intensities throughout the neighbourhood. On part (a) of the figure, a molecule is placed at the center of the pixel, origin of the SiMPa functions with $r = \theta = 0$, showing the symmetric diffraction along the neighbourhood. Regarding a different combination of r and θ , on part (d) for instance, a molecule is placed almost all the distance away from the origin on direction $\theta = 3\pi/2$. Comparing to part (a), the closest neighbouring pixels are boosted, in term of having higher intensity values, whereas the furthest away ones consist of lower valued

intensities.




Figure 3.2.4: Evaluation of the SiMPa functions for different combinations of r and θ . The side length of the pixel is $l = 1$ units which mean that each pixel has dimension 1×1 . The radius R of the molecule is $R = 0.02$, the single event intensity is $I = 100$ and the parameter $c = 0.4$ which controls the power of the diffraction. The results from the SiMPa functions are rounded to the closest integer. Four different positions of the molecule are used (a) The molecule is centred at the origin of the pixel, $r = 0$ and $\theta = 0$. (b) The molecule is placed on a distance $r = 0.45$ with angle $\theta = \pi/4$. (c) The molecule is placed on a distance $r = 0.35$ with $\theta = 9\pi/8$. (d) The molecule is placed on a distance $r = 0.45$ with $\theta = 3\pi/2$.

3.2.1 Special cases of the SiMPa functions

In this section we present the special cases of the SiMPa functions regarding positions of molecules that do not fall within a pixel. We address two individual scenarios,

which can occur when the true position of the molecule is in correspondence with one of the boundaries of the SiMPa functions. The first case corresponds to the molecule being located on an edge of the central pixel, while the second lying exactly on an intersection of two edges. The former accounts for a molecule of radius R having its center precisely placed on an edge of the central pixel, while the latter exactly on the intersection of two edges. Both cases will be presented in detail, however our proposed inference, presented later in Section 3.3, relies only on the general form of the SiMPa functions in Equation (3.2.1). The main reason for this decision is the fact that the behaviour of the diffraction in either of these cases is embedded within the general SiMPa functions if the maxima for the distance r is chosen for the appropriate directions, that is $r = (l/2 - R)/|\cos\theta|$ or $r = (l/2 - R)/|\sin\theta|$. Additionally, the probability of having a molecule at exactly these points of the continuum can be considered practically zero.

Molecule on an edge of the central pixel

Here, we discuss the case where the molecule lies on one of the edges of the central pixel of the general SiMPa functions. When this scenario takes place, the intensities diffraction is symmetrically distributed with respect to the edge. On Figure 3.2.5 a graphical representation can be seen, including all the quantities needed in order to define the functions. The intensities diffraction is considered to spread out on the twelve neighbouring pixels, hence the indexing $k = 1, \dots, 12$, with the functions denoted with $\alpha_k^E(r, \theta, c) \equiv \alpha_k^E$. On the left panel of the figure, a molecule has been placed on the origin of this special case which has been defined to be the median of the edge. The molecule can lie anywhere across the edge therefore the only possible directions with respect to the origin are $\theta = \frac{\pi}{2}$ or $\theta = \frac{3\pi}{2}$. On the right panel, an illustration of a molecule is presented, placed on a distance r_1 from the origin with direction $\theta_1 = \pi/2$.




Figure 3.2.5: Special case when a molecule lies on one of the edges of the central pixel of the SiMPa functions. On the left panel of the figure, the distribution of the functions is shown along with the fixed distances from the predetermined origin, whereas on the right an illustration is shown with a molecule placed at a distance r_1 on direction $\theta_1 = \pi/2$ from the origin.

The functions $\alpha_k^E(r, \theta, c) \equiv \alpha_k^E$, $k = 1, \dots, 12$ describe the way the intensities of the k neighbouring pixels will be distributed with respect to the molecule's position along the edge. As introduced in the SiMPa functions, the position depends on a distance r under two possible directions $\theta = \frac{\pi}{2}$ or $\theta = \frac{3\pi}{2}$ away from the origin. In agreement, we require the functions to be bounded $0 \leq \alpha_k^E \leq 1$, and we define them as:

$$\alpha_k^E = \begin{cases} c \left(d_1 - (-1)^k r \cos \left(\frac{\pi}{2} - \theta \right) g(R, l) \right), & \text{for } k = 1, 2, 3, 4 \\ c \left(\left[d_3^2 + (r g(R, l))^2 - (-1)^k r l \cos \left(\frac{\pi}{2} - \theta \right) g(R, l) \right]^{1/2} \right), & \text{for } k = 5, 6, 7, 8 \\ c (d_4 - r \cos(\pi - \theta) g(R, l)), & \text{for } k = 9, 10 \\ 0, & \text{for } k = 11, 12 \end{cases} \quad (3.2.3)$$

where the quantities used are similar to the ones used in the SiMPa functions. The constants $d_1 = \frac{l}{2}$, $d_4 = l$ and $d_3 = \frac{\sqrt{5}l}{2}$ are the fixed distances of the neighbouring pixels from the origin, with $l = 1$ units denoting the side length of the pixels.

As the molecule can lie anywhere along the edge of the pixel, the distance r is restricted between $0 \leq r < \frac{l}{2} - R$ where R is the radius of the molecule. The function $g(R, l) = \frac{1}{1 - 2R/l}$ is the mechanism accounting for identical performance of the diffraction regardless of the size of the molecule (details in Section 3.2.3), and is multiplied with the distance r . The parameter c controls the power of the diffraction with $0 < c \leq 1/2d_2$, when without loss of generality $R = 0$ (details can be found in Section 3.2.4).

In correspondence to the SiMPa functions, we adopt a cosine function to describe the periodic behaviour of the diffraction in terms of the directions. In order to have agreement, the co-function identity $\sin(\theta) = \cos\left(\frac{\pi}{2} - \theta\right)$ will be used as a tool. Without loss of generality, we assume the molecule is a point, i.e. $R = 0$, leading to $g(R, l) = 1$. Let the quantities \textcircled{A} , \textcircled{B} and \textcircled{C} , from α_k^{E} in Equation (3.2.3) with $g(R, l) = 1$,

$$\begin{aligned}\textcircled{A} &= d_1 - (-1)^k r \cos\left(\frac{\pi}{2} - \theta\right) \\ \textcircled{B} &= \left[d_3^2 + r^2 - (-1)^k r l \cos\left(\frac{\pi}{2} - \theta\right) \right]^{1/2} \\ \textcircled{C} &= d_4 - r \cos(\pi - \theta).\end{aligned}\tag{3.2.4}$$

On Figure 3.2.6, the molecule's minimum distance from the neighbouring pixels is presented. For pixels with $k = 2, 4$ and $k = 1, 3$ the distances are l_2 and l_1 respectively, easily calculated as $l_1 = l/2 + r \sin \theta$ and $l_2 = l/2 - r \sin \theta$ with the use of polar coordinates. Using the co-function identity we get the quantity \textcircled{A} . The appropriate sign is assigned by the term $(-1)^k$ for $k = 1, 2, 3, 4$. Regarding the neighbouring pixels with $k = 9, 10$, the molecule's minimum distance l_9 is constant and equal to $d_4 = l$ with respect to r and $\theta = \pi/2$ or $3\pi/2$. Therefore, for agreement across our definitions we use $r \cos(\pi - \theta)$ which is equal to 0 for $\theta = \pi/2$ or $3\pi/2$ as in \textcircled{C} . For the distances l_5 and l_8 for pixels with $k = 5, 6, 7, 8$ we have:

$$\begin{aligned}
l_5 &= \left(\frac{l}{2} + r \sin \theta\right)^2 + l^2 \Rightarrow l_7 = \left(\frac{5l^2}{4} + r^2 + rl\right)^{1/2} \\
l_8 &= \left(\frac{l}{2} - r \sin \theta\right)^2 + l^2 \Rightarrow l_7 = \left(\frac{5l^2}{4} + r^2 - rl\right)^{1/2}.
\end{aligned}$$

These quantities can be summarised as $l_k = \left[d_3^2 + r^2 - (-1)^k r l \cos\left(\frac{\pi}{2} - \theta\right)\right]^{1/2}$, $k = 5, 6, 7, 8$, as in \textcircled{B} with $d_3 = (\sqrt{5}l)/2$. The cosine function $\cos\left(\frac{\pi}{2} - \theta\right)$ can only take the values 1 or -1 when $\theta = \pi/2$ or $3\pi/2$ while the term $(-1)^k$ applies the appropriate sign for $k = 5, 6, 7, 8$.




Figure 3.2.6: Minimum distances on the special case where the molecule lies in on of the edges of the central pixel. The molecule is placed on an edge of the central pixel. The blue solid lines named by l_k , $k = 1, 2, 5, 8, 9$, denote the minimum distances of the molecule from the respective neighboring pixels.

We propose to use this special case of the SiMPa functions in the exact same way as described previously, according to the single event intensity I . The functions

α_k^E , $k = 1, \dots, 12$ will therefore describe the proportion of I each pixel in the neighbourhood should have. As $0 \leq \alpha_k^E(r, \theta, c) \leq 1$ we can calculate the quantities $(1 - \alpha_k^E(r, \theta, c))I$, $k = 1, \dots, 12$, for a specific combination of r and θ in order to obtain the behaviour of the entire neighbourhood. An evaluation for different combination of r and θ can be in Figure 3.2.7, where we use $I = 100$ and $c = 0.4$.




Figure 3.2.7: Implementation of special case when the molecule is placed on an edge of a pixel. The side length of the pixel is $l = 1$ units which mean that each pixel has dimension 1×1 . The radius R of the molecule is $R = 0.02$, the single event intensity is $I = 100$ and the power of spread parameter $c = 0.4$. The results from this special case are rounded to the closest integer. Four different positions of the molecule are used (a) $r = 0$ and $\theta = 0$. (b) $r = 0.15$ and $\theta = \pi/2$. (c) $r = 0.3$ and $\theta = \pi/2$. (d) $r = 0.47$ and $\theta = 3\pi/2$.

Molecule on the intersection of two edges of the central pixel

Here, we address the case where a molecule lies exactly on the intersection of two edges of the central pixel of the SiMPa functions. This special case is the simplest one as the molecule cannot be placed anywhere but the intersection of the edges. Therefore, a predetermined position of the molecule need to be obtained. We assume this position has $r = \theta = 0$ in order to have agreement with our definitions, hence the intersection works as the origin of the functions. The intensities diffraction is distributed along the sixteen neighbouring pixels and can only be altered by different choices of the power of diffraction parameter c . The proportion of the single event intensity I will be described by the intersection case functions $a_k^{\text{IS}}(c) \equiv a_k^{\text{IS}}$, $k = 1, \dots, 16$ and will be used in exactly the same way as the SiMPa functions (and the edge special case). We similarly require $0 \leq a_k^{\text{IS}} \leq 1$, achieved by bounding c between $0 < c \leq 1/2d_2$ (details in Section 3.2.4). The intersection special case can be seen in Figure 3.2.8 along with the proposed use of the functions using the single event intensity I . We define this special case with the functions,

$$a_k^{\text{IS}} = \begin{cases} cd_4, & \text{for } k = 1, \dots, 8 \\ cd_5, & \text{for } k = 9, \dots, 12 \\ 0, & \text{for } k = 13, \dots, 16, \end{cases} \quad (3.2.5)$$

where $d_4 = l$ and $d_5 = \sqrt{2}l$ with $l = 1$ units to be the side length of the pixel. As can be seen by the a_k^{IS} , $k = 1, \dots, 16$ functions there is no contribution of r and θ on this case as the molecule has a fixed position exactly on the intersection of the edges.




Figure 3.2.8: Special case when a molecule lies on the intersection of two edges of the central pixel. On the left panel (a) of the figure the origin of this case can be seen along with the fixed distances d_4 and d_5 from the neighboring pixels. The appropriate indexing of the a_k^{IS} , $k = 1, \dots, 16$ is presented. On the right panel (b) of the figure the use of these functions in terms of the single event intensity I is implemented.

3.2.2 Extended Single Molecule Pattern (SiMPa_X) functions

In this section the extended single molecule pattern (SiMPa_X) functions are presented, which are an extension of the single molecule pattern (SiMPa) functions in terms of the number of neighbouring pixels. We assume the intensities diffraction spreads along a larger neighbourhood than the 3×3 pixel region we introduced in the SiMPa functions. The molecule is still allowed to exist anywhere within the central pixel and its position is measured based on a distance r and direction θ from the center of the central pixel (same origin as in SiMPa functions). The SiMPa_X functions allow the neighbourhood to be a 5×5 pixel region, which can be graphically seen in Figure 3.2.9. The exact same usage concept is considered, by introducing the single event intensity I , with the SiMPa_X functions denoting the proportion of I each pixel

in the neighbourhood is assigned.




Figure 3.2.9: Extended Single Molecule Pattern (SiMPa_X) functions. The molecule can lie anywhere within the pixel and the intensities diffraction spread along the 5×5 pixel neighbourhood. The distribution of the SiMPa_X functions a_k , $k = 0, \dots, 24$ in the neighbourhood, is shown, with the origin being the center of the pixel, The quantities d_1 , d_2 , d_6 and d_7 are the fixed distances from the origin.

The extended single molecule pattern (SiMPa_X) functions are denoted by $a_k(r, \theta, c) \equiv a_k$, $k = 0, \dots, 24$ with the indexing chosen according to the fixed distances of each neighbouring pixel from the origin. The appropriate distribution of the functions can be seen in Figure 3.2.9 along with the fixed distances d_1 , d_2 , d_6 , d_7 and d_8 from the origin. We similarly require $0 \leq a_k \leq 1$ and we define the SiMPa_X as,

$$a_k = \begin{cases} c \left(d_1 - r \cos \left(\frac{(k-1)\pi}{4} + \theta \right) g(R, l) \right), & k = 1, 3, 5, 7 \\ c \left(\left((rg(R, l))^2 + d_2^2 + 2d_2 s_k r \cos \left(\frac{(k+1)\pi}{4} - \theta \right) g(R, l) \right)^{1/2} \right), & k = 2, 4, 6, 8 \\ c \left(d_6 - r \cos \left(\frac{(k-1)\pi}{4} + \theta \right) g(R, l) \right), & k = 9, 11, 13, 15 \\ c \left(\left[(rg(R, l))^2 + d_8^2 + 3\sqrt{2} l s_k r \cos \left(\frac{(k+1)\pi}{4} - \theta \right) g(R, l) \right]^{1/2} \right), & k = 10, 12, 14, 16 \\ c \left(\left[(rg(R, l))^2 + d_7^2 + (-1)^k l r \cos (3 \cos \theta - s_k^* \sin \theta) g(R, l) \right]^{1/2} \right), & k = 17, 18, 19, 20 \\ c \left(\left[(rg(R, l))^2 + d_7^2 + (-1)^k l r \cos (\cos \theta - s_k^* 3 \sin \theta) g(R, l) \right]^{1/2} \right), & k = 21, 22, 23, 24 \\ 0, & k = 0 \end{cases}, \quad (3.2.6)$$

where all the parameters are similar to the ones in SiMPa functions from Equation 3.2.1. Recalling from Section 3.2, the constant $l = 1$ units denotes the side length of the pixel and R is the radius of the molecule with $R < l/2$. We require $r \geq 0$ with $r|\cos(\theta)| < \frac{l}{2} - R$ and $r|\sin(\theta)| < \frac{l}{2} - R$ in order for the molecule to lie within the central pixel. The constants d_1 , d_2 , d_6 , d_7 and d_8 are the fixed distances of the neighbouring pixels from the origin with $d_1 = \frac{l}{2}$, $d_2 = \frac{l}{\sqrt{2}}$, $d_6 = \frac{3l}{2}$, $d_7 = \frac{\sqrt{10}l}{2}$ and $d_8 = \frac{\sqrt{18}l}{2}$. The function $g(R, l)$ is the mechanism to maintain similar performance regardless of the size of the molecule with respect to the pixels' side length, $g(R, l) = \frac{1}{1 - 2R/l}$ (Section 3.2.3). The parameter c controls the power of diffraction across the entire neighbourhood and should obey the restriction $0 < c \leq 1/(d_8 + d_2)$ when $R = 0$ without loss of generality (Section 3.2.4).

Accordingly to the SiMPa functions, the minimum distance of the position of the molecule within the central pixel and its neighbouring pixels is obtained. For the matter of calculations, we assume that $R = 0$ hence $g(R, l) = 1$, without loss of generality. Let the quantities \textcircled{A}^* , \textcircled{B}^* , \textcircled{C}^* and \textcircled{D}^* from the SiMPa_X functions in Equation (3.2.6) to be:

$$\begin{aligned}
\textcircled{A}^* &= d_6 - r \cos \left(\frac{(k-1)\pi}{4} + \theta \right) \\
\textcircled{B}^* &= \left[r^2 + d_8^2 + 3\sqrt{2}l_s r \cos \left(\frac{(k+1)\pi}{4} - \theta \right) \right]^{1/2} \\
\textcircled{C}^* &= \left[r^2 + d_7^2 + (-1)^k l r \cos (3 \cos \theta - s_k^* \sin \theta) \right]^{1/2} \\
\textcircled{D}^* &= \left[r^2 + d_7^2 + (-1)^k l r \cos (\cos \theta - s_k^* 3 \sin \theta) \right]^{1/2}. \quad (3.2.7)
\end{aligned}$$

All these quantities are obtained in a similar manner to the SiMPa functions. We assume a molecule lies in the central pixel on a distance r units away from the origin with direction θ . We denote the minimum distance of the molecule from the neighbouring pixels with l_k where k is the respective index of the pixel. The different distances are presented in Figure 3.2.10.




Figure 3.2.10: Minimum distances of the neighboring pixels of a molecule placed r units on direction θ away from the origin. The distances are denoted with the orange and blue solid lines from the respective pixels. (a) Minimum distances l_k of the molecule from the corner pixels (blue solid lines) and from pixels with $k = 9, 11, 13, 15$ (orange solid lines). (b) Minimum distances l_k of the molecule from neighboring pixels with $k = 17, 18, \dots, 24$ (blue solid lines).

Regarding the neighbouring pixels with indexes $k = 9, 11, 13, 15$, the distances l_k , shown with the orange solid lines in part (a) of the Figure 3.2.10, are obtained using polar coordinates,

$$\begin{aligned} l_9 &= \frac{3l}{2} - r \cos(\theta) \\ l_{11} &= \frac{3l}{2} + r \cos(\theta) = \frac{3l}{2} - r \sin(-\theta) = \frac{3l}{2} - r \cos\left(\frac{\pi}{2} + \theta\right) \\ l_{13} &= \frac{3l}{2} + r \cos(\theta) = \frac{3l}{2} - r \cos(\theta + \pi) \\ l_{15} &= \frac{3l}{2} - r \sin(\theta) = \frac{3l}{2} - r \cos\left(\frac{\pi}{2} - \theta\right) = \frac{3l}{2} - r \cos\left(\frac{3\pi}{2} + \theta\right), \end{aligned}$$

which can be summed up in $l_k = d_6 - r \cos\left(\frac{(k-1)\pi}{4} + \theta\right) \equiv \textcircled{A}^*$ in Equation (3.2.7)

when $k = 9, 11, 13, 15$, with $d_6 = \frac{3l}{2}$.

In order to obtain the minimum distances l_k when $k = 10, 12, 14, 16$, presented with the blue solid lines in part (a) of Figure 3.2.10, we apply Pythagoras theorem along with the sum-to-product identity and some basic co-functions. Analytically what we get for l_{10} ,

$$\begin{aligned} l_{10}^2 &= \left(\frac{3l}{2} - r \cos \theta\right)^2 + \left(\frac{3l}{2} + r \sin \theta\right)^2 \\ &= \frac{18l^2}{4} + r^2 - 3rl(\cos \theta - \sin \theta) \\ &= \frac{18l^2}{4} + r^2 - 3rl\left(\cos \theta - \cos\left(\frac{\pi}{2} - \theta\right)\right) \\ &= \frac{18l^2}{4} + r^2 - 3rl\left(-2 \sin\left(\frac{\pi}{4}\right) \sin\left(\theta - \frac{\pi}{4}\right)\right) \\ &= \frac{18l^2}{4} + r^2 + 3\sqrt{2}rl \cos\left(\frac{3\pi}{4} - \theta\right). \end{aligned}$$

Therefore, $l_{10} = \left(r^2 + d_8^2 + 3\sqrt{2}rl \cos \left(\frac{3\pi}{4} - \theta \right) \right)^{1/2}$, with $d_8 = \sqrt{18}l/2$. Similarly,

$$\begin{aligned} l_{12}^2 &= \left(\frac{3l}{2} + r \sin \theta \right)^2 + \left(\frac{3l}{2} + r \cos \theta \right)^2 \Rightarrow l_{12} = \left(r^2 + d_8^2 + 3\sqrt{2}rl \cos \left(\frac{\pi}{4} - \theta \right) \right)^{1/2} \\ l_{14}^2 &= \left(\frac{3l}{2} + r \cos \theta \right)^2 + \left(\frac{3l}{2} - r \sin \theta \right)^2 \Rightarrow l_{14} = \left(r^2 + d_8^2 - 3\sqrt{2}rl \cos \left(\frac{3\pi}{4} - \theta \right) \right)^{1/2} \\ l_{16}^2 &= \left(\frac{3l}{2} - r \cos \theta \right)^2 + \left(\frac{3l}{2} - r \sin \theta \right)^2 \Rightarrow l_{16} = \left(r^2 + d_8^2 - 3\sqrt{2}rl \cos \left(\frac{\pi}{4} - \theta \right) \right)^{1/2}. \end{aligned}$$

Using a common index k with $k = 10, 12, 14, 16$ we summarise these functions into $l_k = \left(r^2 + d_8^2 + 3\sqrt{2}l s_k r \cos \left(\frac{(k+1)\pi}{4} - \theta \right) \right)^{1/2} \equiv \mathbb{B}^*$ in Equation (3.2.7) with $k = 10, 12, 14, 16$ and $d_8 = \sqrt{18}l/2$. The quantity s_k is the same declaring the appropriate sign of the cosine function with $s_k = +1, -1, +1, -1$ for $k = 10, 12, 14, 16$ respectively.

Finally, the minimum distances from pixels with $k = 17, \dots, 24$, shown with the blue solid lines in part (b) of Figure 3.2.10, are calculated as:

$$\begin{aligned} l_{17}^2 &= \left(\frac{3l}{2} - r \cos \theta \right)^2 + \left(\frac{l}{2} + r \sin \theta \right)^2 = \frac{10l^2}{4} + r^2 - rl(3 \cos \theta - \sin \theta) \\ l_{18}^2 &= \left(\frac{3l}{2} + r \cos \theta \right)^2 + \left(\frac{l}{2} + r \sin \theta \right)^2 = \frac{10l^2}{4} + r^2 + rl(3 \cos \theta + \sin \theta) \\ l_{19}^2 &= \left(\frac{3l}{2} - r \cos \theta \right)^2 + \left(\frac{l}{2} - r \sin \theta \right)^2 = \frac{10l^2}{4} + r^2 - rl(3 \cos \theta + \sin \theta) \\ l_{20}^2 &= \left(\frac{3l}{2} + r \cos \theta \right)^2 + \left(\frac{l}{2} - r \sin \theta \right)^2 = \frac{10l^2}{4} + r^2 + rl(3 \cos \theta - \sin \theta) \\ l_{21}^2 &= \left(\frac{l}{2} - r \cos \theta \right)^2 + \left(\frac{3l}{2} + r \sin \theta \right)^2 = \frac{10l^2}{4} + r^2 - rl(\cos \theta - 3 \sin \theta) \end{aligned}$$

$$\begin{aligned}
l_{22}^2 &= \left(\frac{l}{2} + r \cos \theta\right)^2 + \left(\frac{3l}{2} + r \sin \theta\right)^2 = \frac{10l^2}{4} + r^2 + rl(\cos \theta + 3 \sin \theta) \\
l_{23}^2 &= \left(\frac{l}{2} - r \cos \theta\right)^2 + \left(\frac{3l}{2} - r \sin \theta\right)^2 = \frac{10l^2}{4} + r^2 - rl(\cos \theta + 3 \sin \theta) \\
l_{24}^2 &= \left(\frac{l}{2} + r \cos \theta\right)^2 + \left(\frac{3l}{2} - r \sin \theta\right)^2 = \frac{10l^2}{4} + r^2 + rl(\cos \theta - 3 \sin \theta).
\end{aligned}$$

Grouping the distances l_k of pixels with $k = 17, 18, 19, 20$ and the ones with $k = 21, 22, 23, 24$ we end up with $l_k = [r^2 + d_7^2 + (-1)^k lr \cos(3 \cos \theta - s_k^* \sin \theta)]^{1/2} \equiv \mathbb{C}^*$ when $k = 17, 18, 19, 20$ and $l_k = [r^2 + d_7^2 + (-1)^k lr \cos(\cos \theta - s_k^* 3 \sin \theta)]^{1/2} \equiv \mathbb{D}^*$ when $k = 21, 22, 23, 24$ in Equation (3.2.7), with $s_k^* = +1, -1, -1, +1$ and $d_7 = \frac{\sqrt{10}l}{2}$.

3.2.3 Function of relative side length of pixel l and molecule's radius R , $g(R, l)$

In this section, we focus on the function of the relative side length of the pixel l and the radius of the molecule R which we defined as $g(R, l) = \frac{1}{1 - 2R/l}$. This function works as a correction mechanism to maintain identical performance of the SiMPa functions as the size of the molecule is different. We make the assumption that a molecule's projection on the 2-dimensional space is a circular object, therefore we measure its size by its radius. It should be noted that this is not a necessity and can be relaxed, by using corresponding appropriate function g . In order to achieve identical behaviour of the SiMPa functions, we want robust performance no matter what the size of the molecule is, a case achieved by taking into consideration the side length of the pixel.




Figure 3.2.11: A molecule with different radius placed on the same position according to the SiMPa functions. On the left panel the molecule is placed on a distance $r = r_1$ units away from the origin with an angle $\theta = \pi/4$, which is the maximum it can lie on in that direction having this size. On the right panel a larger molecule is placed on a distance $r = r_2$ units away from the origin on the same angle, which is again the larger distance it can lie accordingly to this angle and its size.

In Figure 3.2.11, two molecules with different radius can be seen placed away from origin on distances r_1 and r_2 with same direction $\theta = \pi/4$. These distances, with $r_1 > r_2$, represent the maximum distances the molecules can lie onto, under direction $\theta = \pi/4$ and their corresponding sizes. Therefore, we can state that there is a dependency on the molecule size and the maximum distance it can be placed in order to remain within the central pixel. This can also be seen by the restrictions on different choices of r and θ , where we require $r|\cos(\theta)| \leq \frac{l}{2} - R$ and $r|\sin(\theta)| \leq \frac{l}{2} - R$.

We introduce the correction function $g(R, l)$, which is multiplied with the distance r in the SiMPa functions, achieving equivalent performance regardless of the molecule's size. We define this function to be the ratio of the maximum distance a molecule can lie onto within a pixel in a arbitrary chosen direction, when its radius is $R = 0$ (single point), over the maximum distance when its radius is R (a circle with radius R), that is $g(R, l) = \frac{\max_{\theta}\{r\}_{R=0}}{\max_{\theta}\{r\}_{R=R}}$. The side length of the pixel cannot be immediately identified but $g(R, l)$ can be expressed as a function of l and R . As r and

θ are continuous parameters this quantity can be evaluated in a number of different combinations. However this ratio will be the same no matter what the choice of r and θ is due to the restrictions described above. Therefore,

$$\begin{aligned}
 g(R, l) &= \frac{\max_{\theta}\{r\}_{R=0}}{\max_{\theta}\{r\}_{R=R}} = \frac{(l/2) / \max\{|\cos(\theta)|, |\sin(\theta)|\}}{(l/2 - R) / \max\{|\cos(\theta)|, |\sin(\theta)|\}} = \frac{l/2}{l/2 - R} \\
 &\Rightarrow g(R, l) = \frac{1}{1 - 2R/l}, \quad R < l/2.
 \end{aligned} \tag{3.2.8}$$

This function $g(R, l)$ in Equation (3.2.8) does not depend either on the angle θ or the distance r but only on the molecule's size and the side length of the pixel.

3.2.4 Power of diffraction parameter c

In this section, we focus on the diffraction parameter c and form an analysis to study the limitations existing on its choice. As discussed on Section 3.2, the role of the SiMPa functions is to specify the proportion of a single event intensity I each pixel in the predetermined neighborhood is assigned. Therefore, we require these functions to be bounded, $0 \leq a_k(r, \theta, c) \leq 1$, so when the quantities $(1 - a_k(r, \theta, c))I$ are obtained, they contain the respective intensities for each pixel with $k = 0, \dots, 8$.

The SiMPa functions, as in Equation (3.2.1), depend on the relative position of a molecule with respect to the predetermined origin, described by the distance r and direction θ , which is multiplied by the diffraction parameter c . The role of this parameter is two-fold. It transforms the position of the molecule into a proportion and adjusts the power of the diffraction across the neighbourhood. Higher values of c boost the SiMPa functions, hence indicate a faster decay of the intensities diffraction within the neighbourhood, in contrast to lower values corresponding to a slower decay. In order for our requirement $0 \leq a_k(r, \theta, c) \leq 1$ to hold, c should be bounded,

since we allow a molecule to lie anywhere within the central pixel. The boundary will be acquired using the higher distance a molecule can have from every neighbouring pixel, when it is placed on the furthest distance from the origin under the respective angle. We will work with the SiMPa_X functions as the SiMPa functions are a subset of them, therefore by recalling from Equation (3.2.6),

$$\alpha_k = \begin{cases} c \overbrace{\left(d_1 - r \cos \left(\frac{(k-1)\pi}{4} + \theta \right) g(R, l) \right)}^{A_1}, & k = 1, 3, 5, 7 \\ c \overbrace{\left(\left((rg(R, l))^2 + d_2^2 + 2d_2s_k r \cos \left(\frac{(k+1)\pi}{4} - \theta \right) g(R, l) \right)^{1/2} \right)}^{A_2}, & k = 2, 4, 6, 8 \\ c \overbrace{\left(d_6 - r \cos \left(\frac{(k-1)\pi}{4} + \theta \right) g(R, l) \right)}^{A_3}, & k = 9, 11, 13, 15 \\ c \overbrace{\left(\left[(rg(R, l))^2 + d_8^2 + 3\sqrt{2}ls_k r \cos \left(\frac{(k+1)\pi}{4} - \theta \right) g(R, l) \right]^{1/2} \right)}^{A_4}, & k = 10, 12, 14, 16 \\ c \overbrace{\left(\left[(rg(R, l))^2 + d_7^2 + (-1)^k l r \cos(3 \cos \theta - s_k^* \sin \theta) g(R, l) \right]^{1/2} \right)}^{A_5}, & k = 17, 18, 19, 20 \\ c \overbrace{\left(\left[(rg(R, l))^2 + d_7^2 + (-1)^k l r \cos(\cos \theta - s_k^* 3 \sin \theta) g(R, l) \right]^{1/2} \right)}^{A_6}, & k = 21, 22, 23, 24 \\ 0, & k = 0 \end{cases},$$

where $d_1 = \frac{l}{2}$, $d_2 = \frac{l}{\sqrt{2}}$, $d_6 = \frac{3l}{2}$, $d_7 = \frac{\sqrt{10}l}{2}$ and $d_8 = \frac{\sqrt{18}l}{2}$. From the inequality $0 \leq \alpha_k(r, \theta, c) \leq 1$ we get that:

$$\begin{aligned} 0 \leq \alpha_k(r, \theta, c) \leq 1 &\Rightarrow 0 \leq c \leq \min\{1/A_1, 1/A_2, 1/A_3, 1/A_4, 1/A_5, 1/A_6\} \\ &\Rightarrow 0 \leq c \leq 1/\max\{A_1, A_2, A_3, A_4, A_5, A_6\}. \end{aligned} \quad (3.2.9)$$

Therefore, we are seeking for the combination of r and θ that maximises the quantities

A_1, A_2, \dots, A_6 . Then the upper bound for the parameter c is the value that minimises the quantities $1/A_1, 1/A_2, \dots, A/A_6$. Without loss of generality, we will search for this combination of r and θ on the first quadrant of our coordinate system, while additionally assume that the molecule is a single point, hence its radius is $R = 0$, leading to $g(R, l) = 1$. As the SiMPa functions take into consideration the distances of the molecule from all the neighbouring pixels in the predetermined region, the maximum value of A_1, A_2, \dots, A_6 will be obtained when the molecule is placed the furthest away from the origin. This corresponds to the highest value r can take under the appropriate direction θ . According to the restrictions, every combination of r and θ should satisfy:

$$\begin{aligned} r|\cos\theta| < l/2 - R \quad \& \quad r|\sin\theta| < l/2 - R \quad \xrightarrow{R=0} \quad r < \frac{l/2}{|\cos\theta|} \quad \& \quad r < \frac{l/2}{|\sin\theta|} \\ \Rightarrow \quad r < \min \left\{ \frac{l/2}{|\cos\theta|}, \frac{l/2}{|\sin\theta|} \right\}, \end{aligned}$$

which means that the maximum value the distance r can take will be obtained when this boundary is maximised for all possible values of θ , i.e

$$r_{\max} < \max_{\forall\theta} \left\{ \min \left\{ \frac{l/2}{|\cos\theta|}, \frac{l/2}{|\sin\theta|} \right\} \right\}.$$

By construction of the SiMPa functions, the origin is defined to be the center of the pixel and the molecule is allowed to lie anywhere within that pixel (central pixel). This implies the maximum distance r_{\max} is obtained when the molecule is placed all the way up the diagonal line from the origin, as shown in Figure 3.2.12. In terms of our coordinate system, this is achieved when $\theta = \pi/4$, therefore

$$r < \max_{\theta=\pi/4} \left\{ \min \left\{ \frac{l/2}{|\cos\theta|}, \frac{l/2}{|\sin\theta|} \right\} \right\} = l/\sqrt{2} \Rightarrow r_{UL} = l/\sqrt{2}.$$

α_{14}	α_{24}	α_{15}	α_{23}	α_{16}
α_{20}	α_6	α_7	α_8	α_{19}
α_{13}	α_5	$(0,0)$	α_1	α_9
α_{18}	α_4	α_3	α_2	α_{17}
α_{12}	α_{22}	α_{11}	α_{21}	α_{10}

Figure 3.2.12: A molecule is placed at the maximum distance on the diagonal line from the origin of the SiMPa_X functions, based on the first quadrant of the coordinate system. This is achieved when $\theta = \pi/4$.

We evaluate the quantities A_1, A_2, A_3, A_4, A_5 and A_6 , by applying r_{UL} under the direction of $\theta = \pi/4$ for every k on each quantity respectively. According to Equation (3.2.9), we are interested in the maximum of these quantities therefore,

$$\begin{aligned}
 \max_{\theta=\pi/4, k=1,3,5,7} (A_1) &= d_1 + r_{UL} \max_{\theta=\pi/4} \left\{ -\cos \left(\frac{(k-1)\pi}{4} + \theta \right) \right\} \\
 &= l/2 + (l/\sqrt{2})0.7071 \stackrel{l=1}{=} 1, \\
 \max_{\theta=\pi/4, k=2,4,6,8} (A_2) &= \left(r_{UL}^2 + d_2^2 + 2d_2 r_{UL} \max_{\theta=\pi/4} \left\{ s_k \cos \left(\frac{(k+1)\pi}{4} - \theta \right) \right\} \right)^{(1/2)} \\
 &= \left((l/\sqrt{2})^2 + (l/\sqrt{2})^2 + 2(l/\sqrt{2})^2 \right)^{(1/2)} \\
 &= 2(l/\sqrt{2}) \stackrel{l=1}{=} 1.414214, \\
 \max_{\theta=\pi/4, k=9,11,13,15} (A_3) &= d_6 + r_{UL} \max_{\theta=\pi/4} \left\{ -\cos \left(\frac{(k-1)\pi}{4} + \theta \right) \right\} \\
 &= 3l/2 + (l/\sqrt{2})0.7071 \stackrel{l=1}{=} 2,
 \end{aligned}$$

$$\begin{aligned}
\max_{\theta=\pi/4, k=10,12,14,16}(A_4) &= \left(r_{UL}^2 + d_8^2 + 3\sqrt{2}lr_{UL} \max_{\theta=\pi/4} \left\{ s_k \cos \left(\frac{(k+1)\pi}{4} - \theta \right) \right\} \right)^{(1/2)} \\
&= \left((l/\sqrt{2})^2 + (\sqrt{18}l/2)^2 + 3\sqrt{2}l(l/\sqrt{2}) \right)^{(1/2)} \stackrel{l=1}{=} 2.828427, \\
\max_{\theta=\pi/4, k=17,18,19,20}(A_5) &= \left(r_{UL}^2 + d_7^2 + lr_{UL} \max_{\theta=\pi/4} \{ (-1)^k (3 \cos \theta - s_k^* \sin \theta) \} \right)^{(1/2)} \\
&= \left((l/\sqrt{2})^2 + (\sqrt{10}l/2)^2 + l(l/\sqrt{2}) * 2.828427 \right)^{(1/2)} \stackrel{l=1}{=} 2.236068, \\
\max_{\theta=\pi/4, k=21,22,23,24}(A_6) &= \max_{\theta=\pi/4, k=17,18,19,20}(A_5) \stackrel{l=1}{=} 2.236068.
\end{aligned}$$

Regarding the SiMPa functions in Equation (3.2.1), only the quantities A_1 and A_2 are considered, that is when $k = 1, \dots, 8$. According to Equation 3.2.9,

$$\begin{aligned}
0 \leq \alpha_k(r, \theta, c) \leq 1 &\Rightarrow 0 \leq c \leq 1/\max\{A_1, A_2\} \\
&\Rightarrow 0 \leq c \leq 1/\max\left\{ \max_{\theta=\pi/4, k=1,3,5,7}(A_1), \max_{\theta=\pi/4, k=2,4,6,7}(A_2) \right\} \\
&\Rightarrow 0 \leq c \leq 1/\max_{\theta=\pi/4, k=2,4,6,7}(A_2) \\
&\Rightarrow 0 \leq c \leq 0.7071.
\end{aligned}$$

or this can be written as $0 \leq c \leq 1/2d_2$. We note here that for consistency, the same boundaries for c are used for the special cases in Section 3.2.1, even if the edge case can allow a higher upper bound.

For the extended case of the single molecule pattern, X-SiMPa, functions in Equation (3.2.6), every quantity A_1, \dots, A_6 is included leading to

$$\begin{aligned}
0 \leq \alpha_k(r, \theta, c) \leq 1 &\Rightarrow 0 \leq c \leq 1/\max\{A_1, A_2, A_3, A_4, A_5, A_6\} \\
&\Rightarrow 0 \leq c \leq 1/\max_{\theta=\pi/4, k=10,12,14,16}(A_4) \\
&\Rightarrow 0 \leq c \leq 0.3535568,
\end{aligned}$$

or equivalently $0 \leq c \leq 1/(d_8 + d_2)$.

3.2.5 Single frame generation using the SiMPa functions

In this section we present a pseudo-algorithm for producing synthetic data in terms of individual frames, using the SiMPa functions. No dependency is present between the frames and the user is able to choose every parameter introduced in the SiMPa functions, while also add error of the preferred level. This error can be considered to be the background noise every imaging technique is subject to and is considered to come from a Gaussian distribution, a choice related to our modelling in the following Section 3.3.2. Generation of stacks of frames with consecutive frames being dependent is presented in Chapter 4, after we introduce the state-space model based on SiMPa functions. The pseudo-algorithm is presented below, generating individual synthetic frames according to the SiMPa functions.

Pseudo-algorithm - Generation of frames using the SiMPa functions

- (i) Choose values for the quantities F, n, m, K, R, l, d as well values for the parameters c and I where,
 - F : number of individual frames to be generated
 - n and m : number of rows and columns for each frame respectively
 - K : total number of molecules
 - R : radius of the molecule
 - l : side length of pixel
 - d : proportion quantity, $0 < d < 1$
 - c : power of the intensities diffraction
 - I : intensity of a single event
- (ii) Generate distances r and angles θ that satisfy the restrictions $r|\cos(\theta)| < \frac{l}{2} - R$ and $r|\sin(\theta)| < \frac{l}{2} - R$ and then randomly assign the combinations to each one of the molecules in K .

- (iii) Choose randomly the number of active molecules on each frame, denoted by N_F and randomly draw K_F combinations of r and θ from K for each one of the molecules.
- (iv) For every frame f from $1 : F$ repeat
 1. Create a lattice of size $n \times m$ with a baseline intensity $d \times I$.
 2. Generate $N_F[f]$ random numbers u between 1 and $n \times m$ for each of the molecules chosen to be active on frame f .
 3. Transform the random numbers u into positions inside the lattice by applying $Rows = \text{Remainder}(u - 1, n) + 1$ and $Columns = \text{Quotient}(u - 1, n) + 1$.
 4. For every molecule k_f in $1 : N_F[f]$ repeat
 - (a) Evaluate the SiMPa functions using the Equation (3.2.1) with r and θ from $K_F[k_f]$ and obtain the quantities $(1 - \alpha_k)I$ for $k = 0, \dots, 8$.
 - (b) Place the quantities $(1 - \alpha_k)I$ on the lattice, where the quantity $(1 - \alpha_0)I$ is located at the position $(Rows[k_f], Column[k_f])$.
 5. Add independent and identically distributed background error to every pixel on the lattice, for instance by using a Gaussian distribution with mean μ_b and precision τ_b .

3.3 Individual frame inference based on SiMPa

In this section we introduce the complete model for drawing inference on individual frames based on the SiMPa functions. Firstly, a brief introduction to Markov random fields (MRF) is presented (Section 3.3.1), a key tool in our modelling in order to identify pixel regions where the potentially active molecules lie onto. These regions, commonly known as regions of interest (ROI), are basically a set of pixels of higher total intensity. Based on the regions of interest capturing the important parts within a frame, we next obtain the general form of our Bayesian modelling based on the SiMPa functions (Section 3.3.2). This mainly consists of a moving region (MR) which scans neighbourhoods within the frame, aiming to capture intensities diffraction, or Point Spread Functions (PSF), of nearby active molecules. A moving region is considered to consist of a number of PSF's, each one described by the SiMPa functions, hence it represents our proposed data generation mechanism, or likelihood. Every PSFs is fully characterised by the SiMPa functions, describing the proportions of a single event intensity I each pixel should have according to the molecule's position (Section 3.2). A Bayesian framework facilitates coherent propagation of uncertainty for each one of the parameters in the model.

Next, we present the probabilistic scheme of drawing inference on pixels including active molecules, forming what we call the counting scheme, and localisation of them within these pixels, both based on our novel pattern-configuration-realisation concept presented in Section 3.3.3. Briefly, a pattern denotes a structure formed within a MR consisting of a combination of active and inactive pixels, while a configuration accounts for the contribution of potential PSFs based on the SiMPa functions to the creation of the pattern. A realisation is a position of a molecule according to a discretisation of r and θ . The pattern-configuration-realisation concept within a Bayesian framework forms our individual frame probabilistic inference scheme, which is provided with formal definitions. Two similar pseudo-algorithms are provided in Section 3.3.5, one summarising the individual frame counting scheme, where the

location of the molecules is fixed at the origin of the SiMPa functions, i.e. $r = \theta = 0$, with the second additionally allowing for precise localisation with respect to a discretisation of r and θ .

3.3.1 Markov random fields (MRF)

In this section we present a brief introduction to Markov random fields (MRF), a fundamental element of our inference scheme. Let $\{X_t\}, t = 1, 2, \dots$ to be a sequence of random variables taking values on a finite set. The univariate process X_t is a discrete time Markov chain taking values on this finite set, if it satisfies the Markov property, that is $P(X_{t+1} = x_{t+1} | X_1 = x_1, \dots, X_t = x_t) = P(X_{t+1} = x_{t+1} | X_t = x_t)$ for all t . This property implies that the full conditional distribution of X_t moving to the next state depends only on the present state of X_t .




Figure 3.3.1: Neighbourhood systems of 1st and 2nd order. On the left panel of the figure the 1st order neighbourhood structure is presented where on the right panel the 2nd order neighbourhood. The black filled in circle denotes the central pixel of each neighbourhood where the empty circles the neighbouring pixel respectively.

Since our focus is on 2-dimensional frames, each one described as a lattice of pixels, we consider an MRF (Definition 3.3.1), the extension of the Markov chain in 2D (Geman and Geman [1987]). We denote with S a lattice consisting of square pixels with dimension $N = n \times m$, where n and m represent the rows and columns respectively. Each coordinate pair (i, j) corresponds to the location of a pixel within the lattice,

hence $s_{ij} \in S$ for all $i = 1, \dots, n$ and $j = 1, \dots, m$. The neighbourhood of the pixel s_{ij} is defined as the set $\partial_{s_{ij}} = \{q \in S : q \sim s_{ij}\}$, where \sim denotes that q belongs to the neighbourhood of s_{ij} . Suitable neighbourhood systems can be considered for different behaviours. Two of the most common are the first (4-neighbour system) and second-order (8-neighbour system) neighbourhood structures presented in Figure 3.3.1.

In terms of functions of pixels in a lattice, these neighbourhood structures can be defined as,

$$\partial_{s_{1\text{st-order}}} \equiv \partial_{s_{ij}} = \begin{cases} s_{i-1,j}, & \text{if } 2 \leq i \leq n, 1 \leq j \leq m \\ s_{i,j+1}, & \text{if } 1 \leq i \leq n, 1 \leq j \leq m-1 \\ s_{i+1,j}, & \text{if } 1 \leq i \leq n-1, 1 \leq j \leq m \\ s_{i,j-1}, & \text{if } 1 \leq i \leq n, 2 \leq j \leq m, \end{cases} \quad (3.3.1)$$

$$\partial_{s_{2\text{nd-order}}} \equiv \partial_{s_{ij}} = \begin{cases} s_{i-1,j+1}, & \text{if } 2 \leq i \leq n, 1 \leq j \leq m-1 \\ s_{i-1,j}, & \text{if } 2 \leq i \leq n, 1 \leq j \leq m \\ s_{i-1,j-1}, & \text{if } 2 \leq i \leq n, 2 \leq j \leq m \\ s_{i,j+1}, & \text{if } 1 \leq i \leq n, 1 \leq j \leq m-1 \\ s_{i,j}, & \text{if } 1 \leq i \leq n, 1 \leq j \leq m \\ s_{i,j-1}, & \text{if } 1 \leq i \leq n, 2 \leq j \leq m \\ s_{i+1,j+1}, & \text{if } 1 \leq i \leq n-1, 1 \leq j \leq m-1 \\ s_{i+1,j}, & \text{if } 1 \leq i \leq n-1, 1 \leq j \leq m \\ s_{i+1,j-1}, & \text{if } 1 \leq i \leq n-1, 2 \leq j \leq m-1, \end{cases} \quad (3.3.2)$$

where the conditions on each neighbouring pixel imply the possible choices of pixels which result in a complete corresponding neighbourhood structure.

Definition 3.3.1. [Li, 2009] The process $\{X_s\} = \{X_s : s \in S\}$ with $S = \{1, \dots, N (= n \times m)\}$ is considered to be a Markov random field with respect to a neighbourhood system ∂s if it satisfies the following conditions,

- (i) positivity condition, $P(X = x) = P(X_1 = x_1, \dots, X_s = X_s) > 0$ for every realization x
- (ii) Markov property via local conditionals, $P(X_s = x_s | X_q = x_q, \forall q \in S) = P(X_s = x_s | X_{\partial s} = x_{\partial s})$ for all pixels $s \in S$.

3.3.2 Bayesian modelling of individual frames using the SiMPa functions

In this section we introduce our Bayesian modelling for individual frames based on the SiMPa functions. As previously discussed on Chapter 2, data from SMLM are obtained as a stack of sparse frames, containing a number of active molecules, processed in order to reconstruct a final super-resolution image from the determined positions of the molecules. The procedure can be considered to consist of two parts, the spatial and temporal components. The former accounts for acquiring all available information within each frame, capturing behaviours and structures on the spatial domain, which can allow separate identification and localisation of active molecules. The latter is responsible for linking the active molecules over time, where time corresponds to the number of frames in the stack, taking into account the dynamics of the blinking process. In order to carry out appropriate inference on the stack of frames, one can claim that both components should be taken into consideration. In Chapter 4, we consider a state-space model based on SiMPa functions, temporally linking the spatial structure of individual frames. Here, we present the general form of our Bayesian model based on the SiMPa functions, capable of conducting individual frame inference on the spatial domain or stacks of frames without formal time dependence.

Starting with appropriate notation, we define \mathbf{y} to be an observed lattice of intensities

for an individual frame, that is $\mathbf{y} = \{y_{ij}\}$, where $i = 1, \dots, n$, $j = 1, \dots, m$ with n and m being the number of rows and columns of the lattice respectively. We introduce a latent mechanism which we assume is responsible for the generation of the intensities within the lattice and we denote by \mathbf{x} , with $\mathbf{x} = \{x_{ij}\}$, $i = 1, \dots, n$, $j = 1, \dots, m$. This unobserved lattice is considered to be a Markov Random Field (MRF), with its main role being the identification of the regions that can potentially contain an one (or more) active molecules, commonly referred to regions of interest (ROI). As an active molecule is described by its diffraction along a region of pixels, or Point Spread Function (PSF), the MRF will serve as a tool to obtain the potential pixels which are part of one (or more) PSFs. We consider two different states for a pixel, the 'On' and 'Off' state respectively, with the former denoting a PSF pixel of one (or more) active molecules, or signal, and the latter a pixel only subject to noise. Therefore,

$$x_{ij} = \begin{cases} 1, & \text{if the pixel is declared to be in 'On' state} \\ -1, & \text{if the pixel is declared to be in 'Off' state.} \end{cases} \quad (3.3.3)$$

We assume that all the necessary information about a pixel's state is fully described by the MRF, hence the intensities $\mathbf{y} = \{y_{ij}\}$ are conditionally independent given the field $\mathbf{x} = \{x_{ij}\}$. Therefore, regardless of the distributional assumption for the intensities, the likelihood reads,

$$f(\mathbf{y}|\mathbf{x}, \phi) = \prod_{i,j} f(y_{ij}|x_{ij}, \phi), \text{ for a set of parameters } \phi. \quad (3.3.4)$$

The choice of the distribution for y_{ij} is not unique, however we adopt a Gaussian distribution in agreement to what is most frequently used to describe the PSF of an active molecule. A number of examples are presented in Chapter 2 where Gaussian PSFs are fitted to ROIs. The novel SiMPa functions are introduced in the mean of the Gaussian distribution, according to the concept described in Section 3.2, accounting for the proportion of the single event intensity I a pixel should have based on the

position of the molecule within the pixel.

Denoting by $\boldsymbol{\phi}$ the set including all the parameters involved, the probability density of the intensities is,

$$f(y_{ij}|x_{ij}, \boldsymbol{\phi}) \sim N(y_{ij}|\mu(x_{ij}), \tau(x_{ij})), \quad (3.3.5)$$

with τ denoting the precision, hence $\text{Var}(y_{ij}|x_{ij}, \boldsymbol{\phi}) = 1/\tau(x_{ij})$. The mean and precision of $f(y_{ij}|x_{ij}, \boldsymbol{\phi})$ depend on the state of the respective pixel, identified by the underlying field, and with respect to the SiMPa functions are defined as,

$$\mu(x_{ij}) = \begin{cases} I(1 - \alpha_O(r, \theta, c)), & \text{if } x_{ij} = 1 \text{ ('On')} \\ I d, & \text{if } x_{ij} = -1 \text{ ('Off')} \end{cases} \quad \text{and } \tau(x_{ij}) = \tau_b, \quad (3.3.6)$$

where, recalling from Equation 3.2.1 in Section 3.2, the SiMPa functions are defined as,

$$\alpha_k = \begin{cases} c \left(d_1 - r \cos \left[\frac{(k-1)\pi}{4} + \theta \right] g(R, l) \right), & k = 1, 3, 5, 7 \\ c \left(\left((rg(R, l))^2 + d_2^2 + 2d_2 s_k r \cos \left(\frac{(k+1)\pi}{4} - \theta \right) g(R, l) \right)^{1/2} \right), & k = 2, 4, 6, 8, \\ 0, & k = 0 \end{cases}$$

and $\alpha_O(r, \theta, c)$ a single evaluation of the SiMPa functions α_k , $k = 0, \dots, 8$, which we discuss next.

The precision $\tau(x_{ij})$ in Equation (3.3.6) is defined to be the same no matter the state of the field and equal to $\tau(x_{ij}) = \tau_b$, where τ_b denotes the background error or camera noise. This choice can be supported by the fact that the camera noise can be naturally considered to be generated by the same mechanism for every pixel on the lattice (Section 2.2.2).

The mean $\mu(x_{ij})$ in Equation (3.3.6), is formed with respect to the the SiMPa func-

tions. Based on our proposed way of applying the SiMPa functions, as introduced in Section 3.2, the functions account for the proportion of a single event intensity I each pixel across the neighbourhood should be assigned, based on the position of the molecule within the central pixel. It should be highlighted here, that as the field identifies regions of interest, a declared 'On' pixel does not necessarily correspond to an active molecule directly. If this was the case, every pixel identified as 'On' by the field would be assigned with α_0 from the SiMPa functions, which implies an active molecule lying exactly on that pixel. On the contrary, as the SiMPa functions correspond to the behaviour of a PSF of an active molecule within a predetermined region, an identified 'On' pixel is assumed to belong to a PSF linked by the quantity $\alpha_O(r, \theta, c)$ in Equation (3.3.6). The notation α_O denotes an appropriate choice from $\alpha_0, \dots, \alpha_8$ based on the part of the PSF contributing to the pixel with respect to the molecule's position within the central pixel of the PSF.

Specifically, when a pixel is identified as 'On' its mean is defined by the quantity $I (1 - \alpha_O(r, \theta, c))$, whereas when a pixel is identified as 'Off', its mean is only related to background intensity as no PSF is contributing. The quantity $I d$ precisely represents the background intensity when assumed to be proportional to the single event intensity I , with d denoting the proportion. In a way, a fixed value for the proportion d corresponds to a uniform background, that is same average number of photons on each pixels as discussed in Section 2.2.2, with active molecules emitting a larger amount of photons. For that reason, we assume that the background intensity is proportional to the single event intensity.

The basic concept of our modelling lies on the idea that the intensities diffraction of a molecule on the lattice is fully described by the SiMPa functions. Consequently, a mechanism able to scan the lattice and obtain these diffractions, or the point spread functions (PSF) of active molecules, is considered. We introduce a moving region (MR), or a sliding window, which we allow to travel along the lattice adding the intensities of a predetermined region. Reasonably, we assume the predetermined region of this moving region to be the same size as the neighbourhood structure used

in the SiMPa functions. Since we denote the observed intensities by $\mathbf{y} = \{y_{ij}\}$, the moving region is defined as $\text{MR}_{ij} = \sum_{pq \in \partial y_{ij}} y_{pq}$ with i, j being the row and column of the pixel respectively and ∂y_{ij} the 8-neighbour system of the pixel as in Equation 3.3.2. Following the definition of the probability density in Equation (3.3.5), the probability density of the moving region given the states in the neighbourhood of x_{ij} is,

$$f \left(\text{MR}_{ij} = \sum_{pq \in \partial y_{ij}} y_{pq} \mid \partial x_{ij}, \phi \right) \sim \text{N} \left(\text{MR}_{ij} \mid \mu_{\text{MR}}, \tau_{\text{MR}} \right), \quad (3.3.7)$$

with τ denoting the precision and ϕ the set of all parameters included. Consequently, the mean and the precision of $f(\text{MR}_{ij} \mid x_{ij}, \phi)$ depend on the state of the field and consist of the appropriate sum of the Equations (3.3.6) using the conditional independence of \mathbf{y} given \mathbf{x} . Therefore Equation (3.3.7) reads in,

$$\text{MR}_{ij} \mid \partial x_{ij}, \phi \sim \text{N} \left(I \left(dN_{\{\partial x_{ij}=-1\}} + \sum_{v=1}^{N_{\{\partial x_{ij}=1\}}} (1 - \alpha_O(r_v, \theta_v, c)) \right), \tau_b / N_p \right), \quad (3.3.8)$$

where $N_{\{\partial x_{ij}=1\}}$ and $N_{\{\partial x_{ij}=-1\}}$ denote the number of pixels identified as 'On' and 'Off' in the neighbourhood of x_{ij} respectively. The variance of the likelihood is the sum of the individual variances with $\sigma_{\text{MR}}^2 = N_p \sigma_b^2$ or in terms of precision, $\tau_{\text{MR}} = \tau_b / N_p$ where N_p is the number of pixels included on the 8-neighbour system, here $N_p = 9$. The mean is the sum of the corresponding intensities described by the SiMPa functions with respect to the states of the pixels in the moving region.

Then, since \mathbf{y} are conditional independent given \mathbf{x} , the likelihood is calculated by multiplying the conditionally independent terms based on the MR which scans the frame, therefore,

$$f(\mathbf{MR} \mid \mathbf{x}, \phi) = \prod_{i,j} f(\text{MR}_{ij} \mid \partial x_{ij}, \phi), \quad (3.3.9)$$

where $\mathbf{MR} = \{\text{MR}_{ij}, i = 1, \dots, n \ \& \ j = 1, \dots, m\}$.

Briefly, let a moving region partially capturing a single PSF of a molecule according to the states of the pixels in the field. Then, the mean of the likelihood for this specific moving region is considered to be the sum of the number of 'Off' pixels, $N_{\{\partial x_{ij}=-1\}}$, multiplied by the background intensity, $I d$, and the appropriate SiMPa functions for the number of 'On' pixels, $N_{\{\partial x_{ij}=1\}}$, multiplied by the single event intensity, I .

One of our main goals, as discussed in Section 2.3, is to allow inference on closely located active molecules, which have their PSFs overlapping. In order to achieve this, we modify and generalise the mean μ_{MR} to be able to account for more than one PSF, in terms of allowing an identified 'On' pixel to be a product of more than a single PSF. Therefore,

$$\text{MR}_{ij} | \partial x_{ij}, \phi \sim \text{N} \left(\underbrace{I \left(d N_{\{\partial x_{ij}=-1\}} + \sum_{v=1}^{N_{\text{PSF}}} (n_{\text{PSF}_v} - \alpha_{\text{O}, n_{\text{PSF}_v}}(r_v, \theta_v, c)) \right)}_{\mu_{\text{MR}}}, \underbrace{\tau_b / N_p}_{\tau_{\text{MR}}} \right). \quad (3.3.10)$$

The part regarding the number of 'Off' pixels is unchanged, whereas the quantity accounting for the number of 'On' pixels is altered in such a way to account for multiple PSFs, the number of which is indicated by N_{PSF} . For each PSF falling within the moving region the SiMPa functions are obtained with respect to the number of pixels each one contributes to the MR, denoted by n_{PSF} . The sum of every SiMPa function for each separate PSF, results in the total contribution inside the moving region.

Example 3.3.1. For instance, let a moving region identified to have two 'On' pixels according to field. As presented in Figure 3.3.2, the grey filled pixels represent the 'On' pixels while the dashed lines two separate PSF's overlapping with each other,

hence $N_{\text{PSF}} = 2$, both contributing to the MR. The filled in circles within the central pixels of the PSF's denote the respective molecule's position. The number of 'Off' pixels contributing to the total intensity of the MR is $N_{\{x_{\partial ij}=-1\}} = 7$. In order to calculate the total contribution of the PSFs to the 'On' pixels, we individually take into consideration each PSF's segment, corresponding to $n_{\text{PSF}_1} = 2$ 'On' pixels for PSF_1 and $n_{\text{PSF}_2} = 1$ for PSF_2 .




Figure 3.3.2: Example of two Point Spread Functions (PSF) contributing in a moving region (MR). The black solid lines represent the moving region with the grey filled in pixels to have been identified as 'On' by the field. The dashed lines represent two different (black and blue) PSF's contributing to the moving region resulting in the two 'On' pixels case within the MR. The coloured circles denote the corresponding molecules of each PSF's respectively.

Regarding the prior distributions for each one of the parameters involved, we consider a Gaussian distribution for the single event intensity I and a Gamma distribution for the background precision (error) τ_b . The diffraction parameter c is bounded, as discussed in detail in Section 3.2.4, hence we use a Rescaled Beta distribution. The MRF \mathbf{x} is modelled using an auto-logistic model, first introduced by Besag [1974], with

$$\pi(\mathbf{x}|\{\beta_0, \beta_f\}) = \frac{\exp\{\beta_0 V_0(\mathbf{x}) + \beta_f V_f(\mathbf{x})\}}{Z(\{\beta_0, \beta_f\})} \quad (3.3.11)$$

The statistic $V_0(\mathbf{x})$ is defined as $V_0(\mathbf{x}) = \sum_{i,j} x_{ij}$ so positive or negative values of β_0 's tend to control the number of 'On' or 'Off' states, respectively. The statistic $V_f(\mathbf{x})$ is declared in physics as the energy function and is defined as $V_f(\mathbf{x}) = \sum_{pq \in \partial x_{ij}} x_{ij} x_{pq}$ with ∂x_{ij} denoting the neighbourhood of x_{ij} , β_f is a parameter defining the strength of neighbouring interaction with large positive values leading to realizations of \mathbf{x} having homogeneous patches of -1's or 1's and $Z(\{\beta_0, \beta_f\})$ is the normalizing constant. The full conditional distribution of x_{ij} using the 8-neighbour structure can be obtained using the Hammersley-Clifford theorem as,

$$\pi(x_{ij} | \partial x_{ij}, \{\beta_0, \beta_f\}) \propto \exp \left\{ \beta_0 x_{ij} + \beta_f \sum_{pq \in \partial x_{ij}} x_{ij} x_{pq} \right\}. \quad (3.3.12)$$

The parameters I , c , τ_b , β_0 and β_f are considered as global parameters since they rely on the entire frame, whereas the combinations of r and θ are local parameters corresponding to a neighbourhood of the PSF. The complete Bayesian framework for the update of the global parameters is done via the full-conditional distributions (presented in Section 3.3.4), along with inference on the local parameters r and θ based on our novel pattern-configuration concept which is introduced in the following Section 3.3.3 and analysed thoroughly. The general prior setting for the global and local parameters, reads in,

$$\begin{aligned} \text{Global Parameters} & \left\{ \begin{array}{l} I | \mu_I, \tau_I \sim \text{N}(I | \mu_I, \tau_I) \\ c | a_c, b_c \sim \text{RescaledBeta}(c | a_c, b_c, tr_c) \\ \tau_b | a_\tau, b_\tau \sim \text{Gamma}(\tau_b | a_\tau, b_\tau) \\ \pi(\mathbf{x} | b_0, b_f) \sim \text{Auto-logistic}(\mathbf{x} | \beta_0, \beta_f) \\ b_0 | m_{b_0}, \tau_{b_0} \sim \text{N}(b_0 | m_{b_0}, \tau_{b_0}) \\ b_f | m_{b_f}, \tau_{b_f} \sim \text{N}(b_f | m_{b_f}, \tau_{b_f}). \end{array} \right. \quad (3.3.13) \\ \text{Local Parameters} & \left\{ \pi(\{\text{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{ij}) \sim \text{Uniform}(1/N_{\text{Conf}_{ij}}) \right. \quad (3.3.14) \end{aligned}$$

3.3.3 Pattern-Configuration-Realisation concept

In this section we introduce our novel probabilistic individual frame inference scheme via a Bayesian framework. As discussed in Section 3.3.2, the mean μ_{MR} of a moving region (MR) depends on the states ('On' and 'Off') of the corresponding pixels in the field, which determine their contribution to the total intensity of the MR.

A main assumption, which is adopted throughout the thesis, is the restriction of no more than two PSFs allowed to overlap with each other. This can arguably be a reasonable assumption based on the natural of SMLM, where only a sparse subset of molecules is active on a single frame (Section 2.2), however an extension to more than two PSFs is straightforward (discussion on Chapter 6). In a way, it is probable to have two closely located molecules blinking simultaneously, resulting in their PSFs overlapping, but not likely to have three or more molecules active within a small distance at the same time.

We address two main aspects taken into consideration in order to conduct inference on the positions of the active molecules. The first one corresponds to the number of pixels, n_{PSF} , every PSF contributes to a MR, while the second one to the positions of the corresponding active molecules, described by r and θ , on each one of the PSFs. The following definitions are key components of our proposed probabilistic inference.

Definition 3.3.2. A *pattern* is a constructed structure within a moving region, consisting of a combination of 'On' and 'Off' pixels as identified by the field.

Let MR_{ij} be the moving region around a pixel (i, j) , $i = 1, \dots, n$ and $j = 1, \dots, m$ where n and m are the number of rows and columns of the frame, and ∂x_{ij} the states of the corresponding pixels as identified by the field. Then, a *pattern*, denoted with $P(\partial x_{ij})$, is a matrix of 1's and -1's with 1 and -1 denoting to 'On' and 'Off' states respectively.

Definition 3.3.3. A *configuration* is a combination of a number of PSFs, overlapping or not, which can result in a specific pattern with respect to a moving region. Let MR_{ij} be the moving region around a pixel (i, j) , $i = 1, \dots, n$ and $j = 1, \dots, m$ where n and m are the number of rows and columns of the frame, and $P(\partial x_{ij})$ the corresponding pattern. Then, a *configuration* for the MR_{ij} , denoted with $\{\text{Conf}\}_{ij}$ (or Conf), is a collection of the parts of the associated PSFs in order to form the pattern $P(\partial x_{ij})$, hence $\{\text{Conf}\}_{ij} = (n_{\text{PSF}_1}, n_{\text{PSF}_2}, \dots, n_{\text{PSF}_{N_{\text{PSF}}}})$, where N_{PSF} the number of PSFs and $n_{\text{PSF}_v} \leq N_{\text{PSF}}$ for every $v = 1, \dots, N_{\text{PSF}}$.

Number of pixels - n_{PSF}

The number of pixels, n_{PSF} , a PSF contributes in a MR is a quantity not directly observable, if overlapping PSFs are present. A pattern, as introduced in Definition 3.3.2, consists of a structure of 'On' and 'Off' pixels as identified by the field, as a potential product of different parts of one or more PSFs. With respect to our main assumption, of allowing up to two PSFs to overlap, there is a finite number of configurations every pattern can give rise to. The number of PSFs in a configuration, as in Definition 3.3.3, is denoted by N_{PSF} , in agreement with Equation (3.3.10).

Example 3.3.2. A pattern associated with all possible configurations under our assumption is presented in Figure 3.3.3. A MR has been assigned to have three 'On' pixels according to the field, under the structure of two 'On' on the top right corner

and one 'On' on the bottom left corner, hence the pattern $P(\partial x_{ij}) = \begin{bmatrix} -1 & 1 & 1 \\ -1 & -1 & -1 \\ 1 & -1 & -1 \end{bmatrix}$.

With respect to main assumption of no more than two PSFs overlapping, there are six different configurations which can potentially form this specific pattern within the MR. For instance, a possible configuration consists of two individual PSFs without any overlapping, denoted by Conf_1 , where the total number of PSFs is $N_{\text{PSF}} = 2$. The number of pixels from the first PSF is $n_{\text{PSF}_1} = 2$ and from the second one is $n_{\text{PSF}_2} = 1$. Another case, allows four individual PSFs overlapping in a specific way in

pairs of two, as in Conf_6 with $N_{\text{PSF}} = 4$. For this configuration, the number of pixels from the first PSF is $n_{\text{PSF}_1} = 2$ and from its overlapping one is $n_{\text{PSF}_2} = 1$, while the second one has $n_{\text{PSF}_3} = 1$ with its overlapping one to have $n_{\text{PSF}_4} = 1$. All six different configurations under this pattern with their respective pixels, n_{PSF} , contributing to the MR are shown with different colouring for every different PSF.




Figure 3.3.3: Example of a *pattern* along with every possible *configurations* that can be formed from, with respect to the assumption of up to two PSFs allowed to overlap. On the left panel of the figure, a pattern is declared with the grey filled pixels on a MR, with the dashed (coloured) lines representing PSFs forming the different configurations. The right panel contains all the possible configurations with their corresponding pixels contributing.

An algorithm has been constructed in order to obtain all different configurations for every different pattern that can be formed within a MR, with respect to our assumption of allowing up to two PSFs to overlap. The number of all the different patterns for every 'On' case are shown in Table 3.3.1, along with the total number of configurations that can exist on each pattern. It should be noted that every pattern has a different number of configurations based on its structure.

'On' Pixels	1	2	3	4	5	6	7	8	9
Different Patterns	4	14	28	45	52	46	28	9	1
Number of Configurations	8	48	208	648	1584	2684	3264	2492	2152

Table 3.3.1: Number of different patterns that can be formed when the MR size is 3×3 pixels (same as PSF's size) and total number of configurations that can exist (with respect to the assumption of up to two PSF allowed to coincide). Each pattern has a different number of configurations which is not shown here.

Recalling from Equation (3.3.10) the probability density of the MR with respect to the SiMPa functions is,

$$\text{MR}_{ij} | \partial x_{ij}, \phi \sim \text{N} \left(I \left(dN_{\{\partial x_{ij} = -1\}} + q \sum_{v=1}^{N_{\text{PSF}}} (n_{\text{PSF}_v} - \alpha_{O, n_{\text{PSF}_v}}(r_v, \theta_v, c)) \right), \tau_b / N_p \right).$$

Each pattern formed by the field on every MR, is evaluated at the respective configuration, described by the appropriate number of pixels, n_{PSF} , for every PSF contributing. This is highlighted in the Equation (3.3.10) with the red color. As can be seen, the SiMPa functions, $\alpha_{O, n_{\text{PSF}_v}}(r, \theta, c)$, depends on the number of pixels from each PSF, n_{PSF} . This implies that we want to take into consideration only the parts of the PSF contributing to the MR, based on the SiMPa functions evaluated at the position of the molecule denoted by r and θ . For instance, if a MR is identified to have four 'On' pixels produced by a single PSF, i.e. $n_{\text{PSF}} = 4$, then we need the corresponding four pixels from the SiMPa functions of that PSF, according to the position of the molecule. The way to deal with the SiMPa functions for every MR is what we address next.

Positions of molecules - r and θ

The second aspect corresponds exactly to the choice of the appropriate SiMPa functions, consequently inference on position of the molecules denote by the distance r on direction θ . As discussed in Section 3.2, the local parameters r and θ live on the

continuous space and satisfy the restrictions $r|\cos(\theta)| < \frac{l}{2} - R$ and $r|\sin(\theta)| < \frac{l}{2} - R$, where R is the radius of the molecule and l the side length of the pixel. However, even though this is the natural way to proceed, it might not be feasible to measure a molecule's position with such precision. One can argue that the measurements obtained by SMLM can allow localisation reaching up to a certain resolution, hence limiting the precision and accuracy on estimates of r and θ . Additionally, we seek a methodology in which inference on r and θ can be both efficient and not extremely computationally intensive.

A way to tackle this issue is by introducing a discretisation for r and θ . This practically means that instead of trying to measure the exact molecule's position within the pixel, which can be declared naturally intractable, we want to locate it with a precision described by a combination of ranges for r and θ . We consider a discretisation of θ to be all angles from 0 to 2π with a step of $\pi/8$, that is $\theta = k\frac{\pi}{8}$, $k=0,1,\dots,16$. Regarding the distances r , three different values are explored on every angle θ with $r = 0$ (origin of the SiMPA functions), $r = \frac{r_{\max,\theta}}{2}$ (half distance from the origin when angle is θ) and $r = r_{\max,\theta}$ (maximum distance from the origin when angle is θ). Under this modification, one is trying to approach the molecule's position in a discrete space defined by a total number of angles $N_\theta = 16$ and total number of distances $N_r = 33$, i.e. $r = \theta = 0$ and two distances (half and maximum) for every θ .

Definition 3.3.4. A *realisation* is a specific combination of r and θ from a predetermined discretisation, determining the position of a molecule within a pixel and denoted with $\{r, \theta\}$.

Definition 3.3.5. A *complete localisation set* for any moving region MR_{ij} , $i = 1, \dots, n$ (rows of frame) and $j = 1, \dots, m$ (columns of frame), with respect to the respective pattern ∂x_{ij} , is the collection of a configuration, $\{\text{Conf}\}_{ij}$, associated with appropriate realisations, $\{\mathbf{r}, \boldsymbol{\theta}\}_{ij}$, denoted by $\{\text{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{ij} = \left\{ \text{Conf}_{ij}, \{r, \theta\}_{ij(1)}, \dots, \{r, \theta\}_{ij(v)} \right\}$, where $v = 1, \dots, N_{\text{PSF}}$ to be denoting the number of PSFs on the configuration Conf_{ij}

of pixel (i, j) . A *complete counting set* is the simplification of the complete localisation set where each realisation is at the origin of the SiMPa functions, i.e. $\mathbf{r} = \boldsymbol{\theta} = \mathbf{0}$, denoted by $\{\text{Conf}, \mathbf{r} = \mathbf{0}, \boldsymbol{\theta} = \mathbf{0}\}_{ij} \equiv \{\text{Conf}\}_{ij}$.

A main property of the SiMPa functions is the periodicity of their behaviour. The SiMPa functions performance under a specific direction θ_1 , is identical to the symmetric direction θ_2 within the pixel, with respect to the axis of the coordinate system. The only difference is that the whole system is rotated depending on the symmetry of the corresponding angles. A graphical illustration can be seen in Figure 3.3.4, where a direction $\theta_1 = \pi/8$ (a) and $\theta_2 = 5\pi/8$ (b) both on a distance $\frac{r_{\max, \theta}}{2}$ have been chosen. On part (a) of the figure, the SiMPa functions evaluated at r on angle θ_1 are shown, with part (b) containing the symmetric realisation on the 2nd quadrant with θ_2 along with the equivalent SiMPas from the realisation with θ_1 on part (a).

Therefore, as the chosen discretisation consists of multiples of $\pi/8$, we can reduce the amount of total direction θ and distances r to the ones which exist only on the first quadrant of the coordinate system, hence $N_\theta = 5$ ($0, \pi/8, \pi/4, 3\pi/8, \pi/2$) with $N_r = 15$. Based on to the rotation property of the SiMPa functions, we are able to identify r and θ only by evaluating their performance on the first quadrant. This stands for any chosen discretisation as well when the continuous space is considered and can noticeably reduce the complexity of the localisation procedure, however not without complications, as discussed later in the Section.




Figure 3.3.4: Example of the periodicity of the SiMPa functions when the system is rotated. (a) A molecule has been placed r units with direction $\theta_1 = \pi/8$ from the origin. (b) A molecule has been placed r units with direction $\theta_2 = 5\pi/8$ from the origin. The equivalence of the SiMPa functions when the angle θ_1 from (a) is outlined.

With respect to the assumption of allowing only up to two PSFs to overlap, there is a finite number of configurations that can be formed by different patterns, with the total number of configuration remaining unchanged as shown in Table 3.3.1. On the contrary, as every configuration consists of a number of PSFs which contribute to the MR, each one is associated with a realisation which alters the behaviour of the SiMPa functions. Therefore, the appropriate parts of the PSFs, described by n_{PSF} , of the corresponding realisation need to be taken into consideration. As already discussed, due to the periodicity of the SiMPa functions only the first quadrant of the coordinate system is considered, however two immediate questions arise.

Firstly, given a realisation for each PSF involved in a MR, how can we obtain the appropriate regions which have fell within the MR. These regions are represented by the number of pixels n_{PSF} every PSF contributes, associated with the intensities of the realisation based on the SiMPa functions. Secondly, when allowing an additional

realisation, in terms of adding a combination of r and θ within the model, how much the complexity of the model is increased with the effect of complicating the position inference.

The complexity is considered in terms of the number of realisations within each configuration. Sticking to the discretisation $\theta = k\frac{\pi}{8}$, $k=0,1,2,3,4$ (1st quadrant) and $r = 0$ (origin of the SiMPa functions), $r = \frac{r_{\max,\theta}}{2}$ (half distance from the origin when angle is θ) and $r = r_{\max,\theta}$ (maximum distance from the origin when angle is θ), we investigate the behaviour of a single molecule, or equivalent an individual PSF. The symmetric directions of θ are also included based on the periodicity of the SiMPa functions.

For a single realisation, the number of possible combinations a PSF's parts can land on a MR, denoted with $N_{\{\text{MR},n_{\text{PSF}}\}}$, depends on the number of pixels, n_{PSF} , this PSF contributes to the MR. On Table 3.3.2, we present the relation of $N_{\{\text{MR},n_{\text{PSF}}\}}$ and n_{PSF} for a single realisation of an individual PSF, apart from the case where $r = 0$ for any direction θ . In that case, $N_{\{\text{MR},n_{\text{PSF}}\}}$ is reduced from the numbers on the table, due to the symmetric diffraction of the SiMPa functions without implying a different position of the molecule but $r = \theta = 0$ (as can be seen in Figure 3.2.4 (a)). Therefore, $N_{\{\text{MR},n_{\text{PSF}}\}} = 1$ regardless of n_{PSF} .

n_{PSF}	1	2	3	4	5	6	7	8	9
$N_{\{\text{MR},n_{\text{PSF}}\}}$	4	8	4	4	0	4	0	0	1

Table 3.3.2: Number of possible combinations a PSF's parts can fall within a MR, $N_{\{\text{MR},n_{\text{PSF}}\}}$ for a single realisation. This number depends on the number of pixels the PSF contributes to the MR, denoted by n_{PSF} .

From Table 3.3.2, for instance, when $n_{\text{PSF}} = 3$, there are $N_{\{\text{MR},n_{\text{PSF}}\}} = 4$ different PSF parts that might have fell on a MR according to one realisation. These parts can be one of $\{\alpha_8, \alpha_1, \alpha_2\}$, $\{\alpha_2, \alpha_3, \alpha_4\}$, $\{\alpha_4, \alpha_5, \alpha_6\}$ or $\{\alpha_6, \alpha_7, \alpha_8\}$ evaluated at the corresponding realisation. We note here, that there are no cases for $N_{\{\text{MR},n_{\text{PSF}}\}}$ when

$n_{\text{PSF}} = 5, 7$ or 8 , since it is not feasible a 3×3 PSF to have 5,7 or 8 pixels landing within a 3×3 MR.

Consequently, we can calculate the number of possible combinations multiple PSFs parts can fall within the MR to be $\prod_{v=1}^{N_{\text{PSF}}} N_{\{\text{MR}, n_{\text{PSF}_v}\}}$ for a single realisation, leading to a generalisation of the different cases a configuration can be formed, from a number of realisations on the 1st quadrant and their respective symmetries, to be,

$$N_{\text{conf}} = \prod_{v=1}^{N_{\text{PSF}}} (N_{r\theta} N_{\{\text{MR}, n_{\text{PSF}_v}\}} + \mathbf{1}_{0 \in r}), \quad (3.3.15)$$

where $N_{r\theta}$ is the total number of realisations on the first quadrant of the coordinate system excluding those with $r = 0$ and $\mathbf{1}_{0 \in r} = 1$, when $r = 0$ (origin of the SiMPa functions) and $\mathbf{1}_{0 \in r} = 0$ otherwise. As a reminder, we focus on the first quadrant due to the periodicity of the SiMPa functions, where the symmetric realisations have been considered.

Revisiting Example 3.3.2 based on Figure 3.3.3, the configuration Conf_3 , for instance, is produced by three PSFs, these are PSF_1 , PSF_2 and PSF_3 , each one contributing $n_{\text{PSF}_1} = 2$, $n_{\text{PSF}_2} = 1$ and $n_{\text{PSF}_3} = 2$ pixels to the MR respectively. Every individual PSF can contain a molecule placed in a different position within the central pixel. If we consider two realisations from the discretisation, say $\theta_1 = \pi/8$ and $r_1 = \frac{r_{\max, \theta = \pi/8}}{2}$ and $r = \theta = 0$, then for the configuration Conf_3 we get $N_{\text{conf}} = (4+1) \times (8+1) \times (8+1) = 405$ by using Table 3.3.2 and Equation 3.3.15. This gives us the total cases the configuration Conf_3 can be formed, if each one of the molecules involved are placed on the chosen realisations or their symmetric positions, that is $k \frac{\pi}{8}$, $k=1,5,7,15$.

Next, we focus on the way to identify the quadrant each molecule belongs to, hence their position declared by r and θ , which relies on the relative position of the MR and the corresponding number of pixels, n_{PSF} , each PSF contributes. We define four separate areas around a MR a PSF can lie onto which contribute to the decision

on the quadrant of the molecules. Depending on the number of pixels, n_{PSF} , a PSF contributes to the MR, we consider two different cases. In Figure 3.3.5, the areas when $n_{\text{PSF}} = 1, 2, 4$ are presented, using only the case where $n_{\text{PSF}} = 1$, whereas in Figure 3.3.6 the equivalent ones when $n_{\text{PSF}} = 3, 6$, using the case where $n_{\text{PSF}} = 3$. In both figures, the red solid lines associated with the numbering label represent the areas relative to the MR, with the blue solid boxes denoting the quadrants of the central pixel of each PSF. The PSFs are drawn with the black dashed with the SiMPa functions contributing to have been outlined.




Figure 3.3.5: Implementation of case where a single pixel of a PSF lands on an MR. The molecule lies anywhere inside the first quadrant of the coordinate system, denoted with the blue colored 1st. Regarding the position of the PSF with respect to the MR (regions denoted with the red lines and circled numbers identified by the configuration chosen), different quantities of the SiMPa functions account for the respective quadrants the molecule lies in.




Figure 3.3.6: Implementation of case where three pixels of a PSF lands on an MR. The molecule lies anywhere inside the first quadrant of the coordinate system, denoted with the blue coloured label '1st'. Regarding the position of the PSF with respect to the MR (regions denoted with the red lines and circled numbers identified by the configuration chosen), different quantities of the SiMPa functions account for the respective quadrants the molecule lies in.

As already discussed, we assume a molecule lies somewhere within the 1st quadrant. Focusing on the case where $n_{\text{PSF}} = 1$ in Figure 3.3.5, the number of possible combinations a PSF's parts can fall within a MR is $N_{\{\text{MR}, n_{\text{PSF}_v}\}} = 4$ (also declared on Table 3.3.2). Based on the SiMPa functions chosen and the corresponding location of the PSF, we are able to identify symmetric realisations, producing identical performance, just by rotating the SiMPa functions by multiples of $\pi/2$ with respect to the origin.

This implies identification of the appropriate direction θ under the unchanged distance r on the symmetric realisation. The choice of the quadrants for the cases with $n_{\text{PSF}} = 1, 2, 4$ are summarised on Table 3.3.3. These tables show all combinations of the SiMPa functions when a realisation from the first quadrant is chosen for all four different areas in Figure 3.3.5, along with the symmetric realisations when the system is rotated.

1 Pixel	4 Pixels	①	②	③	④	2 Pixels	①	②	③	④
α_2	$\alpha_0\alpha_1\alpha_2\alpha_3$	1 st	2 nd	3 rd	4 th	$\alpha_1\alpha_2$	1 st	2 nd	3 rd	4 th
α_4	$\alpha_0\alpha_3\alpha_4\alpha_5$	2 nd	3 rd	4 th	1 st	$\alpha_2\alpha_3$	1 st	2 nd	3 rd	4 th
α_6	$\alpha_0\alpha_5\alpha_6\alpha_7$	3 rd	4 th	1 st	2 nd	$\alpha_3\alpha_4$	2 nd	3 rd	4 th	1 st
α_8	$\alpha_0\alpha_1\alpha_7\alpha_8$	4 th	1 st	2 nd	3 rd	$\alpha_4\alpha_5$	2 nd	3 rd	4 th	1 st
						$\alpha_5\alpha_6$	3 rd	4 th	1 st	2 nd
						$\alpha_6\alpha_7$	3 rd	4 th	1 st	2 nd
						$\alpha_7\alpha_8$	4 th	1 st	2 nd	3 rd
						$\alpha_1\alpha_8$	4 th	1 st	2 nd	3 rd

Table 3.3.3: Combinations of the SiMPa functions when a PSF contributes one or four pixels on a MR, along with the corresponding quadrants the molecule lies in, with respect to rotation and PSF's relative position with the MR.

On a similar manner, the quadrants when $n_{\text{PSF}} = 3$ or 6 can be obtained, using the four different areas in Figure 3.3.6. All the combinations of the SiMPa functions for a molecule on the 1st quadrant, along with the symmetric realisation are summarised on Table 3.3.4.

3 Pixels	6 Pixels	①	②	③	④
$\alpha_1\alpha_2\alpha_8$	$\alpha_0\alpha_1\alpha_2\alpha_3\alpha_7\alpha_8$	1 st	2 nd	3 rd	4 th
$\alpha_2\alpha_3\alpha_4$	$\alpha_0\alpha_1\alpha_2\alpha_3\alpha_4\alpha_5$	2 nd	3 rd	4 th	1 st
$\alpha_4\alpha_5\alpha_6$	$\alpha_0\alpha_3\alpha_4\alpha_5\alpha_6\alpha_7$	3 rd	4 th	1 st	2 nd
$\alpha_6\alpha_7\alpha_8$	$\alpha_0\alpha_1\alpha_5\alpha_6\alpha_7\alpha_8$	4 th	1 st	2 nd	3 rd

Table 3.3.4: Combinations of the SiMPa functions when a PSF contributes three or six pixels on a MR, along with the corresponding quadrants the molecule lies in, with respect to rotation and PSF's relative position with the MR.

As an example, for the area ① in Figure 3.3.5 the choice of α_2 corresponds to a realisation on the 1st quadrant, since we only consider realisations in there. Applying the rotation, α_4 corresponds to the symmetric realisation on the 2nd quadrant, α_6 on the 3rd and α_8 on the 4th. Similarly, for the area ② in Figure 3.3.6, $\{\alpha_6, \alpha_7, \alpha_8\}$ denotes a realisation on the 1st quadrant, where rotation gives $\{\alpha_8, \alpha_1, \alpha_2\}$ for 2nd quadrant, $\{\alpha_2, \alpha_3, \alpha_4\}$ for 3rd quadrant and $\{\alpha_4, \alpha_5, \alpha_6\}$ for 4th quadrant. Therefore, we can argue that identification of the relative location of the PSF and the MR along with the corresponding SiMPa functions, can be done by choosing the appropriate quadrant of the molecule, with respect to realisations on the 1st quadrant.

3.3.4 Individual frame inference

In this section we introduce our novel pattern-configuration-realisation concept within a probabilistic scheme, leading to our proposed individual frame inference scheme based on the SiMPa functions. In its simplest form we consider $r = \theta = 0$, in which the localisation of any molecule is always fixed at the center of the pixel (origin of SiMPa functions), resulting in what we call *individual frame counting scheme*. Allowing specific positions for molecules within pixels based on the possible realisations (Definition 3.3.4) creates our *individual frame localisation scheme*. In both schemes, the update of the parameters is based on the full-conditional distributions. A special case Gibbs step is introduced for the configurations formed in a every MR, as well as to draw positions of molecules, described by r and θ . The entire parameter update based on the full conditional distributions is presented in Section 3.3.4.1, followed by the probabilistic individual frame counting and localisation schemes in Section 3.3.4.2.

3.3.4.1 Parameter updating via full conditionals

A Metropolis-Hastings (M-H) within Gibbs sampling algorithm is implemented with standard Gibbs steps for the single event intensity I , the background noise τ_b and

M-H for the Markov random field \mathbf{x} and power of diffraction parameter c . Let

$$\begin{aligned}\phi &= \left\{ I, \tau_b, c, \mathbf{x}, \{\text{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{ij} \right\}, \\ \mu_{\text{MR}} &= I \left(dN_{\{\partial x_{ij}=-1\}} + \sum_{v=1}^{\text{NPSF}} (n_{\text{PSF}_v} - \alpha_{O, n_{\text{PSF}_v}}(r_v, \theta_v, c)) \right), \\ \tau_{\text{MR}} &= \tau_b / N_p,\end{aligned}$$

where $\{\text{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{ij}$ the complete localisation set for a moving region MR_{ij} , $i = 1, \dots, n$ and $j = 1, \dots, m$ the rows and columns of the frame respectively (Definition 3.3.5).

First, we obtain the full conditional distributions for the global parameters I , τ_b , c and \mathbf{x} , using the prior distributions in Equation (3.3.13) and the MR likelihood in Equation (3.3.9),

Update of I

The full conditional of the single event intensity I is a Gaussian distribution,

$$\begin{aligned}p(I | \phi_{\{-I\}}, \mu_I, \tau_I) &\propto \prod_{\forall i,j} \text{N}(\text{MR}_{ij} | \mu_{\text{MR}}, \tau_{\text{MR}}) \times \text{N}(I | \mu_I, \tau_I) \\ &\sim \text{N} \left(\frac{\tau_I \mu_I + \frac{\tau_b}{N_p} \sum_{\forall ij} \text{MR}_{ij} B_{ij}}{\tau_I + \frac{\tau_b}{N_p} \sum_{\forall ij} (B_{ij})^2}, \tau_I + \frac{\tau_b}{N_p} \sum_{\forall ij} (B_{ij})^2 \right), \quad (3.3.16)\end{aligned}$$

where $B_{ij} = dN_{\{\partial x_{ij}=-1\}} + \sum_{v=1}^{\text{NPSF}} (n_{\text{PSF}_v} - \alpha_{O, n_{\text{PSF}_v}}(r_v, \theta_v, c))$.

Update of τ_b

The full conditional of precision τ_b is a Gamma distribution,

$$p(\tau_b | \phi_{\{-\tau_b\}}, a_{\tau_b}, b_{\tau_b}) \propto \prod_{\forall i,j} \text{N}(\text{MR}_{ij} | \mu_{\text{MR}}, \tau_{\text{MR}}) \times \text{Gamma}(\tau_b | a_{\tau_b}, b_{\tau_b})$$

$$\sim \text{Gamma} \left(a_\tau + \frac{N_{\text{MR}}}{2}, \frac{\sum_{\forall ij} (\text{MR}_{ij} - \mu_{\text{MR}})^2}{2N_p} + b_\tau \right), \quad (3.3.17)$$

where $N_{\text{MR}} = n \times m$ is the number of pixels in the frame with n and m the number of rows and columns respectively.

Update of c

The full conditional of power of diffraction parameter c is not in a closed form therefore a Metropolis-Hastings step is considered, with

$$p(c | \phi_{\{-c\}}, a_c, b_c, tr_c) \propto \prod_{\forall i,j} \text{N}(\text{MR}_{ij} | \mu_{\text{MR}}, \tau_{\text{MR}}) \times \text{RescaledBeta}(c | a_c, b_c, tr_c), \quad (3.3.18)$$

where tr_c is the rescale parameter. We will use a rescaled Beta with mode at the current value as the proposal distribution.

Update of \mathbf{x}

Regarding the update of the field \mathbf{x} , the full conditional distribution of x_{ij} reads in,

$$p(x_{ij} | \phi_{\{-x_{ij}\}}, \beta_0, \beta_f) \propto f(y_{ij} | \mu(x_{ij}), \tau(x_{ij})) \pi(x_{ij} | \partial x_{ij}, \{\beta_0, \beta_f\}), \quad (3.3.19)$$

where

$$\mu(x_{ij}) = \begin{cases} I (1 - \alpha_0(r, \theta, c)), & \text{if } x_{ij} = 1 \text{ ('on')} \\ I d, & \text{if } x_{ij} = -1 \text{ ('off')} \end{cases} \quad \text{and } \tau(x_{ij}) = \tau_b,$$

as in Equation (3.3.6) and $\pi(x_{ij} | \partial x_{ij}, \{\beta_0, \beta_f\}) \propto \exp \left\{ \beta_0 x_{ij} + \beta_f \sum_{pq \in \partial x_{ij}} x_{ij} x_{pq} \right\}$ as in Equation (3.3.12). Therefore, a pixel is identified as 'On' in the field with

probability,

$$p_{\text{on}} = \frac{\exp\left\{-\frac{\tau_b}{2}(y_{ij} - \mu(x_{ij} = 1))^2\right\} \exp\left\{\beta_0 + \beta_f \sum_{pq \in \partial x_{ij}} x_{pq}\right\}}{\sum_{k=1,-1} \exp\left\{-\frac{\tau_b}{2}(y_{ij} - \mu(x_{ij} = k))^2\right\} \exp\left\{k \left(\beta_0 + \beta_f \sum_{pq \in \partial x_{ij}} x_{pq}\right)\right\}}, \quad (3.3.20)$$

where $p_{\text{on}} = \Pr [x_{ij} = 1 \mid \boldsymbol{\phi}_{\{-x_{ij}\}}, \beta_0, \beta_f]$.

Update of $\{\beta_0, \beta_f\}$

The update of the hyperparameters β_0 and β_f of the field is done using the pseudo-likelihood estimation in order to avoid the calculation of the normalising constant $Z(\{\beta_0, \beta_f\}) = \sum_{\mathbf{x}} \exp\{\beta_0 V_0(\mathbf{x}) + \beta_f V_f(\mathbf{x})\}$. Calculating the normalising constant is very computationally intensive as it consists of 2^{RC} terms. Therefore,

$$\begin{aligned} p(\{\beta_0, \beta_f\} \mid \mathbf{x}) &\propto \pi(\mathbf{x} \mid \{\beta_0, \beta_f\}) \pi(\beta_0) \pi(\beta_f) \\ &= \frac{\exp\{\beta_0 V_0(\mathbf{x}) + \beta_f V_f(\mathbf{x})\}}{Z(\{\beta_0, \beta_f\})} \pi(\beta_0) \pi(\beta_f) \\ &\approx \prod_{i,j} \pi(x_{ij} \mid \partial x_{ij}, \{\beta_0, \beta_f\}) \pi(\beta_0) \pi(\beta_f), \end{aligned} \quad (3.3.21)$$

where the pseudo-likelihood estimation [Besag, 1974], reads in,

$$\pi(\mathbf{x} \mid \{\beta_0, \beta_f\}) \approx \prod_{i,j} \pi(x_{ij} \mid \partial x_{ij}, \{\beta_0, \beta_f\}). \quad (3.3.22)$$

3.3.4.2 Configurations and/or realisations updating

Given our Bayesian approach, the updating of the configurations for the localisation and counting schemes is probabilistic, and it takes into account the appropriate patterns within each MR. The former corresponds to the individual frame localisation

scheme which relies on configurations associated with realisations (Definitions 3.3.3 and 3.3.4), while the latter to the individual frame counting scheme in which every realisation is set to be at the origin of the SiMPa functions, hence $r = \theta = 0$.

For every moving region MR_{ij} and according to Definition 3.3.5, we use the complete localisation set $\{\text{Conf}, r, \theta\}_{ij}$ and the complete counting set $\{\text{Conf}\}_{ij}$ for the individual frame localisation and counting schemes respectively. For the former the full conditional distribution of a complete localisation set is obtained as,

$$\begin{aligned} P_{ij} &\equiv p\left(\{\text{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{ij} \mid \text{MR}_{ij}, \mathbf{x}, I, c, \tau_b\right) \\ &\propto f\left(\text{MR}_{ij} \mid \mathbf{x}, I, c, \tau_b, \{\text{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{ij}\right) \pi(\{\text{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{ij}) \\ &\propto \text{N}(\text{MR}_{ij} \mid \mu_{\text{MR}}, \tau_{\text{MR}}) \text{Uniform}(1/N_{\text{conf}_{ij}}), \end{aligned} \quad (3.3.23)$$

where a uniform prior is assigned for each complete localisation set (as in Equation (3.3.14)). For every MR_{ij} , the corresponding pattern ∂x_{ij} (Definition 3.3.2) is included within the field \mathbf{x} and the quantity $N_{\text{conf}_{ij}}$ denotes the number of possible formations of the corresponding configuration based on the chosen discretisation of r and θ , defined in Equation (3.3.15). Similarly, the full conditional distribution of a complete counting set is obtained as,

$$\begin{aligned} C_{ij} &\equiv p\left(\{\text{Conf}\}_{ij} \mid \text{MR}_{ij}, \mathbf{x}, I, c, \tau_b\right) \\ &\propto f\left(\text{MR}_{ij} \mid \mathbf{x}, I, c, \tau_b, \{\text{Conf}\}_{ij}\right) \pi(\{\text{Conf}\}_{ij}) \\ &\propto \text{N}(\text{MR}_{ij} \mid \mu_{\text{MR}}, \tau_{\text{MR}}) \text{Uniform}(1/N_{\text{conf}_{ij}}). \end{aligned} \quad (3.3.24)$$

The mean μ_{MR} and τ_{MR} in Equations (3.3.23) and (3.3.24) can be obtained from Equation (3.3.10), with μ_{MR} evaluated respectively at the red coloured parts below,

$$\text{Localisation:} \quad \mu_{\text{MR}} = I\left(dN_{\{\partial x_{ij}=-1\}} + \sum_{v=1}^{\text{NPSF}} (n_{\text{PSF}_v} - \alpha_{\text{O}, n_{\text{PSF}_v}}(r_v, \theta_v, c))\right).$$

$$\text{Counting: } \mu_{\text{MR}} = I\left(dN_{\{\partial x_{ij}=-1\}} + \sum_{v=1}^{N_{\text{PSF}}} (n_{\text{PSF}_v} - \alpha_{0, n_{\text{PSF}_v}}(r_v = 0, \theta_v = 0, c))\right).$$

The full conditional distributions P_{ij} and C_{ij} are evaluated at every possible formation in $N_{\text{Conf}_{ij}}$, denoted with $N_{\text{Conf}_{ij}(u)}$ and $N_{\text{Conf}_{ij}(w)}$, for the counting and localisation schemes respectively. Hence, both P_{ij} and C_{ij} have a finite support and thus each full conditional distribution is multinomial with probabilities proportional to $P_{ij(1)}, \dots, P_{ij(u)}$ and $C_{ij(1)}, \dots, C_{ij(w)}$ respectively. Both the probabilistic schemes are introduced within a Gibbs step with probabilities obtained as,

Localisation scheme:

$$\left\{ \begin{array}{l} P_{ij(1)} = p\left(\{\text{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{ij(1)} \mid \text{MR}_{ij}, \mathbf{x}^{(q)}, I^{(q)}, c^{(q)}, \tau_b^{(q)}\right) \\ \vdots \\ P_{ij(u)} = p\left(\{\text{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{ij(u)} \mid \text{MR}_{ij}, \mathbf{x}^{(q)}, I^{(q)}, c^{(q)}, \tau_b^{(q)}\right) \end{array} \right\} \rightarrow \left\{ \begin{array}{l} P_{ij(1)} / \sum_{h=1}^u P_{ij(h)} \\ \vdots \\ P_{ij(u)} / \sum_{h=1}^u P_{ij(h)} \end{array} \right\}.$$

Counting scheme:

$$\left\{ \begin{array}{l} C_{ij(1)} = p\left(\{\text{Conf}\}_{ij(1)} \mid \text{MR}_{ij}, \mathbf{x}^{(q)}, I^{(q)}, c^{(q)}, \tau_b^{(q)}\right) \\ \vdots \\ C_{ij(w)} = p\left(\{\text{Conf}, \}_{ij(w)} \mid \text{MR}_{ij}, \mathbf{x}^{(q)}, I^{(q)}, c^{(q)}, \tau_b^{(q)}\right) \end{array} \right\} \rightarrow \left\{ \begin{array}{l} C_{ij(1)} / \sum_{h=1}^w C_{ij(h)} \\ \vdots \\ C_{ij(w)} / \sum_{h=1}^w C_{ij(h)} \end{array} \right\}.$$

As a numeric illustration, we revisit Example 3.3.1 where the possible configurations are shown in Figure 3.3.7. Let a discretisation of r and θ to consist of $\theta = \{0, \pi/8, \pi/4, 3\pi/8\}$ (1st quadrant only) and $r = \{0, r_{\max, \theta}/2\}$, then $N_{r\theta}$, the total number of realisations on the 1st quadrant without the ones with $r = 0$, is $N_{r\theta} = 3$. In order to implement our localisation scheme, we need to evaluate P_{ij}

for every $\{\text{Conf}, r, \theta\}_{ij}$, hence the total number of complete localisation set, using Equation (3.3.15), Table 3.3.2 and configurations in Figure 3.3.7, are

$$\left. \begin{aligned} N_{\text{Conf}_1} &= (8 \times 3 + 1) = 25 \\ N_{\text{Conf}_2} &= (8 \times 3 + 1) \times (8 \times 3 + 1) = 625 \\ N_{\text{Conf}_3} &= (8 \times 3 + 1) \times (4 \times 3 + 1) = 325 \end{aligned} \right\} \rightarrow N_{\text{Conf}_{(u)}} = 975.$$

This wraps up in $P_{ij(1)}/\sum_{h=1}^{975} P_{ij(h)}, \dots, P_{ij(975)}/\sum_{h=1}^{975} P_{ij(h)}$ which allows to draw a localisation for MR_{ij} . Due to the periodicity described before, the positions drawn are realisations from $\theta = k\frac{\pi}{8}$ with $k = 0, \dots, 16$ with the same distances $r = \{0, r_{\max, \theta}/2\}$, by taking into account the relative positions of the MR to the drawn complete localisation set.




Figure 3.3.7: Simple example of a *pattern* along with every possible *configuration* it can arise from, with respect to the assumption of no more than two PSFs allowed to overlap. On the left panel of the figure, a pattern is declared with the grey filled pixels on a MR, with the dashed (coloured) lines representing PSFs forming the different configurations. The right panel contains all the possible configurations with their corresponding pixels contributing.

3.3.4.3 Challenges and issues of localisation scheme

Two main complications based on our proposed concept to decide on positions of molecules should be addressed, both corresponding to an identifiability problem based on the symmetries of the SiMPa functions. The first is related to cases where the PSFs of a configuration contribute the same number of pixels to a MR, where the second appears when the predetermined realisations produce either a symmetric diffraction with respect to the SiMPa's origin or the number of PSF pixels contributing to the MR is insufficient.

In part (a) of Figure 3.3.8, we present an example of the identifiability issue according to a configuration consisting of PSFs that have the same contribution to the MR. The pattern consists of two 'On' pixels, one on each corner of the MR, with the associated configuration represented by two individual PSFs each one having a single pixel contribution to the MR. Two separate set of molecules are outlined on the PSFs, denoted with the blue and red coloured points, along with their associated SiMPa functions based on the rotation scheme from Figure 3.3.5 when considering directions only on the 1st quadrant. The blue set consists of a molecule at $r = \theta = 0$ (PSF₁) and a second at $r = \frac{r_{\max, \theta=5\pi/4}}{2}$ with $\theta = 5\pi/4$ (PSF₁), while the red set of a molecule at $r = \frac{r_{\max, \theta=\pi/4}}{2}$ with $\theta = \pi/4$ (PSF₁) and a second at $r = \theta = 0$ (PSF₂). The identification issue arises due to the fact that both set of molecules have the same likelihood according to the MR, preventing our proposed scheme to distinguish between the appropriate realisations associated with the configuration, hence potentially leading to inaccurate localisation. Focusing on the mean of the MR, μ_{MR} , in Equation (3.3.10),

$$\begin{aligned}
 \mu_{\text{MR}} &= I\left(dN_{\{\partial x_{ij}=-1\}} + \sum_{v=1}^{\text{N}_{\text{PSF}}} (n_{\text{PSF}_v} - \alpha_{\text{O}, n_{\text{PSF}_v}}(r_v, \theta_v, c))\right) \\
 &= I\left(d \times 7 + (1 - c\alpha_4(r = 0, \theta = 0)) + (1 - c\alpha_8\left(r = \frac{r_{\max, \theta=5\pi/4}}{2}, \theta = \frac{5\pi}{4}\right))\right) \\
 &= I\left(d \times 7 + (1 - c\alpha_4(r = 0, \theta = 0)) + (1 - c\alpha_4\left(r = \frac{r_{\max, \theta=\pi/4}}{2}, \theta = \frac{\pi}{4}\right))\right)
 \end{aligned}$$

$$\begin{aligned}
&= I\left(d \times 7 + \left(1 - c\alpha_4\left(r = \frac{r_{\max, \theta = \pi/4}}{2}, \theta = \frac{\pi}{4}\right) + (1 - c\alpha_8(r = 0, \theta = 0))\right)\right) \\
&= \mu_{\text{MR}},
\end{aligned}$$

where the second line of the equation is evaluated at the realisations without the use of rotation, with their appropriate equivalent rotated ones shown on the third line. The equality of μ_{MR} corresponds to identical contribution to the total intensity of the MR by either sets of molecules, which within the probabilistic scheme translates into identical probabilities.




Figure 3.3.8: Different examples of identification issues regarding localisation of molecules using the SiMPa functions. (a) Two set of molecules, denoted with the blue and red colours respectively, contributing exactly the same total intensity to the MR with respect to the SiMPa functions. (b) Two separate molecules, blue and red coloured, contributing the same intensity to the MR according to the SiMPa functions.

In part (b) of Figure 3.3.8, we present a similar case of the identifiability issue formed when a predetermined realisation produces a symmetric diffraction based on

the SiMPa functions. The pattern consists of one 'On' pixel corresponding to a configuration of an individual PSF, with two different molecules placed at $r = \frac{r_{\max, \theta = \pi/2}}{2}$ with $\theta = \pi/2$ (blue coloured) and $r = \frac{r_{\max, \theta = 3\pi/2}}{2}$ with $\theta = 3\pi/2$ (red coloured) respectively. Additionally, their corresponding SiMPa functions are outlined based on the rotation scheme in Figure 3.3.5. Similarly to part (a), the intensity contribution to the MR of the either PSFs based on the SiMPa functions is identical, hence same probabilities on both realisations will be assigned based on our probabilistic scheme, with the μ_{MR} ,

$$\begin{aligned}
 \mu_{\text{MR}} &= I\left(dN_{\{\partial x_{ij} = -1\}} + \sum_{v=1}^{N_{\text{PSF}}} (n_{\text{PSF}_v} - \alpha_{\text{O}, n_{\text{PSF}_v}}(r_v, \theta_v, c))\right) \\
 &= I\left(d \times 8 + (1 - c\alpha_8(r = \frac{r_{\max, \theta = \pi/2}}{2}, \theta = \frac{\pi}{2}))\right) \\
 &= I\left(d \times 8 + (1 - c\alpha_1(r = \frac{r_{\max, \theta = 0}}{2}, \theta = 0))\right) \\
 &= I\left(d \times 8 + (1 - c\alpha_8(r = \frac{r_{\max, \theta = 0}}{2}, \theta = 0))\right) = \mu_{\text{MR}}.
 \end{aligned}$$

Such an identification issue can be addressed in cases where either the direction θ , with $\theta \in [0, \pi/2)$ (1st quadrant), produces a symmetric diffraction with respect to the origin of the SiMPa functions, as discussed, or when any direction θ from the 1st quadrant is accompanied with insufficient number of PSF pixels. Regarding the latter, consider the scenario in part (b) of Figure 3.3.8, with a single molecule located at $r = \frac{r_{\max, \theta = \pi/8}}{2}$ with $\theta = \pi/8$ and a second one at $r = \frac{r_{\max, \theta = 3\pi/8}}{2}$ with $\theta = 3\pi/8$. These realisations are indistinguishable when the MR captures only a single pixel of the PSF, in terms of both realisations having the same likelihood, however this is not the case if the MR falls onto two pixels of the PSF for instance.

These identifiability issues arise due to the symmetry/periodicity of the SiMPa functions, and can prevent accurate localisation in some of the cases where different configurations associated with realisations yield identical likelihoods. However, given the probabilistic approach, such symmetries are likely to be broken for a different

MR and thus the posterior probabilities of the correct positions should be higher.

3.3.5 Bayesian update Pseudo-algorithms

MCMC Individual frame inference scheme - Pseudo-algorithm

1. *Step $q=0$ (Iteration 0):*

Initialisation of parameters sampling from their prior distributions in Equations 3.3.13. This results in obtaining, $I^{(0)}$, $c^{(0)}$, $\tau_b^{(0)}$, $\{\beta_0, \beta_f\}^{(0)}$, $\mathbf{x}^{(0)}$ and set initial r 's and θ 's equal to 0, hence the collection $\{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}^{(0)}$ containing all the localisations based on the pattern-configuration probabilistic scheme for every pixel is set to $\{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}^{(0)} = \{\mathbf{Conf}, \mathbf{0}, \mathbf{0}\}^{(0)}$.

2. *Step $q = 1, \dots, IT$ (Iteration 1 up to IT):*

- (i) Update the parameters $I^{(q)}$, $\tau_b^{(q)}$ (Gibbs steps), $c^{(q)}$, $\{\beta_0, \beta_f\}^{(q)}$ (Metropolis-Hastings steps) from their corresponding full conditional distributions using the moving regions $\mathbf{MR} = \{\text{MR}_{ij}, \forall i = 1, \dots, n \text{ and } j = 1, \dots, m\}$.

$$I^{(q)} \sim p\left(I \mid \mathbf{x}^{(q-1)}, \mathbf{MR}, \tau_b^{(q-1)}, c^{(q-1)}, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}^{(q-1)}\right)$$

from Equation(3.3.16)

$$\tau_b^{(q)} \sim p\left(\tau_b \mid \mathbf{x}^{(q-1)}, \mathbf{MR}, I^{(q)}, c^{(q-1)}, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}^{(q-1)}\right)$$

from Equation(3.3.17)

$$c^{(q)} \sim p\left(c \mid \mathbf{x}^{(q-1)}, \mathbf{MR}, I^{(q)}, \tau_b^{(q)}, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}^{(q-1)}\right)$$

from Equation(3.3.18)

$$\{\beta_0, \beta_f\}^{(q)} \sim p\left(\beta_0, \beta_f \mid \mathbf{x}^{(q-1)}\right)$$

from Equation(3.3.21)

- (ii) Update the field $\mathbf{x}^{(q)}$ pixel-wise from the full conditional in Equation (3.3.19), using the original data $\mathbf{y} = \{y_{ij}, \forall i = 1, \dots, n \text{ and } j = 1, \dots, m\}$,

$$x_{ij}^{(q)} \sim p\left(x_{ij} \mid \mathbf{y}, \mathbf{x}^{(q-1)}, I^{(q)}, \tau_b^{(q)}, c^{(q)}, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}^{(q-1)}, \{\beta_0, \beta_f\}^{(q)}\right)$$

from Equation(3.3.19)

- (iii) Choose between *individual frame localisation scheme* and *individual frame counting scheme*.

Individual frame localisation scheme

For inference on configurations, that is pixels including the molecules based on the pattern-configuration concept, as well their respective realisation, apply the probabilistic localisation scheme based on SiMPa functions, based on the most up to date field $\mathbf{x}^{(q)}$.

For all $\mathbf{MR} = \{\mathbf{MR}_{ij}, \forall i = 1, \dots, n \text{ and } j = 1, \dots, m\}$, where each individual \mathbf{MR}_{ij} has a finite number of configurations that can be formed based on its pattern in $\mathbf{x}^{(q)}$, apply the proposed probabilistic scheme,

$$\begin{aligned} P_{ij(1)}^{(q)} &= p\left(\{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{ij(1)} \mid \mathbf{MR}_{ij}, \mathbf{x}^{(q)}, I^{(q)}, c^{(q)}, \tau_b^{(q)}\right) \\ &\vdots \\ P_{ij(u)}^{(q)} &= p\left(\{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{ij(u)} \mid \mathbf{MR}_{ij}, \mathbf{x}^{(q)}, I^{(q)}, c^{(q)}, \tau_b^{(q)}\right) \end{aligned}$$

with localisation probabilities to be obtained after normalising the densities $P_{ij(1)}^{(q)}, \dots, P_{ij(u)}^{(q)}$, that is

$$P_{ij(1)}^{(q)} \Big/ \sum_{h=1}^u P_{ij(h)}^{(q)}, \dots, P_{ij(u)}^{(q)} \Big/ \sum_{h=1}^u P_{ij(h)}^{(q)}$$

where each one of the $P_{ij(\bullet)}^{(q)}$ is obtained from Equation (3.3.23).

Individual frame counting scheme

For inference on configurations, that is pixels including the molecules based on the pattern-configuration concept apply the probabilistic counting scheme based on SiMPa functions, based on the most up to date field $\mathbf{x}^{(q)}$.

$$\begin{aligned} C_{ij(1)}^{(q)} &= p\left(\{\text{Conf}, \mathbf{r} = \mathbf{0}, \boldsymbol{\theta} = \mathbf{0}\}_{ij(1)} \mid \text{MR}_{ij}, \mathbf{x}^{(q)}, I^{(q)}, c^{(q)}, \tau_b^{(q)}\right) \\ &\vdots \\ C_{ij(w)}^{(q)} &= p\left(\{\text{Conf}, \mathbf{r} = \mathbf{0}, \boldsymbol{\theta} = \mathbf{0}\}_{ij(w)} \mid \text{MR}_{ij}, \mathbf{x}^{(q)}, I^{(q)}, c^{(q)}, \tau_b^{(q)}\right) \end{aligned}$$

with localisation probabilities to be obtained after normalising the densities $C_{ij(1)}^{(q)}, \dots, C_{ij(w)}^{(q)}$, that is

$$C_{ij(1)}^{(q)} / \sum_{h=1}^w C_{ij(h)}^{(q)}, \dots, C_{ij(w)}^{(q)} / \sum_{h=1}^w C_{ij(h)}^{(q)}$$

where each one of the $C_{ij(\bullet)}^{(q)}$ is obtained from Equation (3.3.24).

3. At the end of the MCMC, calculate posterior probabilities for the configurations by applying (Drawn Configurations)/ (Potential Configurations) to obtain pixels including the fluorophores.

If localisation is performed, calculate additional posterior probabilities of realisations within the configurations by applying (Drawn realisations)/ (Potential realisations).

3.4 Implementation

In this section an implementation of the individual frame inference scheme based on the SiMPa functions is presented, with the main purpose of describing the outcome of our novel pattern-configuration-realisation probabilistic procedure. This is based on a toy example with more complex images analysed in the application on synthetic data in Chapter 5. We use the individual frame localisation algorithm from Section 3.3.5 in order to present the output on drawing configurations associated with realisations, i.e. inference on specific positions of molecules within pixels with respect to a discretisation of r and θ . The individual frame counting algorithm is a simplification where every molecule is assumed to be at the origin of the SiMPa functions, corresponding to the center of the pixel with $r = \theta = 0$. Using the algorithm in Section 3.2.5, we generate a single frame of size $n \times m = 40 \times 40$ based on the SiMPa functions, which is displayed in Figure 3.4.1.

Regarding the parameters chosen to generate the frame, the single event intensity was set to $I = 8$ with the background proportion being $d = 0.3$, hence the background intensity $I_0 = d \times I = 2.4$. Independent error from a zero mean Gaussian distribution was added to the measurements, that is $N(0, \tau_b)$, where the background precision was selected to be $\tau_b = 10$. The choice on the pixels containing the molecules within the frame was arbitrary, with their positions in terms of distances r on directions θ , to have been chosen as continuous values in the 2-dimensional space. Under the power of diffraction parameter $c = 0.4$, the PSFs of the molecules were obtained using the SiMPa functions.

Within the frame, a variety of overlapping scenarios can be identified, with PSFs having different number of pixels coinciding, along with a number of well isolated PSFs. On the middle left part, around [Row, Column] = (22, 5), two individual molecules have their PSFs overlapping on four pixels, whereas on the top right part, around [Row, Column] = (33, 26), on two. Also, on the bottom left area, at [Row, Column] = (7, 13), there is a complete overlap of the PSFs of two active molecules

which lie onto the same pixel on different positions.




Figure 3.4.1: Simulated individual frame using the SiMPa functions. The light blue crosses represent the true positions of the molecules with their Point Spread Functions, PSFs, to be shown in the surrounding 3×3 pixel regions.

We apply the individual frame localisation scheme for $IT = 25.000$ iterations after a burn-in period of 10.000 iterations, as presented in Section 3.3.5. Regarding the output of the scheme, the posterior distributions of the parameters are obtained via the appropriate Gibbs and Metropolis-Hastings steps. The posterior distributions, shown by the histograms and the blue solid density lines, along with the prior distributions, denoted with the red solid lines, are presented in Figure 3.4.2, where the selected prior distributions are,

$$\begin{aligned} \pi(I | \mu_I, \tau_I) &\sim \text{N}(I | \mu_I = 8, \tau_I = 2) \\ \pi(c | a_c, b_c) &\sim \text{RescaledBeta}(c | a_c = 2, b_c = 2, tr_c = 0.7071) \\ \pi(\tau_b | a_\tau, b_\tau) &\sim \text{Gamma}(\tau_b | a_\tau = 2, b_\tau = 0.25) \\ \pi(b_0 | m_{b_0}, \tau_{b_0}) &\sim \text{N}(b_0 | m_{b_0} = 0, \tau_{b_0} = 9) \end{aligned}$$

$$\pi(b_f | m_{b_f}, \tau_{b_f}) \sim N(b_f | m_{b_f} = 0, \tau_{b_f} = 9). \quad (3.4.1)$$

and the background intensity proportion $d = 0.3$. Regarding the choice of I , we consider an informative prior distribution centred at the true value, based on our argument that it can be measured from the data during an imaging experiment. For both the background error precision τ_b and power of diffraction parameter c , the prior distributions correspond to non-informative choices whereas β_0 and β_f are Gaussian distributions centred at zero, while values between -2 and 2 generally serve the purpose.




Figure 3.4.2: Prior and posterior distributions for the parameters I , τ_b , c , β_0 and β_f in the simulated individual frame using the SiMPa functions. The red solid lines denote the prior distributions for each parameter while the histograms associated with the blue solid lines denote the posterior distributions of each parameter respectively.

In order to perform the localisation scheme, we select a discretisation of r and θ to be all the directions $\theta = k\frac{\pi}{4}$, $k=0,1,\dots,8$ under the distances $r = 0$ and $r = \frac{r_{\max,\theta}}{2}$, corresponding to the origin of the SiMPa functions and half the distance from the origin on direction θ respectively.

During the MCMC, the update of the field follows the update of the parameters, resulting in a binary mask representing the states of every pixel, that is 'On' or 'Off'. For every moving region, the full conditional distributions of the complete localisation sets are evaluated, based on the patterns formed within the updated field and the corresponding discretisation of r and θ . A complete localisation set consist of a configuration associated with corresponding realisations. As thoroughly described in Section 3.3.3, the full conditional distributions are introduced within a probabilistic scheme in order to draw localisations of the molecules, both in terms of identifying the pixels containing them and their specific positions within in terms of r and θ .

More specifically, each configuration corresponds to a combination of PSFs based on the respective pattern, regardless of the realisations the configuration is associated with. Consequently, each configuration is an identification mechanism of the pixels in which the molecules lie onto, probably more easily spotted when considering the simplification with $r = \theta = 0$. Taking into consideration every configuration for a specific pattern of a moving region, we are able to obtain the set of all potential pixels in which the molecules can lie onto. We call this the potential positions set, where the drawn positions declare the pixels which include the molecules and are selected using their multinomial full conditional distribution. For the individual frame localisation scheme, the drawn positions are additionally accompanied with a realisation, hence we acquire what we call the drawn localisations. The potential localisations correspond to all the different combinations of r and θ in the discretisation a molecule can lie within a pixel.

All these individual quantities are key in order to conduct inference on the positions

of the molecules within a frame. Different counters are established responsible for recording the draws and potential positions on every iteration. Regarding the configuration draws, corresponding to pixels containing the molecules, three individual counters of the size of the frame are considered. In the first one, named 'Visits', we store the potential positions of every moving region, with the second one collecting the corresponding drawn positions, named as 'Positions'. Since we store single draws within 'Positions', as it corresponds to pixels containing a single active molecule, we consider a third counter named 'Double events' accounting for two active molecules in a single pixel. In the 'Double events' counter, we record the drawn positions constituting of two molecules having their PSFs completely overlapping. In order to store the drawn localisations, we establish a counter we name 'Localisations', where we treat each pixel as a super-pixel consisting of its regular division according to the discretisation scheme of r and θ . Each one of the divisions corresponds to a specific realisation, therefore since $\theta = k\frac{\pi}{4}$, $k=0,1,\dots,8$ under the distances $r = 0$ and $r = \frac{r_{\max,\theta}}{2}$ we have 3×3 super-pixels.

At the end of the MCMC, we construct our posterior probability maps based on the stored counts. The ratio of 'Positions' and 'Visits' corresponds to a probability heat-map declaring pixels containing a molecule, accompanied with a probability heat-map regarding pixels containing two molecules when dividing the 'Double events' with the 'Positions' counter. In general, when a pixel is identified as a pixel containing a molecule, we have an additional mechanism to determine the existence of two molecules within the pixel. In a similar manner, we obtain a posterior probability map denoting the localisation of the molecules within the identified pixels, by taking the ratio of the 'Localisation' and the 'Positions' counter.

In Figure 3.4.3, we present the posterior probability maps for the simulated individual frame, with the unchanged and filtered maps on the right and left columns of the figure respectively. The light blue bars, on the right bottom corners of each heat-map, denote the altering probability levels, with the bins representing the probability




Figure 3.4.3: Posterior probability maps for the simulated individual frame based on SiMPa functions. On the left column of the figure the original probability maps are shown whereas on the right column the ones filtered by the average field. The top row contains the probabilities of pixels containing a single molecule with the middle row accounting for double molecules within the pixels. The bottom row consists of the super-pixels probabilities for molecules on the sub-pixels, each one representing a realisation.

mass in the respective regions. Regarding the filtering, we consider two different ways to clear the heat-maps both related to the average field. The average field is the ratio of the sum of fields for every iteration and the number of iterations. We decide to multiply our probability maps with either the average field, which is the case in our implementation, or a binary mask obtained by a threshold on the average field. The main reason to consider such filters relies on the fact that a molecule cannot lie outside the identified 'On' pixels.

Chapter 4

Stack of frames analysis

4.1 Introduction

Active molecules, or fluorophores, captured on frames within a SMLM procedure are not directly observable, however their intensities within their emission regions can be measured. The intensities correspond to the observed signal of the frames as stored by the recorder of the imaging system. The entire sequence of frames forms the underlying structure under investigation, exhibiting both spatial and temporal dependencies. The former corresponds to the diffraction of captured active molecules within frames, where the latter to the properties and dynamics of fluorophores across the time dependent frames. As we introduced in Chapter 3, we model the behaviour of an active molecule's diffraction by using the SiMPa functions, which take into consideration the molecule's location within a pixel. A complete Bayesian modelling approach was provided, which allows for inference on individual frames based on our novel pattern-configuration-realisation concept.

Besides modelling the spatial association of fluorophores, a corresponding mechanism is needed to account for progressing our SiMPa functions based model over time. In this case, time is translated into the total number of frames within the stack. Natu-

rally, even when the most developed cameras are used in order to capture images of the specimen, the possibility of either capturing entire events of active fluorophores or having complete blinking cycles on a single frame is not very high. This is directly related with the molecule's properties and dynamics, the capabilities of the imaging technique and generally the complete imaging procedure [Antolovic et al., 2017], which has been discussed in more detail in Chapter 2. Practically, same or reappearing events can be captured in more than one consecutive or within the stack of frames, which needs to be incorporated throughout the chosen modelling. From a quantitative point of view, having a mechanism capable of identifying fluorophores interchanging between excited and dark states before being photo-bleached is of extreme importance. Counting the number of molecules on a stack of frames, hence in the specimen under investigation, can contribute to gaining substantial information about intermolecular interactions and behaviours.

In this chapter, we present a state space model based on SiMPa functions, which allows for coherent propagation of uncertainty both in space and time. We introduce a probabilistic frame inference scheme based on SiMPa functions, similar to what we presented in Chapter 3, within a Markov switching model driving the stochastic behaviour of the pixel states over time. In Section 4.2, we start with a brief introduction of the general Markov switching model along with its associated properties to be discussed in Section 4.2.1, mainly based on Frühwirth-Schnatter [2006]. In Section 4.2.2 we formulate the Markov switching model based on SiMPa functions, followed with a fully detailed scheme on Bayesian inference for both parameters and states in Section 4.2.3. In Section 4.3 a pseudo-algorithm of a sampling scheme for generating a stack of frames is introduced. A transition matrix accounts for alteration of a molecule's state between frames whereas the SiMPa functions model the spread of the intensities of any active molecules within the frames. The model implementation is illustrated in Section 4.4, using a synthetic data set generated by the sampling scheme, with a detailed discussion on the output and localisation.

4.2 General form of state-space model

4.2.1 The Markov Switching model

In this section, the basic Finite Markov Switching model is presented when time series data is of main interest, as described in [Frühwirth-Schnatter \[2006\]](#). We adopt this term, however a variety of names exists in the literature denoting the exact same model, such as Markov Mixture or hidden Markov models, with applications on different areas for instance in econometrics [[Pagan and Schwert, 1990](#)], engineering [[Rabiner, 1989](#)], neuroimaging [[Smith et al., 2010b](#)], as well as super resolution reconstructing [[Humblot and Mohammad-Djafari, 2006](#)] and single cell imaging [[Hey et al., 2015](#)]. Some basic properties and conditions will be specified, which are essential before introducing the Markov Switching model based on SiMPa functions in [Section 4.2.2](#).

Following [Frühwirth-Schnatter \[2006\]](#), let a vector $\mathbf{y} = (y_1, \dots, y_t, \dots, y_T)$ denote a set of T time series observations, arising from a stochastic process Y_t , with $t = 1, \dots, T$. For the purposes of our work, we assume that these observations are measurements on a continuous scale, without this being a necessary constraint. We allow the stochastic process Y_t , which is observable, to depend on a hidden discrete stochastic process S_t , $t = 1, \dots, T$, affecting its behaviour via a specified relationship. Let the vector $\mathbf{S} = (S_0, \dots, S_t, \dots, S_T)$ denote a realization of the discrete unobserved process S_t for $t = 0, \dots, T$, taking values on a finite space of K states, i.e. $S_t = k$ with $k = 1, \dots, K$. The two following rather strong assumptions regarding S_t and Y_t need to be obeyed in order to define the basic Markov switching model, entities which will be later relaxed to obtain a sensible Markov switching model based on the SiMPa functions.

Assumptions

- (A) The hidden stochastic process S_t needs to be an irreducible aperiodic Markov chain starting from an ergodic distribution.
- (B) The random variables $\mathbf{Y} = (Y_1, \dots, Y_t, \dots, Y_T)$ are conditionally independent

when knowing the states $\mathbf{S} = (S_0, \dots, S_t, \dots, S_T)$, arising from the distribution,

$$Y_t | S_t = k, \boldsymbol{\phi} \sim f(Y_t | \boldsymbol{\phi}_{S_t=k}),$$

where $f(Y_t | \cdot)$ denotes a specific parametric family with $\boldsymbol{\phi}$ being the corresponding parameters.

With respect to conditions **(A)** and **(B)**, the basic Markov switching model is defined via the conditional distribution of Y_t given S_t as,

$$Y_t | S_t, \boldsymbol{\phi} \sim f(Y_t | \boldsymbol{\phi}_{S_t}). \quad (4.2.1)$$

Regarding condition **(A)**, we present a brief discussion on the specific properties of a Markov chain. More detail can be found in [Frühwirth-Schnatter \[2006\]](#) and [Karlin and Taylor \[1981\]](#). Firstly, we allow the stochasticity of S_t to be entirely described by a transition matrix $\boldsymbol{\xi}$, with ξ_{lk} denoting the probability of transitioning from state l to state k between two discrete time points say $t - 1$ and t , that is $\xi_{lk} = \Pr(S_t = k | S_{t-1} = l)$, $\forall l, k \in \{1, \dots, K\}$ and

$$\boldsymbol{\xi} = \begin{pmatrix} \xi_{11} & \dots & \xi_{1K} \\ \vdots & \ddots & \vdots \\ \xi_{K1} & \dots & \xi_{KK} \end{pmatrix}, \quad (4.2.2)$$

where $\forall lk \in \{1, \dots, K\}$. Every row of the transition matrix $\boldsymbol{\xi}$ obeys $\xi_{lk} \geq 0$ and $\sum_{k=1}^K \xi_{lk} = 1$, with the latter corresponding to the sum of the transition probabilities of a specific state to every possible one to be equal to one.

A probability distribution $\boldsymbol{\eta} = (\eta_1, \dots, \eta_k, \dots, \eta_K)$ satisfying the invariance property is called the invariant distribution of S_t . This means that if S_{t-1} is drawn from such an invariant distribution $\boldsymbol{\eta}$ the states of S_t will be also drawn from $\boldsymbol{\eta}$, translating into $\Pr(S_t = k | \boldsymbol{\xi}) = \eta_k$ for $k \in 1, \dots, K$, with $\boldsymbol{\eta}$ existing for every Markov chain

however not uniquely defined. A Markov chain with a transition matrix $\boldsymbol{\xi}$ which achieves uniqueness of $\boldsymbol{\eta}$ and additionally convergence of S_t to $\boldsymbol{\eta}$ regardless of the initial state S_0 , is called ergodic with $\boldsymbol{\eta}$ being its ergodic distribution. If a Markov chain is aperiodic then the ergodicity property is satisfied, with aperiodicity corresponding to lack of periodicity between different states. This translates into having positive transition probabilities between the same states of a Markov chain, that is the diagonal entries of the transition matrix hence $\xi_{kk} > 0, \forall k \in 1, \dots, K$. Lastly, a Markov chain is irreducible if every state of the chain will be reached for a S_t which starts from an arbitrary state $k \in \{1, \dots, K\}$, hence $\underbrace{(\boldsymbol{\xi} \times \boldsymbol{\xi} \times \dots \times \boldsymbol{\xi})}_{h}{}_{lk} > 0$ for some step $h \geq 1$.

The general form of a Markov Switching model is defined by obtaining the joint distribution of the random variables $\mathbf{Y} = (Y_1, \dots, Y_t, \dots, Y_T)$ and their corresponding hidden states $\mathbf{S} = (S_0, \dots, S_t, \dots, S_T)$, modelled by a Markov process. Let $\mathbf{y}_{1:T} = (y_1, \dots, y_t, \dots, y_T)$ to be a realization of \mathbf{Y} associated with their states $\mathbf{S}_{1:T} = (S_1, \dots, S_t, \dots, S_T)$, then the joint distribution can be written as,

$$\begin{aligned}
p(\mathbf{S}_{1:T}, \mathbf{y}_{1:T} | \boldsymbol{\vartheta}) &\propto f(y_T | \mathbf{y}_{1:T-1}, \mathbf{S}_{1:T}, \boldsymbol{\vartheta}) p(S_T | \mathbf{S}_{1:T-1}, \mathbf{y}_{1:T-1}, \boldsymbol{\vartheta}) \\
&\quad \times p(\mathbf{y}_{1:T-1} | \mathbf{S}_{1:T-1}, \boldsymbol{\vartheta}) p(\mathbf{S}_{1:T-1} | \boldsymbol{\vartheta}) \\
&\propto f(y_T | \mathbf{y}_{1:T-1}, \mathbf{S}_{1:T}, \boldsymbol{\vartheta}) p(S_T | \mathbf{S}_{1:T-1}, \mathbf{y}_{1:T-1}, \boldsymbol{\vartheta}) p(\mathbf{S}_{1:T-1}, \mathbf{y}_{1:T-1} | \boldsymbol{\vartheta}) \\
&\propto f(y_T | \mathbf{y}_{1:T-1}, \mathbf{S}_{1:T}, \boldsymbol{\vartheta}) p(S_T | \mathbf{S}_{1:T-1}, \mathbf{y}_{1:T-1}, \boldsymbol{\vartheta}) f(y_{T-1} | \mathbf{y}_{1:T-2}, \mathbf{S}_{1:T-1}, \boldsymbol{\vartheta}) \\
&\quad \times p(S_{T-1} | \mathbf{S}_{1:T-2}, \mathbf{y}_{1:T-2}, \boldsymbol{\vartheta}) p(\mathbf{y}_{1:T-2} | \mathbf{S}_{1:T-2}, \boldsymbol{\vartheta}) p(\mathbf{S}_{1:T-2} | \boldsymbol{\vartheta}) \\
&\propto f(y_T | \mathbf{y}_{1:T-1}, \mathbf{S}_{1:T}, \boldsymbol{\vartheta}) p(S_T | \mathbf{S}_{1:T-1}, \mathbf{y}_{1:T-1}, \boldsymbol{\vartheta}) f(y_{T-1} | \mathbf{y}_{1:T-2}, \mathbf{S}_{1:T-1}, \boldsymbol{\vartheta}) \\
&\quad \times p(S_{T-1} | \mathbf{S}_{1:T-2}, \mathbf{y}_{1:T-2}, \boldsymbol{\vartheta}) p(\mathbf{S}_{1:T-2}, \mathbf{y}_{1:T-2} | \boldsymbol{\vartheta}) \\
&\propto \dots \\
&\propto \prod_{t=1}^T f(y_t | \mathbf{y}_{1:t-1}, \mathbf{S}_{1:t}, \boldsymbol{\vartheta}) p(S_t | \mathbf{S}_{1:t-1}, \mathbf{y}_{1:t-1}, \boldsymbol{\vartheta}) p(S_0 | \boldsymbol{\vartheta}), \tag{4.2.3}
\end{aligned}$$

where $\mathbf{y}_{1:t-1} = (y_1, \dots, y_{t-1})$ and accordingly $\mathbf{S}_{1:t-1} = (S_1, \dots, S_{t-1})$, $\boldsymbol{\vartheta} = (\boldsymbol{\phi}, \boldsymbol{\xi})$ with $\boldsymbol{\phi} = (\boldsymbol{\phi}_1, \dots, \boldsymbol{\phi}_K)$ the parameters for the different states K and $\boldsymbol{\xi}$ the transition matrix as in Equation (4.2.2). The density $p(S_0|\boldsymbol{\vartheta})$ denotes the initial distribution of the states in $\{1, \dots, K\}$. The other two densities which appear by writing down the joint distribution as this product, are commonly known as *one step ahead predictive density* of y_t conditional on the past observations $\mathbf{y}_{1:t-1} = (y_1, \dots, y_{t-1})$ and the states $\mathbf{S}_{1:t-1} = (S_1, \dots, S_{t-1})$ and *conditional distribution* of S_t when given past observations $\mathbf{y}_{1:t-1} = (y_1, \dots, y_{t-1})$ and states $\mathbf{S}_{1:t-1} = (S_1, \dots, S_{t-1})$. The former, denoted by $f(y_t|\mathbf{y}_{1:t-1}, \mathbf{S}_{1:t}, \boldsymbol{\vartheta})$, basically allows dependency between the current observation and the previous ones when the states are known, while the latter, denoted by $p(S_t|\mathbf{S}_{1:t-1}, \mathbf{y}_{1:t-1}, \boldsymbol{\vartheta})$, gives the conditional probability distribution of the current state when past observations and states are available.

According to the choice of modelling, different assumptions can be made about the dependency between the observations themselves or/and the states modelled by $f(y_t|\mathbf{y}_{1:t-1}, \mathbf{S}_{1:t}, \boldsymbol{\vartheta})$ as well the relationship of the states over time $p(S_t|\mathbf{S}_{1:t-1}, \mathbf{y}_{1:t-1}, \boldsymbol{\vartheta})$. Regarding the former, two of the most common assumptions are either to assume only dependency on the current state S_t , e.g. Markov switching autoregressive model [Hamilton, 1989], or on a finite number of previous states, say S_{t-1}, \dots, S_{t-p} , e.g. Markov switching ARMA models [Billio et al., 1999].

4.2.2 Markov Switching model based on SiMPa functions

In this section, we describe the general form of the Markov switching model based on the SiMPa functions. As discussed in detail in Chapter 2, a SMLM frame consists of intensities which represent the visible outcome of a single image. These intensities provide measurements of underlying fluorophores stochastically emitting light when exposed on specific wavelengths of light. After repetition of this process over a number of times different, or the same molecules undergo similar proceedings. When the entire experiment is complete a large number of unprocessed frames form a stack of frames providing information about the underlying structure of the specimen under

investigation.

We consider this stack of frames as a sequence of lattices consisting of intensities over $t = 1, \dots, T$ time points. Every lattice contains a specific set of fluorophores either on excited or dark states, with the active fluorophores on each lattice to potentially arise from the blinking procedure. Since a photo-bleached fluorophore cannot be directly observed, we consider a two-state procedure of 'On' and 'Off', denoting excited and ground states respectively. The overall state of each lattice is dependent over the time points $t = 1, \dots, T$ due to the blinking process with fluorophores allowed to stochastically alter or retain their state. A natural way to describe this imaging procedure is by a dynamic system which evolves over time, where time can be described as the total number of frames the entire stack consists of.

Following similar notation to Chapter 3, let \mathbf{y}_t to be a lattice of observed intensities at time t with $\mathbf{y}_t = \{y_{t,ij}\}$, $i = 1, \dots, n$ and $j = 1, \dots, m$, where n and m the number of rows and columns of the lattice respectively, and $y_{t,ij}$ to be the observed intensity of pixel $\{i, j\}$ at time t . The sequence of observations $\mathbf{y}_{1:T} = (\mathbf{y}_1, \dots, \mathbf{y}_t, \dots, \mathbf{y}_T)$ can be expressed as a time series over $t = 1, \dots, T$ time points, generated by a stochastic process Y_t . Regarding observations from a SMLM experiment, $\mathbf{y}_{1:T}$ can be considered as a realization of Y_t containing the observed intensities for a stack of frames consisting of T individual frames, with \mathbf{y}_t being the t^{th} frame. As a result, a discrete-time time series seems appropriate where t can take values in the discrete set $\{1, \dots, T\}$, hence $t \in \{1, \dots, T\}$. In our case, each one of the time points t represents a frame from the sequence of frames $1, \dots, T$ which are time related.

We allow a latent discrete valued variable $\mathbf{x}_{1:T}$ to be the hidden mechanism according to which the data are being generated, where $\mathbf{x}_{1:T} = (\mathbf{x}_1, \dots, \mathbf{x}_t, \dots, \mathbf{x}_T)$. We name each one of the $\mathbf{x}_t, t \in \{1, \dots, T\}$ a state configuration of the unobserved discrete valued process for frame t . We define a state configuration to be the collection of individual states for every pixel on frame t , hence $\mathbf{x}_t = \{x_{t,ij}\}$ for $i = 1, \dots, n$ and

$j = 1, \dots, m$ to be the number of rows and columns of frame t respectively. The state $x_{t,ij}$ of any pixel $\{i, j\}$ on frame t can take values in the finite discrete state set $\{1, \dots, K\}$ consisting of K states. Under our modelling, we consider $K = 2$ states with,

$$x_{t,ij} = \begin{cases} 1, & \text{if the pixel is declared to be in 'On' state} \\ -1, & \text{if the pixel is declared to be in 'Off' state} \end{cases},$$

where 1 and -1 represent the 'On', or active, and 'Off', or de-active, pixels respectively. We stress here that $x_{t,ij}$ corresponds to the state of the pixel and not the fluorophore itself. In order to have a more accessible notation throughout our modelling, we adopt a vectorised indexing scheme for every frame t , therefore,

$$\begin{aligned} \mathbf{y}_{1:T} &= (\mathbf{y}_1, \dots, \mathbf{y}_t, \dots, \mathbf{y}_T) \quad \text{with} \quad \mathbf{y}_t = \{y_{t,s}\}, \quad s = 1, \dots, N (= n \times m) \\ \mathbf{x}_{1:T} &= (\mathbf{x}_1, \dots, \mathbf{x}_t, \dots, \mathbf{x}_T) \quad \text{with} \quad \mathbf{x}_t = \{x_{t,s}\}, \quad s = 1, \dots, N, \end{aligned}$$

hence,

$$x_{t,s} = \begin{cases} 1, & \text{if pixel } s \text{ is in 'On' state} \\ -1, & \text{if pixel } s \text{ is in 'Off' state} \end{cases}, \quad (4.2.4)$$

Also, we denote with $\mathbf{y}_{1:t} = (\mathbf{y}_1, \dots, \mathbf{y}_t)$ the sequence of observed frames 1 to t for any $t \in \{1, \dots, T\}$ and similarly, $\mathbf{x}_{1:t} = (\mathbf{x}_1, \dots, \mathbf{x}_t)$ the unobserved configurations states of frames 1 to t .

We model the observed intensities $\mathbf{y}_t, \forall t = 1, \dots, T$, using a Gaussian distribution based on the configuration states $\mathbf{x}_t, \forall t = 1, \dots, T$, where the mean of each individual observation $\{y_{t,s}\}$ relies on the state of the discrete hidden underlying variable $\{x_{t,s}\}$,

$\forall s = 1, \dots, N$. The background error level is considered to be generated by the same distribution regardless of the state of the pixel. Therefore,

$$y_{t,s} = \mu_{x_{t,s}} + \epsilon_t, \quad \epsilon_t \sim \text{N}(0, \tau_b), \quad (4.2.5)$$

or equivalently, expressed in terms of the entire frame \mathbf{y}_t and the corresponding configuration states \mathbf{x}_t ,

$$\mathbf{y}_t = \boldsymbol{\mu}_{\mathbf{x}_t} + \epsilon_t, \quad \epsilon_t \sim \text{N}(0, \tau_b), \quad (4.2.6)$$

where the state dependent mean is represented by the SiMPa functions introduced in Chapter 3 and τ_b is the background error precision. As a reminder, the SiMPa functions $\alpha_k(r, \theta, c)$, $k = 0, \dots, 8$, describe the intensities diffraction in a predetermined neighbourhood (here of dimension 3×3 pixels) around the central pixel containing an active fluorophore, lying r units away from its center (origin) on direction θ , with c the parameter describing the power of diffraction. Our proposed use entails the single event intensity I , allowing the SiMPa functions to account for the intensity proportions of I in the predetermined neighbourhood (Section 3.2). Consequently, each active pixel ('On' state) can be considered as a member of a specified diffraction (PSF) where its assigned SiMPa value depends on the pixel the fluorophore lies into, along with the specified position declared by r and θ . We denote this with $\alpha_O(r, \theta, c)$. Thus,

$$\mu_{x_{t,s}} = \begin{cases} I (1 - \alpha_O(r, \theta, c)), & \text{if } x_{t,s} = 1 \\ I d, & \text{if } x_{t,s} = -1 \end{cases}, \quad (4.2.7)$$

where $I d = I_0$ denotes the proportion of the single event intensity I operating as a background baseline intensity, and the SiMPa functions $\alpha_k(r, \theta, c)$ as in Equation (3.2.1) in Section 3.2,

$$\alpha_k = \begin{cases} c \left(d_1 - r \cos \left[\frac{(k-1)\pi}{4} + \theta \right] g(R, l) \right), & k = 1, 3, 5, 7 \\ c \left(\left((rg(R, l))^2 + d_2^2 + 2d_2 s_k r \cos \left(\frac{(k+1)\pi}{4} - \theta \right) g(R, l) \right)^{1/2} \right), & k = 2, 4, 6, 8, \\ 0, & k = 0 \end{cases}$$

where $\alpha_0(r, \theta, c)$ in Equation 4.2.7 denoting any of the $\alpha_k(r, \theta, c)$ for any $k \in \{0, \dots, 8\}$, with respect to r and θ .

Regarding the latent discrete variable $\mathbf{x}_{1:T}$ representing the states of the pixels, we allow a dependency over consecutive time point progressing via a transition probability matrix $\boldsymbol{\xi}$. Additionally, the pixels themselves are spatially associated with each other since they feature intensities which spread from active fluorophores. This local spatial dependency can be conveniently modelled by a Markov Random Field (MRF), as presented in Section 3.3.1 and Definition 3.3.1. In that case, the state of a pixel depends on the states of the neighbours within a 3×3 pixel region, chosen according to the size of the PSF with respect to the SiMPa functions, leading to,

$$\mathbf{x}_t = \boldsymbol{\xi} \mathbf{x}_{t-1} \quad \text{or} \quad \mathbf{x}_t = \begin{pmatrix} \xi_{11} & \xi_{10} \\ \xi_{01} & \xi_{00} \end{pmatrix} \mathbf{x}_{t-1}, \quad (4.2.8)$$

with $\boldsymbol{\xi}$ to be the transition matrix consisting of $K = 2$ states ('On' and 'Off') and \mathbf{x}_{t-1} a MRF with pixel-wise dependency described by a neighbourhood structure via the conditional distributions, that is

$$\Pr(x_{t-1,s} | x_{t-1,q}, \forall q \neq s \in \{1, \dots, N\}) = \Pr(x_{t-1,s} | \partial x_{t-1,s}),$$

where $\partial \bullet$ denotes the neighborhood structure as in Equations (3.3.1) and (3.3.2). Practically, between two time points $t - 1$ to t , we allow a pixel-wise transition

between 'On' and 'Off' states while taking into consideration the states of the neighbourhood at time $t - 1$. The evolution of states depends on a transition scheme affected by the states of the neighbouring pixels, hence pixel-wise we have equivalently to (4.2.8),

$$x_{t,s} = \begin{pmatrix} \xi_{11} & \xi_{10} \\ \xi_{01} & \xi_{00} \end{pmatrix} x_{t-1,s}. \quad (4.2.9)$$

Equations (4.2.6) and (4.2.8) fully describe the general form of the state space model based on SiMPa functions, graphically represented in Figure 4.2.1. We note here that since we model the state of the pixels and not the states of the molecules themselves, a pixel cannot be 'Bleached' (see [Eggeling et al., 1998] for detail on photo-bleaching) but can cycle between 'On' and 'Off' states based on the diffraction of closely active molecules.




Figure 4.2.1: Graphical representation of the state space model based on SiMPa functions. Every lattice of intensities \mathbf{y}_t is a result of an unobserved state configuration \mathbf{x}_t , which alters its individual states over time according to the transition matrix ξ .

More specifically, similarly to Chapter 3, we model the MRF by using an autologistic model [Besag, 1974] with the joint distribution of $\mathbf{x}_{t-1} = \{x_{t-1,s}\}$, $s = 1, \dots, N$,

$$\pi(\mathbf{x}_{t-1}|\{\beta_0, \beta_f\}) = \frac{\exp\{\beta_0 V_0(\mathbf{x}_{t-1}) + \beta_f V_f(\mathbf{x}_{t-1})\}}{Z(\{\beta_0, \beta_f\})}, \quad (4.2.10)$$

where $V_0(\mathbf{x}_{t-1}) = \sum_{\forall s} x_{t-1,s}$ and $V_f(\mathbf{x}_{t-1}) = \sum_{q \in \partial x_{t-1,s}} x_{t-1,s} x_{t-1,q}$ with β_0 and β_f being parameters controlling the number of states and the strength of the neighboring interaction respectively. The quantity $Z(\{\beta_0, \beta_f\})$ denotes the normalizing constant. According to the Hammersley-Clifford theorem [Hammersley and Clifford, 1971], the MRF can be fully characterized from its full conditional distributions by,

$$\pi(x_{t-1,s}|\partial x_{t-1,s}, \{\beta_0, \beta_f\}) \propto \exp\left\{\beta_0 x_{t-1,s} + \beta_f \sum_{q \in \partial x_{t-1,s}} x_{t-1,s} x_{t-1,q}\right\}, \quad (4.2.11)$$

Similarly to Equation (3.3.22), we use the pseudo-likelihood estimation for the auto logistic field in order to reduce the complexity of having to calculate 2^N terms in the normalising constant [Besag, 1974],

$$\pi(\mathbf{x}_{t-1}|\{\beta_0, \beta_f\}) \approx \prod_{s \in \{1, \dots, N\}} \pi(x_{t-1,s}|\partial x_{t-1,s}, \{\beta_0, \beta_f\}). \quad (4.2.12)$$

In order to specify the Markov switching model based on SiMPa functions described by Equation (4.2.6) and (4.2.8), we need to make two assumptions associated with the dependencies over time. Firstly, we introduce some notation with $\boldsymbol{\vartheta} = (\boldsymbol{\phi}, \boldsymbol{\xi})$ to be denoting the entire collection of parameters, where $\boldsymbol{\phi} = (I, c, \tau_b, \boldsymbol{\beta}_{1:T}, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{1:T})$ is the respective collection of within frame parameters and $\boldsymbol{\xi}$ the transition matrix, as the between frame parameters. The vector $\boldsymbol{\beta}_{1:T}$ contains each frame-wise parameter set $\{\beta_{0_t}, \beta_{f_t}\}$ for every frame $t \in \{1, \dots, T\}$, while the collection $\{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{1:T}$ represents the complete localisation set for every frame $1, \dots, T$ with,

$$\boldsymbol{\beta}_{1:T} = (\boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_T) = (\{\beta_{0_1}, \beta_{f_1}\}, \dots, \{\beta_{0_T}, \beta_{f_T}\}),$$

$$\begin{aligned} \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{1:T} &= \left(\{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_1, \dots, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_T \right), \text{ and} \\ \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_t &= \left(\{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{t,s}, \forall s \in \{1, \dots, N\} \right), \end{aligned} \quad (4.2.13)$$

where a complete localisation frame set $\{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_t, \forall t = 1, \dots, T$, consists of all the drawn configurations associated with their corresponding realisations for every pixel $s \in \{1, \dots, N\}$, introduced on the probabilistic inference scheme introduced in Chapter 3 based on our novel pattern-configuration-realisation concept. We adjust the scheme for the Markov switching model based on SiMPa functions, later in Section 4.2.3. Therefore, for a single pixel,

$$\{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{t,s} = \left\{ \text{Conf}_{t,s}, \{r, \theta\}_{t,s(1)}, \dots, \{r, \theta\}_{t,s(v)} \right\}, \quad (4.2.14)$$

where $v = 1, \dots, N_{\text{PSF}}$ to be denoting the number of fluorophores that have been localised based on the configuration $\text{Conf}_{t,s}$ of pixel s . It should be noted that the notation used for the number of localisation of fluorophores is directly equivalent with the number of the PSF's, hence the use of N_{PSF} . For notation simplicity, we will remove the frame indexing when it is not confusing, hence use $\boldsymbol{\phi} = \left(I, c, \tau_b, \boldsymbol{\beta}_{1:T}, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{1:T} \right)$ instead of $\boldsymbol{\phi}_t = \left(I, c, \tau_b, \boldsymbol{\beta}_t, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_t \right)$. Using this notation we make the following assumptions.

Assumptions for MSM based on SiMPa functions

(A1) For fixed \mathbf{x}_t at time t , the intensity observations \mathbf{y}_t are conditionally independent and do not depend on previous intensity observations $\mathbf{y}_{1:t-1}$:

$$f(\mathbf{y}_t | \mathbf{y}_{1:t-1}, \mathbf{x}_{1:t}, \boldsymbol{\vartheta}) = f(\mathbf{y}_t | \mathbf{x}_t, \boldsymbol{\phi}) = \prod_{\forall s \in \{1, \dots, N\}} f(y_{t,s} | x_{t,s}, \boldsymbol{\phi}).$$

(A2) Given that \mathbf{x}_t is a Markov random field, the conditional distribution of \mathbf{x}_t

is independent of past observations $\mathbf{y}_{1:t-1}$ and only depends on the states of the previous states \mathbf{x}_{t-1} , transitioning according to the transition matrix $\boldsymbol{\xi}$ in (4.2.8), and their neighbouring states denoted by $\partial\mathbf{x}_{t-1}$, hence

$$\begin{aligned} p(\mathbf{x}_t | \mathbf{x}_{1:t-1}, \mathbf{y}_{1:t-1}, \boldsymbol{\vartheta}) &= p(\mathbf{x}_t | \mathbf{x}_{t-1}, \boldsymbol{\vartheta}) \\ &= p(\mathbf{x}_t | \mathbf{x}_{t-1}, \boldsymbol{\xi}) p(\mathbf{x}_{t-1} | \partial\mathbf{x}_{t-1}, \boldsymbol{\phi}) \\ &\stackrel{\text{pseudo}}{\approx} \prod_{\forall s \in \{1, \dots, N\}} p(x_{t,s} | x_{t-1,s}, \boldsymbol{\xi}) p(x_{t-1,s} | \partial x_{t-1,s}, \boldsymbol{\phi}). \end{aligned}$$

As the domain of the states $x_{t,s}$, $\forall t, s$ is discrete and can only take values in $K = \{-1, 1\}$, we can straightforward obtain the pixel-wise conditional probabilities of any state $k \in \{-1, 1\}$, that is 'Off' and 'On' state respectively, thus

$$\begin{aligned} \Pr(x_{t,s} = k | \mathbf{x}_{1:t-1}, \mathbf{y}_{1:t-1}, \boldsymbol{\vartheta}) &= \Pr(x_{t,s} = k | \mathbf{x}_{t-1}, \boldsymbol{\vartheta}) \\ &= \sum_{k_p \in \{-1, 1\}} \Pr(x_{t,s} = k | x_{t-1,s} = k_p, \boldsymbol{\xi}) \times \\ &\quad \times \Pr(x_{t-1,s} = k_p | \partial x_{t-1,s}, \boldsymbol{\phi}), \end{aligned}$$

where $\boldsymbol{\vartheta} = \{\boldsymbol{\phi}, \boldsymbol{\xi}\}$ and

$$\Pr(x_{t-1,s} = k_p | \partial x_{t-1,s}, \boldsymbol{\phi}) = \frac{\exp\{\beta_0 k_p + \beta_f \sum_{q \in \partial x_{t-1,s}} k_p x_{t-1,q}\}}{\sum_{k_a \in \{-1, 1\}} \exp\left\{k_a \left(\beta_0 + \beta_f \sum_{q \in \partial x_{t-1,s}} x_{t-1,q}\right)\right\}}, \quad (4.2.15)$$

coming from the autologistic model as in (4.2.11). The condition of having a state to depend on the transition probability of the previous state of the pixel as well the states of its predetermined neighborhood is presented graphically in Figure 4.2.2.




Figure 4.2.2: Transition of a state via a transition matrix and the corresponding neighbourhood region of the MRF. On the left lattice an arbitrary pixel s is shown, denoted with $x_{t-1,s}$ at time $t-1$, either on the 'on' or 'off' state along with its neighbourhood structure, assigned with the blue solid line. Based on assumption (A2), the state of the pixel s at time t , $x_{t,s}$, depends on the transition probability of $x_{t-1,s}$ according to ξ and the state of its neighbourhood.

Similar to Equation (4.2.3), the general form of the Markov switching model based on SiMPa functions is fully determined from the joint density of the hidden states $\mathbf{x}_{1:T}$ and the observed intensities $\mathbf{y}_{1:T}$, hence

$$\begin{aligned}
 p(\mathbf{x}_{1:T}, \mathbf{y}_{1:T} | \boldsymbol{\vartheta}) &\propto f(\mathbf{y}_{1:T} | \mathbf{x}_{1:T}, \boldsymbol{\vartheta}) p(\mathbf{x}_{1:T} | \boldsymbol{\vartheta}) \\
 &\propto \prod_{t=1}^T f(\mathbf{y}_t | \mathbf{y}_{1:t-1}, \mathbf{x}_{1:t}, \boldsymbol{\vartheta}) p(\mathbf{x}_t | \mathbf{x}_{1:t-1}, \mathbf{y}_{1:t-1}, \boldsymbol{\vartheta}) p(\mathbf{x}_0 | \boldsymbol{\vartheta}) \\
 &\stackrel{(A1)}{\propto} p(\mathbf{x}_0 | \boldsymbol{\phi}) \prod_{t=1}^T f(\mathbf{y}_t | \mathbf{x}_t, \boldsymbol{\phi}) p(\mathbf{x}_t | \mathbf{x}_{t-1}, \boldsymbol{\vartheta}) \\
 &\stackrel{(A2)}{\propto} p(\mathbf{x}_0 | \boldsymbol{\phi}) \prod_{t=1}^T f(\mathbf{y}_t | \mathbf{x}_t, \boldsymbol{\phi}) p(\mathbf{x}_t | \mathbf{x}_{t-1}, \boldsymbol{\xi}) p(\mathbf{x}_{t-1} | \partial \mathbf{x}_{t-1}, \boldsymbol{\phi}), \quad (4.2.16)
 \end{aligned}$$

where $p(\mathbf{x}_0 | \boldsymbol{\phi})$ is the initial distribution of the states, which we assume is independent

of the transition matrix $\boldsymbol{\xi}$, hence $p(\mathbf{x}_0|\boldsymbol{\vartheta}) \equiv p(\mathbf{x}_0|\boldsymbol{\phi})$.

In other words, the joint distribution of $\mathbf{y}_{1:T}$ and $\mathbf{x}_{1:T}$ is the product of the likelihood of $\mathbf{x}_{1:T}$ and $\boldsymbol{\vartheta}$, when the intensities $\mathbf{y}_{1:T}$ are observed, and the marginal density of $\mathbf{x}_{1:T}$. The likelihood can be written as the product of the conditional densities of \mathbf{y}_t when the states $\mathbf{x}_{1:t}$ are known and $\mathbf{y}_{1:t-1}$ have been observed $\forall t \in \{1, \dots, T\}$. According to assumption **(A1)** of conditional independence of \mathbf{y}_t 's when conditioning on the states, the likelihood simplified to

$$\begin{aligned} f(\mathbf{y}_{1:T}|\mathbf{x}_{1:T}, \boldsymbol{\phi}) &\propto \prod_{t=1}^T f(\mathbf{y}_t|\mathbf{y}_{1:t-1}, \mathbf{x}_{1:t}, \boldsymbol{\vartheta}) \\ &\stackrel{(A1)}{\propto} \prod_{t=1}^T f(\mathbf{y}_t|\mathbf{x}_t, \boldsymbol{\phi}) \\ &\propto \prod_{t=1}^T \prod_{s \in \{1, \dots, N\}} f(y_{t,s}|x_{t,s}, \boldsymbol{\phi}), \end{aligned} \quad (4.2.17)$$

while the marginal density of $\mathbf{x}_{1:T}$ can be written as the product of the conditional densities of \mathbf{x}_t , $\forall t \in \{1, \dots, T\}$, which depend only on the states of the neighbourhood of the previous time point $t - 1$ with respect to assumption **(A2)**, thus

$$\begin{aligned} p(\mathbf{x}_{1:T}|\boldsymbol{\vartheta}) &\propto p(\mathbf{x}_0|\boldsymbol{\phi}) \prod_{t=1}^T p(\mathbf{x}_t|\mathbf{x}_{1:t-1}, \boldsymbol{\vartheta}) \\ &\stackrel{(A2)}{\propto} p(\mathbf{x}_0|\boldsymbol{\phi}) \prod_{t=1}^T p(\mathbf{x}_t|\mathbf{x}_{t-1}, \boldsymbol{\vartheta}) \\ &\propto p(\mathbf{x}_0|\boldsymbol{\phi}) \prod_{t=1}^T \prod_{s \in \{1, \dots, N\}} p(x_{t,s}|x_{t-1,s}, \boldsymbol{\vartheta}) \\ &\propto p(\mathbf{x}_0|\boldsymbol{\phi}) \prod_{t=1}^T \prod_{s \in \{1, \dots, N\}} p(x_{t,s}|x_{t-1,s}, \boldsymbol{\xi}) p(x_{t-1,s}|\partial x_{t-1,s}, \boldsymbol{\phi}). \end{aligned} \quad (4.2.18)$$

On the last line of Equation (4.2.18), the conditional density of $x_{t,s}$ has been written as the product of the transition density from $x_{t-1,s}$ to $x_{t,s}$ and the probability density of the Markov field, which takes into account the neighbourhood structure.

4.2.3 Bayesian inference for the Markov Switching model based on SiMPa functions

In this section we describe the general form of Bayesian inference for the Markov Switching model based on SiMPa functions. This is based on the full posterior distribution, which is written as the product of the joint likelihood of the data and states and the prior distributions, hence its general form reads in,

$$\begin{aligned}
 p(\mathbf{x}_{1:T}, \boldsymbol{\vartheta} | \mathbf{y}_{1:T}) \equiv p(\mathbf{x}_{1:T}, \boldsymbol{\phi}, \boldsymbol{\xi} | \mathbf{y}_{1:T}) &\propto p(\mathbf{x}_{1:T}, \mathbf{y}_{1:T} | \boldsymbol{\phi}, \boldsymbol{\xi}) p(\boldsymbol{\phi}) p(\boldsymbol{\xi}) \\
 &\stackrel{(4.2.16)}{\propto} f(\mathbf{y}_{1:T} | \mathbf{x}_{1:T}, \boldsymbol{\phi}) p(\mathbf{x}_{1:T} | \boldsymbol{\phi}, \boldsymbol{\xi}) p(\boldsymbol{\phi}) p(\boldsymbol{\xi}).
 \end{aligned}
 \tag{4.2.19}$$

Sampling directly from the full posterior is not feasible, therefore a sequential MCMC sampling scheme is considered. It consists of three main parts which rely on conditional and sequential updates. The first part accounts for updating the transition matrix $\boldsymbol{\xi}$ based on data augmentation as has been described in [Frühwirth-Schnatter \[2006\]](#). The second part consists of parameter updating when the states of the hidden process are considered known. Regarding the unknown parameters $\boldsymbol{\phi} = (I, c, \tau_b, \boldsymbol{\beta}_{1:T}, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{1:T})$, we adopt a similar concept as in Chapter 3 by using the full conditional distributions. While we condition on the states, our probabilistic counting and localising inference scheme presented in Chapter 3 still stands, however consecutive frames must somehow be connected as their states exhibit a temporal dependency, progressing stochastically according to the transition matrix $\boldsymbol{\xi}$. The corresponding parameter updating scheme for $\boldsymbol{\vartheta} = \{\boldsymbol{\phi}, \boldsymbol{\xi}\}$ when the states

are fixed is presented in Section 4.2.3.1 in detail. The third and last part accounts for estimating the states of the unknown process $\mathbf{x}_{1:T}$ when all the parameters of the Markov switching model based on SiMPa functions are sampled and assumed known. We adopt a Forward-Filtering-Backward-Sampling (FFBS) algorithm, in a similar manner as Frühwirth-Schnatter [2006], to update the hidden states. The FFBS algorithm consists of a sequential procedure, initially running a filtering step forward in time, that is for $t = 1, \dots, T$, based on the filtered and one-step ahead probability densities, followed by a multi-move backwards sampling algorithm, that is for $t = T, T - 1, \dots, 1$, using the conditional distribution of the states when the future states are considered known, presented in Section 4.2.3.2. This section concludes with a Pseudo-algorithm in 4.2.3.3, summarising the steps in order to draw inference on the Markov switching model based on the SiMPa functions from a Bayesian perspective via an FFBS state update algorithm.

4.2.3.1 Parameter updating and localisation given the states

In this section we describe the parameter updating scheme as well as fluorophore localisation based on SiMPa functions. The procedures are similar to those in Section 3.3.2, adapted to the appropriate notation and modifications to account for the Markov switching model. Parameter updating is performed using the full conditional distributions of each one of the parameters via a sequential update procedure onto an MCMC sampling scheme, as sampling directly from the full posterior $p(\mathbf{x}_{1:T}, \boldsymbol{\vartheta} | \mathbf{y}_{1:T})$ in Equation 4.2.19 is extremely challenging, if not impossible. Localisation of fluorophores relies on our novel probabilistic inference scheme using the pattern-configuration-realisation concept, as introduced in Section 3.3.3, relying to a discretisation of the continuous parameters r and θ .

One of the key concepts when applying the modelling scheme based on the SiMPa functions, is the use of a mechanism which attempts to capture the behaviour of a fluorophore intensity spread along a predetermined neighbourhood. Since the SiMPa functions describe the diffraction on a 3×3 region of pixels (or 2nd order neighbour-

hood structure), a same sized moving region (MR) scans the lattice adding intensities in an attempt to capture and identify the corresponding diffractions. More specifically, the likelihood as presented in Equation (4.2.17) is described by,

$$f(\mathbf{y}_{1:T}|\mathbf{x}_{1:T}, \boldsymbol{\phi}) \propto \prod_{t=1}^T \prod_{\forall s \in \{1, \dots, N\}} f(y_{t,s}|x_{t,s}, \boldsymbol{\phi}),$$

due to assumption **(A1)** of conditional independence of $\mathbf{y}_{1:T}$ when $\mathbf{x}_{1:T}$ are known with $\boldsymbol{\phi} = (I, c, \tau_b, \boldsymbol{\beta}_{1:T}, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{1:T})$, where from (4.2.5) and (4.2.7),

$$f(y_{t,s}|x_{t,s}, \boldsymbol{\phi}) \sim \text{N}(\mu_{x_{t,s}}, \tau_b), \quad (4.2.20)$$

with τ_b to be denoting precision and

$$\mu_{x_{t,s}} = \begin{cases} I (1 - \alpha_O(r, \theta, c)), & \text{if } x_{t,s} = 1 \\ I d, & \text{if } x_{t,s} = -1 \end{cases}.$$

We summarise the probability density function of the intensity of a specific pixel s at time t in Equation (4.2.20) by introducing the indicator functions $\mathbb{1}_{\{x_{t,s}=1\}}$ and $\mathbb{1}_{\{x_{t,s}=-1\}}$ where,

$$\mathbb{1}_{\{x_{t,s}=1\}} = \begin{cases} 1, & \text{if } x_{t,s} = 1 \\ 0, & \text{if } x_{t,s} = -1 \end{cases} \quad \text{and} \quad \mathbb{1}_{\{x_{t,s}=-1\}} = \begin{cases} 0, & \text{if } x_{t,s} = 1 \\ 1, & \text{if } x_{t,s} = -1 \end{cases},$$

leading to

$$f(y_{t,s}|x_{t,s}, \boldsymbol{\phi}) \sim \text{N}\left(I d \mathbb{1}_{\{x_{t,s}=-1\}} + I (1 - \alpha_O(r, \theta, c)) \mathbb{1}_{\{x_{t,s}=1\}}, \tau_b\right). \quad (4.2.21)$$

The data in the MR can be obtained by adding the intensities in the predetermined

neighbourhood region, hence $\text{MR}_{t,s} = \sum_{q \in \partial y_{t,s}} y_{t,q}$ with,

$$f(\text{MR}_{t,s} | \partial x_{t,s}, \phi) \sim \text{N} \left(\sum_{q \in \partial x_{t,s}} \mu_{x_{t,q}}, 1 / \sum_{q \in \partial x_{t,s}} \tau_b^{-1} \right), \quad (4.2.22)$$

due to conditional independence of \mathbf{y}_t when the states \mathbf{x}_t are considered known, and τ_b is the background error precision. The precision $\tau_{\text{MR}_{t,s}}$ of the probability density of $\text{MR}_{t,s}$ can be simplified into,

$$\tau_{\text{MR}_{t,s}} = 1 / \sum_{q \in \partial x_{t,s}} \tau_b^{-1} = (N_p \tau_b^{-1})^{-1} = \frac{\tau_b}{N_p}, \quad (4.2.23)$$

with N_p being the total number of pixel in the predetermined neighbourhood, here $N_p = 9$, while the mean $\mu_{\text{MR}_{t,s}} = \sum_{q \in \partial x_{t,s}} \mu_{x_{t,q}}$ of the probability density of $\text{MR}_{t,s}$, can be obtained by,

$$\begin{aligned} \mu_{\text{MR}_{t,s}} &= \sum_{q \in \partial x_{t,s}} \mu_{x_{t,q}} = I \left(d \sum_{q \in \partial x_{t,s}} \mathbb{1}_{\{x_{t,q}=-1\}} + \sum_{q \in \partial x_{t,s}} \mathbb{1}_{\{x_{t,q}=1\}} (1 - \alpha_{\text{O}}(r, \theta, c)) \right) \\ &= I \left(dN_{\{\partial x_{t,s}=-1\}} + \sum_{v=1}^{N_{\{\partial x_{t,s}=1\}}} (1 - \alpha_{\text{O},v}(r, \theta, c)) \right) \\ &= I \left(dN_{\{\partial x_{t,s}=-1\}} + (N_{\{\partial x_{t,s}=1\}} - \alpha_{\text{O},N_{\{\partial x_{t,s}=1\}}}(r, \theta, c)) \right), \end{aligned} \quad (4.2.24)$$

where $N_{\{\partial x_{t,s}=1\}}$ and $N_{\{\partial x_{t,s}=-1\}}$ denote the total number of 'On' and 'Off' pixels respectively in the neighbourhood of pixel s on frame t . The term $\alpha_{\text{O},v}(r, \theta, c)$ represents the corresponding SiMPa functions for each 'On' pixel v , while the quantity $\alpha_{\text{O},N_{\{\partial x_{t,s}=1\}}}(r, \theta, c)$ the sum of all the corresponding SiMPa functions for every 'On' pixel $v = 1, \dots, N_{\{\partial x_{t,s}=1\}}$ within the $\text{MR}_{t,s}$. Similarly to the modifications in Section

3.3.3 in Chapter 3, we adjust μ_{MR} in a way to take into account the number of PSF's contributing to the MR, if this information was available, allowing overlapping scenarios to potentially occur. Therefore, the modified generalised version of μ_{MR} as in Equation (3.3.10) becomes,

$$\mu_{\text{MR}_{t,s}} = I\left(dN_{\{\partial x_{t,s}=-1\}} + \sum_{v=1}^{N_{\text{PSF}}} (n_{\text{PSF}_v} - c\alpha_{O,n_{\text{PSF}_v}}(r_v, \theta_v))\right), \quad (4.2.25)$$

where the total number of 'On' pixels $N_{\{\partial x_{t,s}=1\}}$ can be factorised into a combination of PSF's, denoted with N_{PSF} . Each one of the PSF's in N_{PSF} has n_{PSF} 'On' pixels, associated with the total sum of the corresponding SiMPa functions $\alpha_{O,n_{\text{PSF}}}$, where every $n_{\text{PSF}_v} \subseteq N_{\{\partial x_{t,s}=1\}}$ for all $v \in \{1, \dots, N_{\text{PSF}}\}$.

Therefore, based on Equations (4.2.25) and (4.2.23), the density of a moving region $\text{MR}_{t,s}$, $\forall s \in \{1, \dots, N\}$ at frame t in Equation (4.2.22), when the states and parameters are known takes the form,

$$f(\text{MR}_{t,s} | \partial x_{t,s}, \phi) \sim N\left(I\left(dN_{\{\partial x_{t,s}=-1\}} + \sum_{v=1}^{N_{\text{PSF}}} (n_{\text{PSF}_v} - c\alpha_{O,n_{\text{PSF}_v}}(r_v, \theta_v))\right), \frac{\tau_b}{N_p}\right). \quad (4.2.26)$$

We consider the same assumption as in (A1), regarding conditional independence of the $\mathbf{MR}_t = \{\text{MR}_{t,s}, \forall s \in \{1, \dots, N\}\}$ in any frame t when the corresponding state configuration \mathbf{x}_t is known. Similarly, we assume that they do not depend on previous moving regions $\mathbf{MR}_{1:t-1}$, hence,

$$f(\mathbf{MR}_t | \mathbf{MR}_{1:t-1}, \mathbf{x}_{1:t}, \boldsymbol{\vartheta}) = f(\mathbf{MR}_t | \mathbf{x}_t, \phi) = \prod_{s=1}^N f(\text{MR}_{t,s} | \partial x_{t,s}, \phi)$$

and accordingly to (4.2.17) the likelihood based on the moving regions results in,

$$\begin{aligned}
f(\mathbf{MR}_{1:T}|\mathbf{x}_{1:T}, \boldsymbol{\phi}) &\propto \prod_{t=1}^T f(\mathbf{MR}_t|\mathbf{MR}_{1:t-1}, \mathbf{x}_{1:t}, \boldsymbol{\vartheta}) \\
&\propto \prod_{t=1}^T \prod_{s=1}^N f(\mathbf{MR}_{t,s}|\partial x_{t,s}, \boldsymbol{\phi}) \\
&\stackrel{(4.2.26)}{\propto} \prod_{t=1}^T \prod_{s=1}^N \mathcal{N}(\mathbf{MR}_{t,s} | \mu_{\mathbf{MR}_{t,s}}, \tau_{\mathbf{MR}_{t,s}}). \tag{4.2.27}
\end{aligned}$$

We divide the parameters into global parameters, frame-wise global parameters and frame-wise local parameters according to their relation with the stack of frames. We adopt a Bayesian framework in which we assign priors on every parameter in $\boldsymbol{\phi} = \{I, c, \tau_b, \boldsymbol{\beta}_{1:T}, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{1:T}\}$ in the following fashion,

$$\text{Global Parameters} \quad \begin{cases} \pi(I) \equiv \pi(I|\mu_I, \tau_I) \sim \mathcal{N}(I|\mu_I, \tau_I) \\ \pi(\tau_b) \equiv \pi(\tau_b|a_\tau, b_\tau) \sim \text{Gamma}(\tau_b|a_\tau, b_\tau) \\ \pi(c) \equiv \pi(c|a_c, b_c) \sim \text{RescaledBeta}(c|a_c, b_c, tr_c) \end{cases} \tag{4.2.28}$$

$$\begin{array}{l} \text{Frame-wise} \\ \text{Global Parameters} \end{array} \quad \begin{cases} \pi(\beta_{0_t}) \equiv \pi(\beta_{0_t}|m_{b_0}, \tau_{b_0}) \sim \mathcal{N}(\beta_{0_t}|m_{b_0_t}, \tau_{b_0_t}) \\ \pi(\beta_{f_t}) \equiv \pi(\beta_{f_t}|m_{b_f}, \tau_{b_f}) \sim \mathcal{N}(\beta_{f_t}|m_{b_{f_t}}, \tau_{b_{f_t}}) \end{cases}, \tag{4.2.29}$$

$$\begin{array}{l} \text{Frame-wise} \\ \text{Local Parameters} \end{array} \quad \left\{ \pi(\{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{t,s}) \sim \text{Uniform}(1/N_{\text{Conf}_{t,s}}) \right. \cdot \tag{4.2.30}$$

We assign a discrete uniform distribution with equally probability $1/N_{\text{Conf}_{t,s}}$ as the prior of the complete localisation set $\{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{t,s}$ for any pixel $s \in \{1, \dots, N\}$ at frame t (Equation (4.2.30)), where $N_{\text{Conf}_{t,s}}$ is the total number of configurations for pixel s based on the discretisation of r and θ . As mentioned earlier, the complete localisation set $\{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{t,s}$ (see Definition 3.3.5) consists of the configuration associated with the realisations for pixel s at frame t (Definitions 3.3.3 and 3.3.4 in Section 3.3.3). The parameters I, τ_b, c are the global parameters, since they depend

on the entire time series from $t = 1, \dots, T$, the β_t 's $\forall t \in \{1, \dots, T\}$ are global parameters on each respective individual frame t and the complete localisation sets $\{\text{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{t,s}$ are local parameters for every pixel s of every frame t .

Regarding the global parameters I , τ_b and c , we obtain the corresponding full conditional distributions based on the entire moving region observations collection $\mathbf{MR}_{1:T}$ and their respective prior distributions in Equation (4.2.28), therefore,

Full conditional of I

$$\begin{aligned}
p(I \mid \mathbf{MR}_{1:T}, \boldsymbol{\phi}_{-I}) &\equiv p(I \mid \mathbf{MR}_{1:T}, \mathbf{x}_{1:T}, \tau_b, c, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{1:T}) \\
&\propto f(\mathbf{MR}_{1:T} \mid \mathbf{x}_{1:T}, I, \boldsymbol{\phi}_{-I}) \pi(I) \\
&\stackrel{(4.2.28)}{\propto} \prod_{t=1}^T \prod_{s=1}^N \text{N}(\text{MR}_{t,s} \mid \mu_{\text{MR}_{t,s}}, \tau_{\text{MR}_{t,s}}) \text{N}(I \mid \mu_I, \tau_I). \\
&\stackrel{(4.2.27)}{\propto} \prod_{t=1}^T \prod_{s=1}^N \text{N}(\text{MR}_{t,s} \mid \mu_{\text{MR}_{t,s}}, \tau_{\text{MR}_{t,s}}) \text{N}(I \mid \mu_I, \tau_I).
\end{aligned} \tag{4.2.31}$$

Therefore, the full conditional of I is,

$$p(I \mid \mathbf{MR}_{1:T}, \boldsymbol{\phi}_{-I}) \sim \text{N}(I \mid \mu_I^{\text{FC}}, \tau_I^{\text{FC}}), \tag{4.2.32}$$

where the full conditional mean μ_I^{FC} and precision τ_I^{FC} ,

$$\begin{aligned}
\mu_I^{\text{FC}} &= \frac{\tau_I \mu_I + \frac{\tau_b}{N_p} \sum_{t=1}^T \sum_{s=1}^N \text{MR}_{t,s} \overbrace{\left(dN_{\{\partial x_{t,s} = -1\}} + \sum_{v=1}^{\text{N}_{\text{PSF}}} (n_{\text{PSF}_v} - c\alpha_{\text{O}, n_{\text{PSF}_v}}(r_v, \theta_v)) \right)}^{\mu_{\text{MR}_{t,s}}/I}}{\tau_I^{\text{FS}}}, \\
\tau_I^{\text{FC}} &= \tau_I + \frac{\tau_b}{N_p} \sum_{t=1}^T \sum_{s=1}^N \underbrace{\left(dN_{\{\partial x_{t,s} = -1\}} + \sum_{v=1}^{\text{N}_{\text{PSF}}} (n_{\text{PSF}_v} - c\alpha_{\text{O}, n_{\text{PSF}_v}}(r_v, \theta_v)) \right)^2}_{\mu_{\text{MR}_{t,s}}/I},
\end{aligned}$$

with the mean of the moving region $\mu_{\text{MR}_{t,s}}$ for pixel s at frame t is as in Equation

(4.2.25).

Full conditional of τ_b

$$\begin{aligned}
p(\tau_b \mid \mathbf{MR}_{1:T}, \boldsymbol{\phi}_{-\tau_b}) &\equiv p(\tau_b \mid \mathbf{MR}_{1:T}, \mathbf{x}_{1:T}, I, c, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{1:T}) \\
&\propto f(\mathbf{MR}_{1:T} \mid \mathbf{x}_{1:T}, \tau_b, \boldsymbol{\phi}_{-\tau_b}) \pi(\tau_b) \\
&\stackrel{(4.2.28)}{\propto} \prod_{t=1}^T \prod_{s=1}^N \mathcal{N}(\mathbf{MR}_{t,s} \mid \mu_{\mathbf{MR}_{t,s}}, \tau_{\mathbf{MR}_{t,s}}) \text{Gamma}(\tau_b \mid a_{\tau_b}, b_{\tau_b}). \\
&\stackrel{(4.2.27)}{\propto} \prod_{t=1}^T \prod_{s=1}^N \mathcal{N}(\mathbf{MR}_{t,s} \mid \mu_{\mathbf{MR}_{t,s}}, \tau_{\mathbf{MR}_{t,s}}) \text{Gamma}(\tau_b \mid a_{\tau_b}, b_{\tau_b}).
\end{aligned} \tag{4.2.33}$$

Similarly, the full conditional distribution of the precision τ_b is in closed, hence,

$$p(\tau_b \mid \mathbf{MR}_{1:T}, \boldsymbol{\phi}_{-\tau_b}) \sim \text{Gamma}(\tau_b \mid a_{\tau_b}^{\text{FC}}, b_{\tau_b}^{\text{FC}}), \tag{4.2.34}$$

where the full conditional shape $a_{\tau_b}^{\text{FC}}$ and rate $b_{\tau_b}^{\text{FC}}$,

$$\begin{aligned}
a_{\tau_b}^{\text{FC}} &= a_{\tau_b} + (T \times N)/2, \\
b_{\tau_b}^{\text{FC}} &= \frac{\sum_{t=1}^T \sum_{s=1}^N \left(\mathbf{MR}_{t,s} - I \left(dN_{\{\partial x_{t,s} = -1\}} + \overbrace{\sum_{v=1}^{\text{NPSF}} (n_{\text{PSF}_v} - c\alpha_{O, n_{\text{PSF}_v}}(r_v, \theta_v))}^{\mu_{\mathbf{MR}_{t,s}}}} \right) \right)^2}{2N_p} + b_{\tau_b}.
\end{aligned}$$

When it comes to updating, usual Gibbs steps are used for both parameters I and τ_b in order to sample from their full conditional distribution throughout an MCMC update scheme. This is not the case for the update of spread parameter c as well the frame wise global parameters $\boldsymbol{\beta}_{1:T}$.

Full conditional of c

The full conditional distribution of c is the product of the normally distributed density of $\mathbf{MR}_{1:T}$ and the rescaled beta prior distribution, chosen with respect to

the limitations of c as discussed in detail on Section 3.2.4, therefore,

$$\begin{aligned}
p(c \mid \mathbf{MR}_{1:T}, \boldsymbol{\phi}_{-c}) &\equiv p(c \mid \mathbf{MR}_{1:T}, \mathbf{x}_{1:T}, I, \tau_b, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{1:T}) \\
&\propto f(\mathbf{MR}_{1:T} \mid \mathbf{x}_{1:T}, c, \boldsymbol{\phi}_{-c}) \pi(c) \\
&\stackrel{(4.2.28)}{\propto} \prod_{t=1}^T \prod_{s=1}^N \mathcal{N}(\mathbf{MR}_{t,s} \mid \mu_{\mathbf{MR}_{t,s}}, \tau_{\mathbf{MR}_{t,s}}) \text{RescBeta}(c \mid a_c, b_c, tr_c), \\
&\stackrel{(4.2.27)}{\propto} \prod_{t=1}^T \prod_{s=1}^N \mathcal{N}(\mathbf{MR}_{t,s} \mid \mu_{\mathbf{MR}_{t,s}}, \tau_{\mathbf{MR}_{t,s}}) \text{RescBeta}(c \mid a_c, b_c, tr_c),
\end{aligned} \tag{4.2.35}$$

which cannot be expressed in exact form. The fixed parameter tr_c denotes the rescaling of the usual Beta distribution with $tr_c = 0.7071$, when the side length of the pixel is $l = 1$ without loss of generality, as obtained on Section 3.2.4.

Therefore, in order to update the parameter c we adopt a Metropolis-Hastings step within MCMC, using a rescaled Beta distribution with mode at the current value of c as the proposal distribution.

Full conditional of $\{\beta_{0_t}, \beta_{f_t}\}$

$$\begin{aligned}
p(\{\beta_{0_t}, \beta_{f_t}\} \mid \mathbf{x}_t) &\propto \pi(\mathbf{x}_t \mid \beta_{0_t}, \beta_{f_t}) \pi(\beta_{0_t}) \pi(\beta_{f_t}) \\
&= \frac{\exp\{\beta_{0_t} V_0(\mathbf{x}_t) + \beta_{f_t} V_f(\mathbf{x}_t)\}}{Z(\{\beta_{0_t}, \beta_{f_t}\})} \pi(\beta_{0_t}) \pi(\beta_{f_t}) \\
&\approx \prod_{s=1}^N \pi(x_{t,s} \mid \partial x_{t,s}, \{\beta_{0_t}, \beta_{f_t}\}) \mathcal{N}(\beta_{0_t} \mid m_{\beta_{0_t}}, \tau_{\beta_{0_t}}) \mathcal{N}(\beta_{f_t} \mid m_{\beta_{f_t}}, \tau_{\beta_{f_t}}).
\end{aligned} \tag{4.2.36}$$

The full conditional of the frame wise global parameters $\boldsymbol{\beta}_{1:T} = \left\{ \{\beta_{0_t}, \beta_{f_t}\}, \forall t = 1, \dots, T \right\}$ for every t , is not exact, hence a Metropolis-Hastings step is considered. The update is based on Gaussian distributed proposal distributions centred at the current values of $\{\beta_{0_t}, \beta_{f_t}\}$. On the last line of Equation (4.2.36), the pseudo-likelihood approximation [Besag, 1974] is used in order to reduce the complexity of calculations as in Equation (4.2.12).

Localisations update based on full conditionals

The localisation of active fluorophores, or the simplified form of pixel identification where the active fluorophores lie onto, is performed according to our novel probabilistic inference scheme introduced in Section 3.3.3 in Chapter 3, based on the pattern-configuration-realisation concept. The latter relies on the assumption of every identified fluorophore being at the origin of the SiMPa functions ($r = \theta = 0$), while the former associates each one with a realisation as well. As a reminder, a realisation is combination of r and θ from a chosen discretisation (Definition 3.3.4). Both approaches rely on the potential configurations a pattern can be formed from, under the assumption of up to two PSFs allowed to overlap. A pattern is a structure of 'On' and 'Off' pixels within a MR (Definition 3.3.2), with the configuration a combination of PSFs which can result in the corresponding pattern (Definition 3.3.3).

Full conditional of complete localisation set $\{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{t,s}$

According to Equation (4.2.14) and Definition 3.3.5, a complete localisation set for a pixel s at frame t is denoted with $\{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{t,s} = \left\{ \mathbf{Conf}_{t,s}, \{r, \theta\}_{t,s(1)}, \dots, \{r, \theta\}_{t,s(v)} \right\}$, where $\{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_t = \left(\{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{t,s}, \forall s \in \{1, \dots, N\} \right)$ the complete localisation set for frame t , as in Equation (4.2.13). Then we can obtain the full conditional distribution for any pixel s in frame t ,

$$\begin{aligned}
 P_{t,s} &\equiv p \left(\{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{t,s} \mid \text{MR}_{t,s}, \partial x_{t,s}, I, c, \tau_b \right) \\
 &\propto f \left(\text{MR}_{t,s} \mid \partial x_{t,s}, I, c, \tau_b, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{t,s} \right) \pi(\{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{t,s}) \\
 &\stackrel{(4.2.30)}{\propto} \text{N} \left(\text{MR}_{t,s} \mid \mu_{\text{MR}_{t,s}}, \tau_{\text{MR}_{t,s}} \right) \text{Uniform}(1/N_{\text{Conf}_{t,s}}), \tag{4.2.37} \\
 &\stackrel{(4.2.26)}{\propto}
 \end{aligned}$$

where the mean $\mu_{\text{MR}_{t,s}}$, from Equation (4.2.25), is evaluated at the appropriate complete localisation set $\{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{t,s}$, contributing to the red coloured parts below,

$$\mu_{\text{MR}_{t,s}} = I \left(dN_{\partial x_{t,s}=-1} + \sum_{v=1}^{\text{NPSF}} (n_{\text{PSF}_v} - \alpha_{0, n_{\text{PSF}_v}}(r_v, \theta_v, c)) \right).$$

The full conditional distribution $P_{t,s}$ is evaluated at every possible formation in $\mathcal{N}_{\text{Conf}_{t,s}}$, denoted with $\mathcal{N}_{\text{Conf}_{t,s}(u)}$, i.e. $\forall 1, \dots, u$ in the finite configurations set for the corresponding pattern of pixel s on frame t . Since $P_{t,s}$ has a finite support, it is multinomial with probabilities proportional to $P_{t,s(1)}, \dots, P_{t,s(u)}$, which is introduced within a Gibbs step to draw localisations, similar to Equation (4.2.37) in Chapter 3, as,

$$\left\{ \begin{array}{l} P_{t,s(1)} = p \left(\{\text{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{t,s(1)} \mid \text{MR}_{t,s}, \partial x_{t,s}, I, c, \tau_b \right) \\ \vdots \\ P_{t,s(u)} = p \left(\{\text{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{t,s(u)} \mid \text{MR}_{t,s}, \partial x_{t,s}, I, c, \tau_b \right) \end{array} \right\} \rightarrow \left\{ \begin{array}{l} P_{t,s(1)} / \sum_{i=1}^u P_{t,s(i)} \\ \vdots \\ P_{t,s(u)} / \sum_{i=1}^u P_{t,s(i)} \end{array} \right\}.$$

Full conditional of complete counting set $\{\text{Conf}, \mathbf{r} = \boldsymbol{\theta} = \mathbf{0}\}_{t,s}$

The complete counting set, $\{\text{Conf}, \mathbf{r} = \boldsymbol{\theta} = \mathbf{0}\}_{t,s} \equiv \{\text{Conf}\}_{t,s}$, is the simplification of the complete localisation set, allowing inference only on the configurations forming the patterns, and relies on fixed positions of the fluorophores within the pixels hence $r = \theta = 0$ (Definition 3.3.5). In that case, the complexity of the probabilistic scheme reduces substantially, by avoiding an enormous number of probabilities to be calculated for each pixel, which can potentially result in less reliable inference. Similar to Section 3.3.3, this modification corresponds to the probabilistic counting scheme, reading in,

$$\left\{ \begin{array}{l} C_{t,s(1)} = p \left(\{\text{Conf}\}_{t,s(1)} \mid \text{MR}_{t,s}, \partial x_{t,s}, I, c, \tau_b \right) \\ \vdots \\ C_{t,s(w)} = p \left(\{\text{Conf}\}_{t,s(w)} \mid \text{MR}_{t,s}, \partial x_{t,s}, I, c, \tau_b \right) \end{array} \right\} \rightarrow \left\{ \begin{array}{l} C_{t,s(1)} / \sum_{h=1}^w C_{t,s(h)} \\ \vdots \\ C_{t,s(w)} / \sum_{h=1}^w C_{t,s(h)} \end{array} \right\},$$

from the simplified full conditional distribution of any configuration $\{\text{Conf}\}_{t,s}$,

$$\begin{aligned}
C_{t,s} &\equiv p\left(\{\text{Conf}\}_{t,s} \mid \text{MR}_{t,s}, \partial x_{t,s}, I, c, \tau_b\right) \\
&\propto f\left(\text{MR}_{t,s} \mid \partial x_{t,s}, I, c, \tau_b, \{\text{Conf}\}_{t,s}\right) \pi(\{\mathbf{Conf}\}_{t,s}) \\
&\stackrel{(4.2.30)}{\propto} \text{N}\left(\text{MR}_{t,s} \mid \mu_{\text{MR}_{t,s}}, \tau_{\text{MR}_{t,s}}\right) \text{Uniform}(1/N_{\text{Conf}_{t,s}}), \quad (4.2.38) \\
&\stackrel{(4.2.26)}{\propto}
\end{aligned}$$

where a similar Uniform prior is used for the number of configurations $N_{\text{Conf}_{t,s}}$, assigning equal prior probability on each configuration, with $N_{\text{Conf}_{t,s}(w)}$ denoting every possible configuration. As for $P_{t,s}$, $C_{t,s}$ has a finite support and thus it is multinomial with probabilities proportional to $C_{t,s(1)}, \dots, C_{t,s(w)}$. Their corresponding contributions to the mean $\mu_{\text{MR}_{t,s}}$ from Equation (4.2.25) is outlined with the red coloured parts below,

$$\mu_{\text{MR}_{t,s}} = I\left(dN_{\partial x_{t,s}=-1} + \sum_{v=1}^{\text{NPSF}} (n_{\text{PSF}_v} - \alpha_{\text{O}, n_{\text{PSF}_v}}(r_v = 0, \theta_v = 0, c))\right).$$

Full conditional of the transition matrix ξ

The full conditional distribution of ξ is obtained as the product of the marginal density of $\mathbf{x}_{1:T}$ and the prior distribution for ξ , while it does not depend on the observed intensities but only on the configuration states of $\mathbf{x}_{1:T} = (\mathbf{x}_1, \dots, \mathbf{x}_T)$. Therefore,

$$p(\xi \mid \mathbf{x}_{1:T}) \propto p(\mathbf{x}_{1:T} \mid \xi) p(\xi). \quad (4.2.39)$$

Recalling from Equation (4.2.18), the marginal density of $\mathbf{x}_{1:T}$ reads,

$$\begin{aligned}
p(\mathbf{x}_{1:T} \mid \boldsymbol{\vartheta}) &\propto p(\mathbf{x}_0 \mid \phi) \prod_{t=1}^T p(\mathbf{x}_t \mid \mathbf{x}_{t-1}, \xi) p(\mathbf{x}_{t-1} \mid \boldsymbol{\vartheta}) \\
&\propto p(\mathbf{x}_0 \mid \phi) \prod_{t=1}^T \prod_{s=1}^N p(x_{t,s} \mid x_{t-1,s}, \xi) p(x_{t-1,s} \mid \partial x_{t-1,s}, \phi)
\end{aligned}$$

and for the full conditional of ξ as in Equation (4.2.39) only the terms that include ξ contribute when the configurations states $\mathbf{x}_{1:T}$ are known, hence

$$\begin{aligned}
p(\mathbf{x}_{1:T}|\xi) &\propto \prod_{t=1}^T p(\mathbf{x}_t|\mathbf{x}_{t-1}, \xi) \propto \prod_{t=1}^T \prod_{s=1}^N p(x_{t,s}|x_{t-1,s}, \xi) \\
&\propto \prod_{t=1}^T \prod_{s=1}^N \xi_{x_{t-1,s}, x_{t,s}} \\
&\propto \prod_{s=1}^T \prod_{j_1=0}^1 \prod_{j_2=0}^1 \xi_{j_1, j_2}^{n_{j_1, j_2}(\mathbf{x}_{1:T, s})} \\
&\propto \prod_{j_1=0}^1 \prod_{j_2=0}^1 \xi_{j_1, j_2}^{N_{j_1, j_2}(\mathbf{x}_{1:T})}, \tag{4.2.40}
\end{aligned}$$

where $n_{j_1, j_2}(\mathbf{x}_{1:T, s})$ the number of times a state transition occurs from j_1 to j_2 in the trajectory of a pixel $s \in \{1, \dots, N\}$ over all time points from $t = 1, \dots, N$. Similarly, the quantity $N_{j_1, j_2}(\mathbf{x}_{1:T})$ contains the total number of transitions from j_1 to j_2 for every pixel $s \in \{1, \dots, N\}$ hence,

$$n_{j_1, j_2}(\mathbf{x}_{1:T, s}) = \#\{x_{t,s} = j_1, x_{t-1,s} = j_2\}, \tag{4.2.41}$$

$$N_{j_1, j_2}(\mathbf{x}_{1:T}) = \sum_{s=1}^N n_{j_1, j_2}(\mathbf{x}_{1:T, s}) = \sum_{s=1}^N \#\{x_{t,s} = j_1, x_{t-1,s} = j_2\}. \tag{4.2.42}$$

We assume a-priori independence on the rows of the transition matrix ξ in Equation (4.2.8), translating into $\xi_{1, \bullet}$ independent of $\xi_{0, \bullet}$. We assign a Dirichlet distribution on each row $\xi_{j_1, \bullet}$ for $j_1 = 0, 1$ of ξ , similar to two independent Beta distributions since $K = 2$, hence

$$\begin{aligned}
\pi(\xi_{j_1, \bullet} | e_{j_1, 1}, e_{j_1, 0}) &\sim \text{Dirichlet}(\xi_{j_1, \bullet} | e_{j_1, 1}, e_{j_1, 0}) \\
&\propto \prod_{j_2=0, 1} \xi_{j_1, j_2}^{(e_{j_1, j_2} - 1)}. \tag{4.2.43}
\end{aligned}$$

Using Equation (4.2.39), the rows $\boldsymbol{\xi}_{1,\bullet}$ and $\boldsymbol{\xi}_{0,\bullet}$ hold their independence a-posteriori, and their posterior distributions for $j_1 = 0, 1$ can be obtained, as

$$\begin{aligned}
 p(\boldsymbol{\xi}_{j_1,\bullet} | \mathbf{x}_{1:T}) &\propto p(\mathbf{x}_{1:T} | \boldsymbol{\xi}_{j_1,\bullet}) p(\boldsymbol{\xi}_{j_1,\bullet}) \\
 &\stackrel{(4.2.40)}{\propto} \prod_{j_2=0}^1 \zeta_{j_1,j_2}^{N_{j_1,j_2}(\mathbf{x}_{1:T})} \prod_{j_2=0,1} \zeta_{j_1,j_2}^{(e_{j_1,j_2}-1)} \\
 &\stackrel{(4.2.43)}{\propto} \prod_{j_2=0}^1 \zeta_{j_1,j_2}^{(N_{j_1,j_2}(\mathbf{x}_{1:T})+e_{j_1,j_2}-1)},
 \end{aligned}$$

which is in closed expression, recognised as a Dirichlet distribution with,

$$p(\boldsymbol{\xi}_{j_1,\bullet} | \mathbf{x}_{1:T}) \sim \text{Dirichlet}\left(\boldsymbol{\xi}_{j_1,\bullet} \mid e_{j_1,1} + N_{j_1,1}(\mathbf{x}_{1:T}), e_{j_1,0} + N_{j_1,0}(\mathbf{x}_{1:T})\right), \quad (4.2.44)$$

where $N_{j_1,1}(\mathbf{x}_{1:T})$ and $N_{j_1,0}(\mathbf{x}_{1:T})$ the quantities denoted in Equations (4.2.41) and (4.2.42).

Duration of a state

In order to have a mechanism to distinguish between reappearing or different active fluorophores on consecutive frames within the stack, we consider a parameter D_{k_c} accounting for the duration of a state $k_c \in K = \{1, -1\}$. Then,

$$\Pr(D_{k_c} = l_{k_c} \mid \mathbf{x}_{1:T}, \boldsymbol{\vartheta}) = \prod_{s=1}^N \Pr(D_{k_c} = l_{k_c} \mid \mathbf{x}_{1:T,s}, \boldsymbol{\vartheta}), \quad (4.2.45)$$

where l_{k_c} a length of state k_c . Our assumption is that the probability of the duration of a state k_c to be l_{k_c} based on the configuration states $\mathbf{x}_{1:T}$ is the product of the individual probabilities of duration l_{k_c} for every pixel s when their corresponding trajectories of the neighbourhood are known. These probabilities which we denote

by $p_{l_{k_c}}$ can be obtained,

$$\begin{aligned}
p_{l_{k_c}} &\equiv \Pr(D_{k_c} = l_{k_c} \mid \mathbf{x}_{1:T,s}, \boldsymbol{\vartheta}) \\
&\propto \prod_{\tau \in \{1, \dots, T\}} \Pr(x_{\tau+1,s} = \dots = x_{\tau+(l_{k_c}-1),s} = k_c, x_{\tau+l_{k_c},s} = k'_c \mid x_{\tau,s} = k_c, \boldsymbol{\vartheta}) \\
&\propto \prod_{\tau \in \{1, \dots, T\}} \xi_{k_c, k'_c} \Pr(x_{\tau+(l_{k_c}-1),s} = k_c \mid \partial x_{\tau+(l_{k_c}-1),s}, \boldsymbol{\vartheta}) \\
&\quad \times \prod_{m=1}^{l_{k_c}-1} \xi_{k_c, k_c} \Pr(x_{\tau+(m-1),s} = k_c \mid \partial x_{\tau+(m-1),s}, \boldsymbol{\vartheta}) \\
&\propto \prod_{\tau \in \{1, \dots, T\}} \xi_{k_c, k'_c} \phi_{k_c}(x_{\tau+(l_{k_c}-1),s}) \prod_{m=1}^{l_{k_c}-1} \xi_{k_c, k_c} \phi_{k_c}(x_{\tau+(m-1),s}),
\end{aligned} \tag{4.2.46}$$

where $\phi_{k_c}(x_{\tau,s}) = \Pr(x_{\tau,s} = k_c \mid \partial x_{\tau,s}, \boldsymbol{\vartheta})$. Using Equation (4.2.45), the probability of a state k_c having length l_{k_c} based on the configuration states $\mathbf{x}_{1:T}$,

$$\Pr(D_{k_c} = l_{k_c} \mid \mathbf{x}_{1:T}, \boldsymbol{\vartheta}) \propto \prod_{s=1}^N \prod_{\tau \in \{1, \dots, T\}} \xi_{k_c, k'_c} \phi_{k_c}(x_{\tau+(l_{k_c}-1),s}) \prod_{m=1}^{l_{k_c}-1} \xi_{k_c, k_c} \phi_{k_c}(x_{\tau+(m-1),s}). \tag{4.2.47}$$

Therefore, since the number of frames in the stack is discrete with $t = 1, \dots, T$, the length a pixel s can remain on a state k_c will be also discrete with maximum duration to be the total number of frames T , when assuming no prior knowledge. Let L be the maximum length with $L \leq T$, then we can calculate the probability of duration of state k_c being l_{k_c} ,

$$\Pr(D_{k_c} = l_{k_c} \mid \mathbf{x}_{1:T}, \boldsymbol{\vartheta}) = \frac{\prod_{s=1}^N \prod_{\tau \in \{1, \dots, T\}} \xi_{k_c, k'_c} \phi_{k_c}(x_{\tau+(l_{k_c}-1),s}) \prod_{m=1}^{l_{k_c}-1} \xi_{k_c, k_c} \phi_{k_c}(x_{\tau+(m-1),s})}{\sum_{l_a=1}^L \prod_{s=1}^N \prod_{\tau \in \{1, \dots, T\}} \xi_{k_c, k'_c} \phi_{k_c}(x_{\tau+(l_a-1),s}) \prod_{m=1}^{l_a-1} \xi_{k_c, k_c} \phi_{k_c}(x_{\tau+(m-1),s})}.$$

(4.2.48)

We can straightforward obtain the expected duration of state k_c ,

$$E(D_{k_c}) = \sum_{l_a=1}^L l_a \times \Pr(D_{k_c} = l_a \mid \mathbf{x}_{1:T}, \boldsymbol{\vartheta}). \quad (4.2.49)$$

The expected duration in Equation (4.2.49) is our main tool to distinguish between different and reappearing active fluorophores between the frames on the stack. However, the probability in Equation (4.2.48) is extremely challenging to calculate, therefore we consider an approximation based on the average states of the fields, discussed during the implementation in Section 4.4.

The entire parameter updating for the Markov switching model based on the SiMPa functions is summarised on the following Pseudo-algorithm.

Pseudo-algorithm - Parameter updating for the Markov Switching model based on SiMPa functions when the states are known

1. Step $q=0$ (Iteration 0):

Initialisation of parameters by sampling from their prior distributions in Equations (4.2.28), (4.2.29) and (4.2.30), resulting in

$$\begin{aligned} \boldsymbol{\vartheta}^{(0)} &= \left\{ I^{(0)}, c^{(0)}, \tau_b^{(0)}, \boldsymbol{\beta}_{1:T}^{(0)}, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{1:T}^{(0)}, \boldsymbol{\xi}^{(0)} \right\} \\ \text{where } \boldsymbol{\beta}_{1:T}^{(0)} &= \left(\{\beta_{0_1}^{(0)}, \beta_{f_1}^{(0)}\}, \dots, \{\beta_{0_T}^{(0)}, \beta_{f_T}^{(0)}\} \right) \\ \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{1:T}^{(0)} &= \left(\{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_1^{(0)}, \dots, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_T^{(0)} \right) \end{aligned}$$

where the collection $\{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{1:T}^{(0)}$ contains all the localisations based on the pattern-configuration-realisation probabilistic scheme for every pixel in every frame. Without loss of generality we can assume that the initial

localisations for all pixels are at the origin of the SiMPa functions, i.e. $\{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{1:T}^{(0)} = \{\mathbf{Conf}, \mathbf{0}, \mathbf{0}\}_{1:T}^{(0)}$.

As discussed throughout the Chapter, the configuration states $\mathbf{x}_{1:T} = (\mathbf{x}_1, \dots, \mathbf{x}_T)$ are assumed known.

2. *Step $q=1, \dots, IT$ (Iteration 1 up to IT)*

Estimation of parameters in $\boldsymbol{\vartheta}^{(q)} = \{I^{(q)}, c^{(q)}, \tau_b^{(q)}, \boldsymbol{\beta}_{1:T}^{(q)}, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{1:T}^{(q)}, \boldsymbol{\xi}^{(q)}\}$ and probabilistic draw of localisations of molecules based on SiMPa functions conditional on knowing $\mathbf{x}_{1:T}^{(q-1)}$.

- (a) Update the transition matrix $\boldsymbol{\xi}^{(q)}$ from its full conditional distribution for each one of the rows $\boldsymbol{\xi}_{j_1, \bullet}^{(q)}$,

$$\boldsymbol{\xi}_{j_1, \bullet}^{(q)} \sim p\left(\boldsymbol{\xi}_{j_1, \bullet} \mid \mathbf{x}_{1:T}^{(q-1)}\right)$$

from Equation(4.2.44)

- (b) Update the parameters $I^{(q)}, \tau_b^{(q)}$ (Gibbs steps), $c^{(q)}, \boldsymbol{\beta}_{1:T}^{(q)}$ (Metropolis-Hastings steps) from their corresponding full conditional distributions using the entire time series from $1 : T$ and the moving regions $\mathbf{MR}_{1:T} = (\mathbf{MR}_1, \dots, \mathbf{MR}_T)$ where $\mathbf{MR}_t = \{\text{MR}_{t,s}, \forall s \in \{1, \dots, N\}\}$, $\forall t = (1, \dots, T)$.

$$I^{(q)} \sim p\left(I \mid \mathbf{x}_{1:T}^{(q-1)}, \mathbf{MR}_{1:T}, \tau_b^{(q-1)}, c^{(q-1)}, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{1:T}^{(q-1)}\right)$$

from Equation(4.2.32)

$$\tau_b^{(q)} \sim p\left(\tau_b \mid \mathbf{x}_{1:T}^{(q-1)}, \mathbf{MR}_{1:T}, I^{(q)}, c^{(q-1)}, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{1:T}^{(q-1)}\right)$$

from Equation(4.2.34)

$$c^{(q)} \sim p\left(c \mid \mathbf{x}_{1:T}^{(q-1)}, \mathbf{MR}_{1:T}, I^{(q)}, \tau_b^{(q)}, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{1:T}^{(q-1)}\right)$$

$$\begin{aligned} & \text{from Equation(4.2.35)} \\ \{\beta_{0_t}, \beta_{f_t}\}^{(q)} & \sim p\left(\{\beta_0, \beta_f\} \mid \mathbf{x}_{1:T}^{(q-1)}\right) \\ & \text{from Equation(4.2.36)} \end{aligned}$$

(c) Apply either the *localisation scheme* or *counting scheme*.

Localisation scheme

For each frame $t \in (1, \dots, T)$ on iteration q , apply the individual frame localisation scheme based on SiMPa functions, based on the most up to date configuration states of fields $\mathbf{x}_{1:T}^{(q-1)} = (\mathbf{x}_1^{(q-1)}, \dots, \mathbf{x}_T^{(q-1)})$.

For every $\mathbf{MR}_t = \{\text{MR}_{t,s}, s \in \{1, \dots, N\}\}$, evaluate the full conditional distributions in Equation (4.2.37),

$$\begin{aligned} P_{t,s(1)}^{(q)} &= p\left(\{\text{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{t,s(1)} \mid \text{MR}_{t,s}, \partial x_{t,s}^{(q-1)}, I^{(q)}, c^{(q)}, \tau_b^{(q)}\right) \\ &\vdots \\ P_{t,s(u)}^{(q)} &= p\left(\{\text{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{t,s(u)} \mid \text{MR}_{t,s}, \partial x_{t,s}^{(q-1)}, I^{(q)}, c^{(q)}, \tau_b^{(q)}\right) \end{aligned}$$

with localisation probabilities obtained after normalising the densities $P_{t,s(1)}^{(q)}, \dots, P_{t,s(u)}^{(q)}$, that is

$$P_{t,s(1)}^{(q)} \Big/ \sum_{h=1}^u P_{t,s(h)}^{(q)}, \dots, P_{t,s(u)}^{(q)} \Big/ \sum_{h=1}^u P_{t,s(h)}^{(q)}$$

Counting scheme

For each frame $t \in (1, \dots, T)$ on iteration q , apply the individual frame counting scheme based on SiMPa functions, based on the most up to date configuration states of fields $\mathbf{x}_{1:T}^{(q-1)} = (\mathbf{x}_1^{(q-1)}, \dots, \mathbf{x}_T^{(q-1)})$.

For every $\mathbf{MR}_t = \{\text{MR}_{t,s}, s \in \{1, \dots, N\}\}$, evaluate the full conditional

distributions in Equation (4.2.38),

$$\begin{aligned} C_{t,s(1)}^{(q)} &= p\left(\{\text{Conf}, \mathbf{r} = \mathbf{0}, \boldsymbol{\theta} = \mathbf{0}\}_{t,s(1)} \mid \text{MR}_{t,s}, \partial x_{t,s}^{(q-1)}, I^{(q)}, c^{(q)}, \tau_b^{(q)}\right) \\ &\vdots \\ C_{t,s(w)}^{(q)} &= p\left(\{\text{Conf}, \mathbf{r} = \mathbf{0}, \boldsymbol{\theta} = \mathbf{0}\}_{t,s(w)} \mid \text{MR}_{t,s}, \partial x_{t,s}^{(q-1)}, I^{(q)}, c^{(q)}, \tau_b^{(q)}\right) \end{aligned}$$

with localisation probabilities obtained after normalising the densities

$C_{t,s(1)}^{(q)}, \dots, C_{t,s(w)}^{(q)}$, that is

$$C_{t,s(1)}^{(q)} / \sum_{h=1}^w C_{t,s(h)}^{(q)}, \dots, C_{t,s(w)}^{(q)} / \sum_{h=1}^w C_{t,s(h)}^{(q)}$$

3. Store the most up to date parameters in

$$\boldsymbol{\vartheta}^{(q)} = \left\{ I^{(q)}, c^{(q)}, \tau_b^{(q)}, \boldsymbol{\beta}_{1:T}^{(q)}, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{1:T}^{(q)}, \boldsymbol{\xi}^{(q)} \right\}$$

4.2.3.2 State updating given the parameters using data augmentation

In this section we describe the complete state updating based on data augmentation using a Forward-Filtering-Backward-Sampling (FFBS) algorithm. Backward sampling is one of the most popular procedures in order to draw inference on a state-space model with a hidden stochastic process being the underlying mechanism generating the observations. It is based on the idea of firstly moving forward in time, that is $t = 1, \dots, T$, applying a filtering process based on a determined prediction step and then sampling trajectories of the hidden states by moving backwards in time, that is $t = T, \dots, 1$. This allows sampling from hidden state processes under extremely complicated and intractable scenarios and has been increasingly studied over the past decades, e.g. for non linear state space model by Fong et al. [2002] and jump Markov linear systems by Doucet et al. [2000].

In terms of probability modelling, the filtering procedure is carried out using the

filtering density, defined by $p(\mathbf{x}_t|\mathbf{y}_{1:t}, \boldsymbol{\vartheta})$, while the backward sampling scheme is implemented with the commonly known joint smoothing density $p(\mathbf{x}_{1:T}|\mathbf{y}_{1:T}, \boldsymbol{\vartheta})$. Filtering is based on the idea of obtaining probabilities for the configuration states \mathbf{x}_t when observations up to time t are available, that is $\mathbf{y}_{1:t}$, where backward sampling draws the configuration states of \mathbf{x}_t as if the future states $\mathbf{x}_{t+1:T}$ are considered fixed. According to the modelling described in Section 4.2.2 and the assumptions (A1) and (A2), we are able to obtain the corresponding densities. The filtered density of \mathbf{x}_t can be written in the following way,

$$\begin{aligned}
p(\mathbf{x}_t|\mathbf{y}_{1:t}, \boldsymbol{\vartheta}) &= p(\mathbf{x}_t|\mathbf{y}_t, \mathbf{y}_{1:t-1}, \boldsymbol{\vartheta}) \\
&\propto f(\mathbf{y}_t|\mathbf{y}_{1:t-1}, \mathbf{x}_t, \boldsymbol{\vartheta}) p(\mathbf{x}_t|\mathbf{y}_{1:t-1}, \boldsymbol{\vartheta}) \\
&\stackrel{\text{(A1)}}{\propto} f(\mathbf{y}_t|\mathbf{x}_t, \boldsymbol{\phi}) \underbrace{p(\mathbf{x}_t|\mathbf{y}_{1:t-1}, \boldsymbol{\vartheta})}_{\text{one-step ahead density at } t}, \tag{4.2.50}
\end{aligned}$$

where $f(\mathbf{y}_t|\mathbf{x}_t, \boldsymbol{\phi}) = \prod_{\forall s \in \{1, \dots, N\}} f(y_{t,s}|\mathbf{x}_{t,s}, \boldsymbol{\phi})$ is as in Equation (4.2.17). As can be seen in the last line of the equation, the so called one-step ahead density at time t appears, $p(\mathbf{x}_t|\mathbf{y}_{1:t-1}, \boldsymbol{\vartheta})$. This corresponds to predicting the configuration states of \mathbf{x}_t when observations up to $t-1$ are available. It can be factorised into the product of the marginal density of \mathbf{x}_t and the filtered distribution at $t-1$ as,

$$\begin{aligned}
p(\mathbf{x}_t|\mathbf{y}_{1:t-1}, \boldsymbol{\vartheta}) &= \sum_{\forall \mathbf{x}_{t-1}} p(\mathbf{x}_t, \mathbf{x}_{t-1}|\mathbf{y}_{1:t-1}, \boldsymbol{\vartheta}) \\
&\propto \sum_{\forall \mathbf{x}_{t-1}} p(\mathbf{x}_t|\mathbf{x}_{t-1}, \mathbf{y}_{1:t-1}, \boldsymbol{\vartheta}) p(\mathbf{x}_{t-1}|\mathbf{y}_{1:t-1}, \boldsymbol{\vartheta}) \\
&\stackrel{\text{(A2)}}{\propto} \sum_{\forall \mathbf{x}_{t-1}} p(\mathbf{x}_t|\mathbf{x}_{t-1}, \boldsymbol{\vartheta}) p(\mathbf{x}_{t-1}|\mathbf{y}_{1:t-1}, \boldsymbol{\vartheta}) \\
&\stackrel{\text{(A2)}}{\propto} \sum_{\forall \mathbf{x}_{t-1}} p(\mathbf{x}_t|\mathbf{x}_{t-1}, \boldsymbol{\xi}) p(\mathbf{x}_{t-1}|\partial \mathbf{x}_{t-1}, \boldsymbol{\phi}) \underbrace{p(\mathbf{x}_{t-1}|\mathbf{y}_{1:t-1}, \boldsymbol{\vartheta})}_{\text{filtered density at } t-1}, \tag{4.2.51}
\end{aligned}$$

where $p(\mathbf{x}_{t-1}|\partial\mathbf{x}_{t-1}, \phi)$ is modelled by an autologistic model as in Equation (4.2.10). The first line of the equation states that the marginal density can be obtained by summing out all possible configuration states \mathbf{x}_{t-1} from the joint distribution of \mathbf{x}_t and \mathbf{x}_{t-1} .

As can be seen in Equations (4.2.50) and (4.2.51), updating either the filtered or the one-step ahead probabilities requires knowledge of the corresponding one-step ahead or filtered probabilities respectively, implying a recursive acquisition procedure. The way we perform the filtering, as well the sampling which follows, is based on pixel-wise trajectories over the corresponding time points, as can be graphically seen in Figure 4.2.3.




Figure 4.2.3: Two different pixel's, s_1 and s_2 , trajectories of raw data points over all time points, or frames equivalently, where $t = 1, \dots, T$. These trajectories are denoted with $\mathbf{y}_{1:T,s_1}$ and $\mathbf{y}_{1:T,s_2}$ respectively, and correspond to all the observed intensities for pixel s_1 and s_2 across the stack of frames. Similar trajectories account for their states $\mathbf{x}_{1:T,s_1}$ and $\mathbf{x}_{1:T,s_2}$ (not shown).

This allows for a recursive pixel-wise update implementation of the states in every $\mathbf{x}_t = \{x_{t,s}, s \in \{1, \dots, N\}\}, \forall t \in \{1, \dots, T\}$, therefore we are interested in the univariate probabilities $\Pr(x_{t,s} = k_c | \mathbf{y}_{1:t,s}, \boldsymbol{\vartheta})$ and $\Pr(x_{t,s} = k_c | \mathbf{y}_{1:t-1,s}, \boldsymbol{\vartheta})$. These probabilities denote the filtered and one step ahead probabilities of $x_{t,s}$ being on state k_c , con-

ditioning on a trajectory of observations over t and $t - 1$ time points for pixel s respectively. Due to the discrete domain of possible values for the states, that is 'On' and 'Off' declared with 1 and -1 as in Equation (4.2.4), these probabilities are straightforward to obtain. The one step ahead probability of $x_{t,s}$ given observations up to $t - 1$, $\mathbf{y}_{1:t-1,s}$, reads in,

$$\begin{aligned} \Pr(x_{t,s} = k_c | \mathbf{y}_{1:t-1,s}, \boldsymbol{\vartheta}) &= \sum_{k_p=0}^1 \Pr(x_{t,s} = k_c | x_{t-1,s} = k_p, \boldsymbol{\xi}) \times \\ &\times \Pr(x_{t-1,s} = k_p | \partial x_{t-1,s}, \boldsymbol{\phi}) \Pr(x_{t-1,s} = k_p | \mathbf{y}_{1:t-1,s}, \boldsymbol{\vartheta}), \end{aligned} \quad (4.2.52)$$

while the filtered probability of $x_{t,s}$ given observations up to t , $\mathbf{y}_{1:t-1,s}$,

$$\Pr(x_{t,s} = k_c | \mathbf{y}_{1:t,s}, \boldsymbol{\vartheta}) = \frac{f(y_{t,s} | x_{t,s} = k_c, \boldsymbol{\phi}) \Pr(x_{t,s} = k_c | \mathbf{y}_{1:t-1,s}, \boldsymbol{\vartheta})}{\sum_{k_a=0}^1 f(y_{t,s} | x_{t,s} = k_a, \boldsymbol{\phi}) \Pr(x_{t,s} = k_a | \mathbf{y}_{1:t-1,s}, \boldsymbol{\vartheta})}. \quad (4.2.53)$$

We note here that $p(\mathbf{x}_0 | \boldsymbol{\phi}, \boldsymbol{\xi}) \equiv p(\mathbf{x}_0 | \boldsymbol{\phi})$ independent of the transition matrix $\boldsymbol{\xi}$, as our assumption in Equation (4.2.16). Therefore, we sample the initial field \mathbf{x}_0 taking only into consideration the predetermined neighbourhood structure via the autologistic model. The state of every pixel s in $\mathbf{x}_0 = \{x_{0,s}, s \in \{1, \dots, N\}\}$ is obtained the probability in Equation (4.2.15),

$$\Pr(x_{0,s} = k_p | x_{0,\partial s}, \boldsymbol{\phi}) = \frac{\exp\{\beta_0 k_p + \beta_f \sum_{q \in \partial s} k_p x_{0,q}\}}{\sum_{k_a \in \{-1,1\}} \exp\left\{k_a \left(\beta_0 + \beta_f \sum_{q \in \partial s} x_{0,q}\right)\right\}}. \quad (4.2.54)$$

In order to sample the initial field \mathbf{x}_0 knowledge on the states of the neighbouring pixels is required. Therefore, instead of having a completely random configuration state to apply the autologistic model, we consider an initial classification step based on the first frame of the stack. A quartile of the observed intensities on the first frame sets a binary classification of the pixels, which is then used to sample \mathbf{x}_0 from Equation (4.2.54).

In order to sample the configuration states of $\mathbf{x}_{1:T} = (\mathbf{x}_1, \dots, \mathbf{x}_t, \dots, \mathbf{x}_T)$ we adopt a multi-move sampling scheme similar to Frühwirth-Schnatter [2006]. Multi move sampling implies sampling the configuration states of $\mathbf{x}_{1:T}$ simultaneously from the full conditional distribution of $\mathbf{x}_{1:T}$ given $\mathbf{y}_{1:T}$ and the parameters $\boldsymbol{\vartheta}$. Firstly, we write down this full conditional distribution $p(\mathbf{x}_{1:T}|\mathbf{y}_{1:T}, \boldsymbol{\vartheta})$ as,

$$\begin{aligned} p(\mathbf{x}_{1:T}|\mathbf{y}_{1:T}, \boldsymbol{\vartheta}) &\propto p(\mathbf{x}_1, \dots, \mathbf{x}_{T-1}|\mathbf{x}_T, \mathbf{y}_{1:T}, \boldsymbol{\vartheta}) p(\mathbf{x}_T|\mathbf{y}, \boldsymbol{\vartheta}) \\ &\propto p(\mathbf{x}_1, \dots, \mathbf{x}_{T-2}|\mathbf{x}_{T-1}, \mathbf{x}_T, \mathbf{y}_{1:T}, \boldsymbol{\vartheta}) p(\mathbf{x}_{T-1}|\mathbf{x}_T, \mathbf{y}_{1:T}, \boldsymbol{\vartheta}) p(\mathbf{x}_T|\mathbf{y}_{1:T}, \boldsymbol{\vartheta}) \\ &\propto p(\mathbf{x}_1|\mathbf{x}_2, \dots, \mathbf{x}_T, \mathbf{y}_{1:T}, \boldsymbol{\vartheta}) \dots p(\mathbf{x}_{T-1}|\mathbf{x}_T, \mathbf{y}_{1:T}, \boldsymbol{\vartheta}) p(\mathbf{x}_T|\mathbf{y}_{1:T}, \boldsymbol{\vartheta}) \\ &\propto p(\mathbf{x}_T|\mathbf{y}_{1:T}, \boldsymbol{\vartheta}) \prod_{t=0}^{T-1} p(\mathbf{x}_t|\mathbf{x}_{t+1}, \dots, \mathbf{x}_T, \mathbf{y}_{1:T}, \boldsymbol{\vartheta}), \end{aligned}$$

where $p(\mathbf{x}_t|\mathbf{x}_{t+1}, \dots, \mathbf{x}_T, \mathbf{y}_{1:T}, \boldsymbol{\vartheta})$ is the conditional distribution of the configuration of states \mathbf{x}_t when the future configurations of states $\mathbf{x}_{t+1}, \dots, \mathbf{x}_T$ are considered known, and $p(\mathbf{x}_T|\mathbf{y}_{1:T}, \boldsymbol{\vartheta})$ is the filtered density of \mathbf{x}_T at time T .

Since we perform a recursive update of the configuration states, we are mainly interested in the conditional distribution of the configuration states \mathbf{x}_t given the future configurations of states $\mathbf{x}_{t+1}, \dots, \mathbf{x}_T$. Therefore, the conditional distribution of \mathbf{x}_t given $\mathbf{x}_{t+1}, \dots, \mathbf{x}_T$ can be simplified,

$$\begin{aligned} p(\mathbf{x}_t|\mathbf{x}_{t+1}, \dots, \mathbf{x}_T, \mathbf{y}_{1:T}, \boldsymbol{\vartheta}) &\propto f(\mathbf{y}_{t+1}, \dots, \mathbf{y}_T|\mathbf{x}_t, \dots, \mathbf{x}_T, \mathbf{y}_{1:t}, \boldsymbol{\vartheta}) p(\mathbf{x}_t|\mathbf{x}_t, \dots, \mathbf{x}_T, \mathbf{y}_{1:t}, \boldsymbol{\vartheta}) \\ &\propto p(\mathbf{x}_t|\mathbf{x}_{t+1}, \dots, \mathbf{x}_T, \mathbf{y}_{1:t}, \boldsymbol{\vartheta}) \\ &\propto p(\mathbf{x}_{t+1}, \dots, \mathbf{x}_T|\mathbf{x}_t, \mathbf{y}_{1:t}, \boldsymbol{\vartheta}) p(\mathbf{x}_t|\mathbf{y}_{1:t}, \boldsymbol{\vartheta}) \end{aligned}$$

$$\begin{aligned}
&\propto p(\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{y}_{1:t}, \boldsymbol{\vartheta})p(\mathbf{x}_t|\mathbf{y}_{1:t}, \boldsymbol{\vartheta}) \\
\stackrel{\text{(A2)}}{\propto} &p(\mathbf{x}_{t+1}|\mathbf{x}_t, \boldsymbol{\vartheta})p(\mathbf{x}_t|\mathbf{y}_{1:t}, \boldsymbol{\vartheta}) \\
\stackrel{\text{(A2)}}{\propto} &p(\mathbf{x}_{t+1}|\mathbf{x}_t, \boldsymbol{\xi})p(\mathbf{x}_t|\partial\mathbf{x}_t, \boldsymbol{\phi}) \underbrace{p(\mathbf{x}_t|\mathbf{y}_{1:t}, \boldsymbol{\vartheta})}_{\text{filtered density at } t}
\end{aligned}$$

where in the first line of the equation, $f(\mathbf{y}_{t+1}, \dots, \mathbf{y}_T|\mathbf{x}_t, \mathbf{x}_{t+1}, \dots, \mathbf{x}_T, \mathbf{y}_{1:t}, \boldsymbol{\phi})$ is independent of \mathbf{x}_t when \mathbf{x}_{t+1} is available and considered known. Therefore, if the state of $x_{t+1,s}$ was known and equal to k_f the probability the state of $x_{t,s}$ being equal to k_c given the trajectory of observations up to time t , that is $\mathbf{y}_{1:t,s}$, reads as,

$$\begin{aligned}
\Pr(x_{t,s} = k_c | x_{t+1,s} = k_f, \mathbf{y}_{1:t,s}, \boldsymbol{\vartheta}) = \\
\frac{\Pr(x_{t+1,s} = k_f | x_{t,s} = k_c, \boldsymbol{\vartheta})\Pr(x_{t,s} = k_c | \mathbf{y}_{1:t,s}, \boldsymbol{\vartheta})}{\sum_{k_a \in \{-1, 1\}} \Pr(x_{t+1,s} = k_f | x_{t,s} = k_a, \boldsymbol{\vartheta})\Pr(x_{t,s} = k_a | \mathbf{y}_{1:t,s}, \boldsymbol{\vartheta})}. \quad (4.2.55)
\end{aligned}$$

The following Pseudo-algorithm summarises the steps in order to sample a complete path of $\mathbf{x}_{1:T}$ using the corresponding pixel-wise trajectories over time, using the described FFBS procedure. For the multi-move updating of $\mathbf{x}_{1:T}$ the parameters $\boldsymbol{\vartheta} = \{I, c, \tau_b, \boldsymbol{\beta}_{1:T}, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{1:T}, \boldsymbol{\xi}\}$ are considered known.

Pseudo-algorithm for Multi-move State updating

1. *Step* $t=0$ (*Frame 0*):

Initialise \mathbf{x}_0 from initial distribution $p(\mathbf{x}_0|\boldsymbol{\phi})$.

Every pixel $s \in \{1, \dots, N\}$ can be on the 'on' state (=1) with probability,

$$\Pr(x_{0,s} = 1 | \partial x_{0,s}, \boldsymbol{\phi})$$

as defined in Equation (4.2.54), where the field \mathbf{x}_0 is initially sampled using

a broad classification of the first frame based on a quartile of the data.

2. *Step $t=1$ to T (Frame 1 up to Frame T) :*

For every pixel $s \in \{1, \dots, N\}$ calculate recursively the one step ahead and filtered probabilities of the 'On' state ($=1$) with probabilities

$$\Pr(x_{t,s} = 1 | \mathbf{y}_{1:t-1,s}, \boldsymbol{\vartheta}) \quad \text{and} \quad \Pr(x_{t,s} = 1 | \mathbf{y}_{1:t,s}, \boldsymbol{\vartheta})$$

respectively, from Equations (4.2.51) and (4.2.52) using their corresponding trajectories.

The filtered probabilities need to be saved as they are necessary in order to perform the backward sampling in Step 3 and 4.

3. *Step $t=T$ (Frame T):*

Sample the configuration states \mathbf{x}_T of the last frame by using the filtered probability $\Pr(x_{T,s} = 1 | \mathbf{y}_{1:T,s}, \boldsymbol{\vartheta})$ for every pixel $s \in \{1, \dots, N\}$ obtained and saved in Step 2.

4. *Step $t=T-1$ to 1 (Frame $T-1$ up to Frame 1) :*

For every pixel $s \in \{1, \dots, N\}$ sample a state with probability of being on the 'on' state from

$$\Pr(x_{t,s} = 1 | x_{t+1,s} = k_f, \mathbf{y}_{1:t,s}, \boldsymbol{\vartheta})$$

in Equation (4.2.55) where the state of $x_{t+1,s}$, denoted by k_f in the probability, has been obtained and saved in Step 2 along with the filtered probabilities $\Pr(x_{t,s} = 1 | \mathbf{y}_{1:t,s}, \boldsymbol{\vartheta})$ needed.

4.2.3.3 Pseudo-algorithm for Bayesian inference of a Markov Switching model based on SiMPa functions via a Forward-Filtering-Backward-Sampling (FFBS) algorithm

*Pseudo-algorithm - Bayesian Inference Markov Switching model
based on SiMPa functions*

1. *Step $q=0$ (Iteration 0):*

Initialisation of parameters and fields

$$\begin{aligned} \boldsymbol{\vartheta}^{(0)} &= \left\{ I^{(0)}, c^{(0)}, \tau_b^{(0)}, \boldsymbol{\beta}_{1:T}^{(0)}, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{1:T}^{(0)}, \boldsymbol{\xi}^{(0)} \right\} \\ \text{where } \boldsymbol{\beta}_{1:T}^{(0)} &= \left(\{\beta_{0_1}^{(0)}, \beta_{f_1}^{(0)}\}, \dots, \{\beta_{0_T}^{(0)}, \beta_{f_T}^{(0)}\} \right) \\ \mathbf{x}_{1:T}^{(0)} &= (\mathbf{x}_1, \dots, \mathbf{x}_T) \end{aligned}$$

The parameters in $\boldsymbol{\vartheta}^{(0)}$ can drawn from the corresponding prior distributions, while the configurations states $\mathbf{x}_{1:T}^{(0)}$ from the marginal density of $\mathbf{x}_{1:T}$ in Equation (4.2.18),

$$\mathbf{x}_{1:T}^{(0)} \sim p(\mathbf{x}_{1:T} | \boldsymbol{\vartheta}^{(0)})$$

Without loss of generality we assume that the initial values for the complete localisation sets $\{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{1:T}^{(0)} = \{\mathbf{Conf}, \mathbf{0}, \mathbf{0}\}_{1:T}^{(0)}$, hence every fluorophore is at the origin of the SiMPa functions..

2. *Step $q=1, \dots, IT$ (Iteration 1 up to IT)*

Sequential update of the parameter and states

- 2i. Updating of parameters in $\boldsymbol{\vartheta}^{(q)} = \left\{ I^{(q)}, c^{(q)}, \tau_b^{(q)}, \boldsymbol{\beta}_{1:T}^{(q)}, \boldsymbol{\xi}^{(q)} \right\}$ and probabilistic draw of fluorophores localisations based on SiMPa functions, conditional on knowing $\mathbf{x}^{(q-1)}$.

- (a) Update the transition matrix $\boldsymbol{\xi}^{(q)}$ from its full conditional distribution for each one of the rows $\boldsymbol{\xi}_{j_1, \bullet}^{(q)}$,

$$\boldsymbol{\xi}_{j_1, \bullet}^{(q)} \sim p\left(\boldsymbol{\xi}_{j_1, \bullet} \mid \mathbf{x}_{1:T}^{(q-1)}\right)$$

from Equation(4.2.44)

- (b) Update the parameters $I^{(q)}, \tau_b^{(q)}$ (Gibbs steps), $c^{(q)}, \boldsymbol{\beta}_{1:T}^{(q)}$ (Metropolis-Hastings steps) from their corresponding full conditional distributions using the entire time series from $1 : T$ and the moving regions $\mathbf{MR}_{1:T} = (\mathbf{MR}_1, \dots, \mathbf{MR}_T)$ where $\mathbf{MR}_t = \{\text{MR}_{t,s}, \forall s \in \{1, \dots, N\}\}$, $\forall t = (1, \dots, T)$.

It should be noted that the parameters $I^{(q)}, \tau_b^{(q)}, c^{(q)}$ are global parameters while $\boldsymbol{\beta}_{1:T}^{(q)}$ are frame-wise parameters, therefore

$$I^{(q)} \sim p\left(I \mid \mathbf{x}^{(q-1)}, \mathbf{MR}_{1:T}, \tau_b^{(q-1)}, c^{(q-1)}, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{1:T}^{(q-1)}\right)$$

from Equation(4.2.32)

$$\tau_b^{(q)} \sim p\left(\tau_b \mid \mathbf{x}^{(q-1)}, \mathbf{MR}_{1:T}, I^{(q)}, c^{(q-1)}, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{1:T}^{(q-1)}\right)$$

from Equation(4.2.34)

$$c^{(q)} \sim p\left(c \mid \mathbf{x}^{(q-1)}, \mathbf{MR}_{1:T}, I^{(q)}, \tau_b^{(q)}, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{1:T}^{(q-1)}\right)$$

from Equation(4.2.35)

$$\{\beta_{0_t}, \beta_{f_t}\}^{(q)} \sim p\left(\{\beta_0, \beta_f\} \mid \mathbf{x}_t^{(q-1)}\right)$$

from Equation(4.2.36)

- (c) Apply either the *localisation scheme* or *counting scheme*.

Localisation scheme

For each frame $t \in (1, \dots, T)$ on iteration q , apply the individual frame localisation scheme based on SiMPa functions, based on the most up

to date configuration states of fields $\mathbf{x}_{1:T}^{(q-1)} = (\mathbf{x}_1^{(q-1)}, \dots, \mathbf{x}_T^{(q-1)})$.

For every $\mathbf{MR}_t = \{\text{MR}_{t,s}, s \in \{1, \dots, N\}\}$, evaluate the full conditional distributions in Equation (4.2.37),

$$\begin{aligned} P_{t,s(1)}^{(q)} &= p\left(\{\text{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{t,s(1)} \mid \text{MR}_{t,s}, \partial x_{t,s}^{(q-1)}, I^{(q)}, c^{(q)}, \tau_b^{(q)}\right) \\ &\vdots \\ P_{t,s(u)}^{(q)} &= p\left(\{\text{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{t,s(u)} \mid \text{MR}_{t,s}, \partial x_{t,s}^{(q-1)}, I^{(q)}, c^{(q)}, \tau_b^{(q)}\right) \end{aligned}$$

with localisation probabilities obtained after normalising the densities $P_{t,s(1)}^{(q)}, \dots, P_{t,s(u)}^{(q)}$, that is

$$P_{t,s(1)}^{(q)} \Big/ \sum_{h=1}^u P_{t,s(h)}^{(q)}, \dots, P_{t,s(u)}^{(q)} \Big/ \sum_{h=1}^u P_{t,s(h)}^{(q)}$$

Counting scheme

For each frame $t \in (1, \dots, T)$ on iteration q , apply the individual frame counting scheme based on SiMPa functions, based on the most up to date configuration states of fields $\mathbf{x}_{1:T}^{(q-1)} = (\mathbf{x}_1^{(q-1)}, \dots, \mathbf{x}_T^{(q-1)})$.

For every $\mathbf{MR}_t = \{\text{MR}_{t,s}, s \in \{1, \dots, N\}\}$, evaluate the full conditional distributions in Equation (4.2.38),

$$\begin{aligned} C_{t,s(1)}^{(q)} &= p\left(\{\text{Conf}, \mathbf{r} = \mathbf{0}, \boldsymbol{\theta} = \mathbf{0}\}_{t,s(1)} \mid \text{MR}_{t,s}, \partial x_{t,s}^{(q-1)}, I^{(q)}, c^{(q)}, \tau_b^{(q)}\right) \\ &\vdots \\ C_{t,s(w)}^{(q)} &= p\left(\{\text{Conf}, \mathbf{r} = \mathbf{0}, \boldsymbol{\theta} = \mathbf{0}\}_{t,s(w)} \mid \text{MR}_{t,s}, \partial x_{t,s}^{(q-1)}, I^{(q)}, c^{(q)}, \tau_b^{(q)}\right) \end{aligned}$$

with localisation probabilities obtained after normalising the densities

$C_{t,s(1)}^{(q)}, \dots, C_{t,s(w)}^{(q)}$, that is

$$C_{t,s(1)}^{(q)} \bigg/ \sum_{h=1}^w C_{t,s(h)}^{(q)}, \dots, C_{t,s(w)}^{(q)} \bigg/ \sum_{h=1}^w C_{t,s(h)}^{(q)}$$

- 2ii. Updating of configuration states in $\mathbf{x}_{1:T}^{(q)} = (\mathbf{x}_1^{(q)}, \dots, \mathbf{x}_T^{(q)})$ using the recursive FFBS algorithm (Subsection 4.2.3.2) conditional on knowing all the parameters $\boldsymbol{\vartheta}^{(q)} = \{I^{(q)}, c^{(q)}, \tau_b^{(q)}, \boldsymbol{\beta}_{1:T}^{(q)}, \{\mathbf{Conf}, \mathbf{r}, \boldsymbol{\theta}\}_{1:T}^{(q)}, \boldsymbol{\xi}^{(q)}\}$.

Updating the configuration states $\mathbf{x}_{1:T}^{(q)}$ is based on the raw time series data $\mathbf{y}_{1:T}$ and the appropriate pixel-wise trajectories over time points as presented on Section 4.2.3.

- (a) *Step $t=0$ (Frame 0) in iteration q :*

Initialise $\mathbf{x}_0^{(q)} = \{x_{0,s}^{(q)}, \text{ for every } s \in \{1, \dots, N\}\}$ from initial distribution $p(\mathbf{x}_0 | \boldsymbol{\phi})$. The probability of the 'On' state (=1) can be,

$$\Pr(x_{0,s} = 1 | \partial x_{0,s}^{(q)}, \boldsymbol{\phi})$$

as defined in Equation (4.2.54), where the field \mathbf{x}_0 is initially sampled using a broad classification of the first frame based on a quartile of the data.

- (b) *Forward Step $t=1$ to T (Frame 1 up to Frame T) in iteration q :*

For every $\mathbf{x}_t^{(q)} = \{x_{t,s}^{(q)}, \text{ for every } s \in \{1, \dots, N\}\}$ with $t = 1, \dots, T$ calculate recursively the one step ahead and filtered probabilities of every pixel being on the 'On' state (=1) with probabilities,

$$\Pr(x_{t,s} = 1 | \mathbf{y}_{1:t-1,s}, \boldsymbol{\vartheta}^{(q)}) \quad \text{and} \quad \Pr(x_{t,s} = 1 | \mathbf{y}_{1:t,s}, \boldsymbol{\vartheta}^{(q)})$$

for their corresponding trajectories $\mathbf{y}_{1:t-1,s}$ and $\mathbf{y}_{1:t,s}$ respectively, from

Equations (4.2.51) and (4.2.52).

The filtered probabilities are stored, since they are necessary in order to perform the backward sampling in Step (c) and (d), in the variable

$$\mathbf{FP} = (\mathbf{FP}_1, \dots, \mathbf{FP}_t, \dots, \mathbf{FP}_T),$$

$$\text{where } \mathbf{FP}_t = \{\Pr(x_{t,s} = 1 | \mathbf{y}_{1:t,s}, \boldsymbol{\vartheta}^{(q)}), \forall s \in \{1, \dots, N\}\}$$

(c) *Step $t=T$ (Frame T) in iteration q :*

Sample the configuration states $\mathbf{x}_T^{(q)}$ of the last frame by using the filtered probability $\Pr(x_{T,s} = 1 | \mathbf{y}_{1:T,s}, \boldsymbol{\vartheta}^{(q)})$ for every pixel $s \in \{1, \dots, N\}$ obtained and stored in Step (b).

(d) *Backward Step $t=T-1$ to 1 (Frame $T-1$ up to Frame 1) in iteration q :*

For every $\mathbf{x}_t^{(q)} = \{x_{t,s}^{(q)}, \text{ for every } s \in \{1, \dots, N\}\}$ with $t = T, \dots, 1$ sample the configuration states with pixel-wise probability of being on the 'on' state

$$\Pr(x_{t,s} = 1 | x_{t+1,s}^{(q)} = k_f, \mathbf{y}_{1:t,s}, \boldsymbol{\vartheta})$$

in Equation (4.2.55) for the corresponding trajectory $\mathbf{y}_{1:t,s}$. The state of $x_{t+1,s}$, denoted by k_f in the probability, has been obtained and stored in Step (b), along with the filtered probabilities $\Pr(x_{t,s} = 1 | \mathbf{y}_{1:t,s}, \boldsymbol{\vartheta})$.

3. At the end of the MCMC for every $t = 1, \dots, T$, calculate posterior probabilities for the configurations by applying (Drawn Configurations)/ (Potential Configurations) to obtain pixels including the fluorophores.

If localisation is performed, calculate additional posterior probabilities of realisations within the configurations by applying (Drawn realisations)/ (Potential realisations).

4.3 Sampling scheme for sequence of images using the SiMPa functions

Pseudocode - Generation of stack of frames using the SiMPa functions

- (i) Choose the parameters $n, m, R, l, c, I, d, \tau_b, N^{(0)}$
 n, m : number of rows and columns respectively for each frame
 R : radius of the molecules
 l : length of pixel
 c : diffraction parameter
 I : intensity of a single event
 I_0 : background intensity
 τ_b : background error precision
 $N^{(0)}$: total number of alive molecules

Also define the transition matrix P including the probabilities in which states (1='On', 0='Off', D='Bleached') of molecules between frames are altering, i.e.

$$P = \begin{pmatrix} p_{11} & p_{10} & p_{1D} \\ p_{01} & p_{00} & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

- (ii) Choose an r and θ for every molecule in $N^{(0)}$ and draw a random position for each one. We define the following sets:

$$\mathbf{K}^{(0)} = 1 : N^{(0)} = \{1, 2, \dots, N^{(0)}\} \text{ (set including molecules)}$$

$$\mathbf{r} = \{r_1, r_2, \dots, r_{N^{(0)}}\} \text{ (each molecule's distance)}$$

$$\boldsymbol{\theta} = \{\theta_1, \theta_2, \dots, \theta_{N^{(0)}}\} \text{ (each molecule's angle)}$$

$$\mathbf{u} = \{u_1, u_2, \dots, u_{N^{(0)}}\} \text{ (each molecule's pixel on the lattice)}$$

(iii) Choose a number of 'On' molecules denoted by $n_{11}^{(0)}$. We use the following notation:

$$N_a^{(0)} = n_{11}^{(0)} : \text{ number of active molecules}$$

$$N_{da}^{(0)} = N^{(0)} - N_a^{(0)} : \text{ number of de-active molecules}$$

$$N_D^{(0)} = 0 : \text{ number of bleached molecules}$$

$$\mathbf{K}_a^{(0)} : \text{ Choose } N_a^{(0)} \text{ molecules from } \mathbf{K}^{(0)} \text{ to be active}$$

$$\mathbf{K}_{da}^{(0)} : \text{ Choose } N^{(0)} - N_a^{(0)} \text{ molecules from } \mathbf{K}^{(0)} \text{ to be de-active}$$

$$\mathbf{K}_D^{(0)} = \emptyset : \text{ Choose 'Bleached' molecules}$$

(iv) Set frame $f = 1$ and denote by:

$n_{ij}^{(f)}$: the number of molecules moving from state i to state j on frame f .

Transition of every molecule and repetition until $N_D^{(f)} = N^{(0)}$:

$$\mathbf{K}_a^{(f-1)} = \begin{bmatrix} 1 \\ 2 \\ \vdots \\ N_a^{(f-1)} \end{bmatrix} \xrightarrow{T_a^{(f)} \sim \text{Multinomial}(1|p_{11}, p_{10}, p_{1D})} \mathbf{T}_a^{(f)} = \begin{matrix} \mathbf{11} & \mathbf{10} & \mathbf{1D} \\ \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ & \ddots & \\ 0 & 0 & 1 \end{bmatrix} \end{matrix}$$

$$\mathbf{K}_{da}^{(f-1)} = \begin{bmatrix} 1 \\ 2 \\ \vdots \\ N_{da}^{(f-1)} \end{bmatrix} \xrightarrow{T_{da}^{(f)} \sim \text{Multinomial}(1|p_{01}, p_{00})} \mathbf{T}_{da}^{(f)} = \begin{matrix} \mathbf{01} & \mathbf{00} \\ \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ \ddots & \\ 1 & 0 \end{bmatrix} \end{matrix}$$

Then the molecules on each state can be obtained by:

$$\begin{aligned}\mathbf{K}_a^{(f)} &= \left\{ \mathbf{K}_a^{(f-1)} [\mathbf{T}_a^{(f)} [\mathbf{11}] = 1] \right\} \cup \left\{ \mathbf{K}_{da}^{(f-1)} [\mathbf{T}_{da}^{(f)} [\mathbf{01}] = 1] \right\} \\ \mathbf{K}_{da}^{(f)} &= \left\{ \mathbf{K}_a^{(f-1)} [\mathbf{T}_a^{(f)} [\mathbf{10}] = 1] \right\} \cup \left\{ \mathbf{K}_{da}^{(f-1)} [\mathbf{T}_{da}^{(f)} [\mathbf{00}] = 1] \right\} \\ \mathbf{K}_D^{(f)} &= \left\{ \mathbf{K}_a^{(f-1)} [\mathbf{T}_a^{(f)} [\mathbf{1D}] = 1] \right\} \cup \left\{ \mathbf{K}_D^{(f-1)} \right\}\end{aligned}$$

and the total number of them on each state at frame f can be determined by the sums of the respective columns of the transition matrices $\mathbf{T}_a^{(f)}$ and $\mathbf{T}_{da}^{(f)}$. Therefore,

$$\begin{aligned}N^{(f)} &= N^{(f-1)} - n_{1D}^{(f)} \\ N_a^{(f)} &= n_{11}^{(f)} + n_{01}^{(f)} \\ N_{da}^{(f)} &= N^{(f)} - N_a^{(f)} \\ N_D^{(f)} &= N_D^{(f-1)} + n_{1D}^{(f)}\end{aligned}$$

(v) For every frame f :

1. Create a lattice of size $n \times m$ with a baseline intensity I_0 .
2. Calculate the quantities $\left(1 - \alpha_k \left(\mathbf{r}[\mathbf{K}_a^{(f)}], \boldsymbol{\theta}[\mathbf{K}_a^{(f)}]\right)\right) I$, $k = 0, \dots, 8$ for every molecule in $\mathbf{K}_a^{(f)}$.
3. Place the PSF's, based on SiMPa functions, of the $\mathbf{K}_a^{(f)}$ molecules on the lattice, with the center pixels to be defined by $\mathbf{u}[\mathbf{K}_a^{(f)}]$.
5. Add independent and identically distributed background error with precision τ_b to every pixel on the lattice.

4.4 Implementation

In this section we implement the Markov switching model based on SiMPa functions using a toy example data set. Using the algorithm in Section 4.3 we are able to obtain a sequence of frames which are time related according to a transition matrix P . This transition matrix accounts for the state transition of fluorophores across the sequence, with their diffraction described by the SiMPa functions. For this implementation, we consider a subset of the stack of frames constructed for our application on synthetic data in Chapter 5. Briefly, this is derived from an underlying structure of a circle within a circle with a total of $F = 4730$ low resolution frames constituting the entire stack. The low resolution frames are constructed according to the SiMPa functions using the single event intensity $I = 8$ with a baseline intensity $I_0 = 2.4$, the power of spread parameter $c = 0.4$ and the background error precision $\tau_b = 10$ (see Section 5.2). The chosen subset of the stack of frames consists of $T = 9$ frames, corresponding to f_{401}, \dots, f_{409} , and is displayed in Figure 4.4.1.

We apply the Markov switching model based on SiMPa functions for $IT = 25.000$ iterations after a burn-in period of 10.000 iterations, following the pseudo-algorithm in Section 4.2.3.3. The posterior distributions of the parameters are obtained via the appropriate Gibbs and Metropolis-Hastings steps, while the localisation update relies on our novel pattern-configuration-realisation probabilistic scheme. Here, we consider the probabilistic counting scheme where every molecule is assumed to lie at the origin of the SiMPa functions, hence $r = \theta = 0$. The prior distributions for I , τ_b , c , β_{0_t} and β_{f_t} for $t = 401, \dots, 409$ are identical to the ones considered for the individual frame implementation in Equation (3.4.1), and additionally independent Dirichlet prior distributions for each row of the pixel state transition matrix ξ ,

$$\begin{aligned} \pi(\xi_{1,\bullet} | e_{1,1}, e_{1,0}) &\sim \text{Dirichlet}(\xi_{1,\bullet} | e_{1,1} = 30, e_{1,0} = 50), \\ \pi(\xi_{0,\bullet} | e_{0,1}, e_{0,0}) &\sim \text{Dirichlet}(\xi_{0,\bullet} | e_{0,1} = 30, e_{0,0} = 50), \end{aligned}$$

according to Equation (4.2.43). Because the transition matrix ξ consists of $K = 2$




Figure 4.4.1: Subset of consecutive frames from the entire stack generated for application on 'Circle within circle' synthetic data. The light blue crosses represent the true positions of the molecules, with their square pixel regions to be denoting the intensities diffraction according to the SiMPa functions.

states, that is 'On'(1) and 'Off'(0), we can equivalently have independent Beta prior

distributions,

$$\begin{aligned}\pi(\boldsymbol{\xi}_{11} | e_{1,1}, e_{1,0}) &\sim \text{Beta}(\boldsymbol{\xi}_{11} | e_{1,1} = 30, e_{1,0} = 50), \\ \pi(\boldsymbol{\xi}_{00} | e_{0,1}, e_{0,0}) &\sim \text{Beta}(\boldsymbol{\xi}_{00} | e_{0,1} = 30, e_{0,0} = 50).\end{aligned}$$

Similar to the implementation for an individual frame (Section 3.4), the output of the algorithm includes the posterior distribution of the parameters I , τ_b , c , $\boldsymbol{\xi}$ and $\boldsymbol{\beta}$, and posterior probability heatmaps associated with the localisation of the molecules. In Figure 4.4.2, we present the corresponding prior and posterior distributions for I , τ_b and c , as well as ξ_{11} and ξ_{00} . The posterior distribution for the frame wise parameters $\{\beta_{0_t}, \beta_{f_t}\}$ can be found in Appendix A.1.




Figure 4.4.2: Prior and posterior distributions for parameters in the synthetic stack implementation. The red solid lines denote the prior distributions for each parameter while the histograms associated with the blue solid lines denote the posterior distributions of each parameter respectively.

Regarding the localisation of molecules, here corresponding to identification of pixels that contain active molecules (probabilistic counting scheme), the posterior probability heatmaps are shown in Figures 4.4.3 and 4.4.4. The former contains the probabilities of pixels containing active molecules constructed based on the ratio of the counters 'Positions' and 'Visits', while the latter the probabilities of pixels including two active molecules, as a result of the ratio of 'Double events' counter and 'Positions'.

As a reminder from Section 3.4, the 'Positions' counter stores the drawn configurations for every MR, consequently the drawn pixels containing the active molecules, while at the same time all the corresponding potential configurations for the MR are stored in the 'Visits' counter. Cases where the drawn configurations consist of PSFs that completely overlap, corresponding to two active molecules on the same pixel, are stored in the 'Double events' counter. All these counters are updated for every MR during all MCMC iterations, providing posterior heatmaps when considering the appropriate ratios. Due to the probabilistic nature of our modelling to draw positions of molecules, cases can exist where MRs identify pixels containing events outside of the determined 'On' pixels by the field. Since this cannot be true we filter our probability maps using the average field at the end of the MCMC. The average field corresponds to the sum of the fields obtained on every iteration divided by the number of iterations, presented in Figure 4.4.5.

Considering the posterior probability heatmap in Figure 3.4, we can analyse the capability and power of our probabilistic scheme to resolve partially overlapping PSFs, as well as identifying pixels that contain two simultaneously active molecules. The former can be seen in multiple cases within different frames, for instance on Frames 403 (bottom left corner) and 404 (top left corner), while the latter is apparent on Frame 405 in Figure 4.4.4, where the probabilities at the region on the right bottom corner are quite high. A more detailed discussion about the probabilistic inference is presented in Chapter 5 where we perform the application on synthetic data and conduct a sensitivity analysis.




Figure 4.4.3: Posterior probability heatmaps filtered by the average field for frames in Figure 4.4.1. Higher probabilities of pixels containing molecules are declared with white color levels, as shown in the light blue probability scale bar on each probability map. The white bins represent the density around the regions.




Figure 4.4.4: Double event posterior probability heatmaps filtered by the average field for frames in Figure 4.4.1. Higher probabilities are declared with white color levels as shown in the light blue probability scale bar on each probability map. The white bins represent the density around the regions.




Figure 4.4.5: Average field per frame in Figure 4.4.1. The larger proportions a pixel is identified as 'On' are declared with the light blue color levels as shown in the scale bar on each frame. Each figure represents the average state of each frame after the end of the MCMC.

Now, we focus on inference from the Markov switching model based on SiMPa functions, accounting for simultaneous update of the states of the time related frames within the stack. The configuration states $\mathbf{x}_{401:409}$ are obtained using the FFBS algorithm from Section 4.2.3.3 and we are able to calculate the probabilities of a state, say 'On', having a specific length or duration d_{On} using Equation (4.2.48). Due to the high complexity of this probability, we decide to use an approximation based on the average fields. More specifically, at the end of the MCMC we construct a binary mask of the average fields in Figure 4.4.5 using an arbitrary threshold of 0.75. Practically, this means we consider the pixels that were identified active 75% during the MCMC, therefore the binary mask of field t reads in $Q_t = \left\{ \sum_{q=1}^{IT} x_t^{(q)} / IT > 0.75 \right\}$, $\forall t \in \{401, \dots, 409\}$. Then, we approximate $\Pr(D_{\text{On}} = d_{\text{On}} \mid \mathbf{x}_{1:T}, \boldsymbol{\vartheta})$ in Equation (4.2.48) with,

$$\Pr(D_{\text{On}} = d_{\text{On}} \mid \mathbf{x}_{1:T}, \boldsymbol{\vartheta}) \approx \frac{\sum_{s=1}^N \#\{Q_{1:T,s:(s+d_{\text{On}})} = 1, Q_{1:T,s+(d_{\text{On}}+1)} = 0\}}{\sum_{d_a} \sum_{s=1}^N \#\{Q_{1:T,s:(s+d_a)} = 1, Q_{1:T,s+(d_a+1)} = 0\}}$$

where $Q_t = \{Q_{t,s}, \forall s \in \{1, \dots, N\}\}$ and $Q_{1:T,s}$ denotes pixel's s binary mask trajectory. Here, we considered the frames f_{401}, \dots, f_{409} therefore $1 : T = 401 : 409$. On a similar manner we approximate the probability $\Pr(D_{\text{Off}} = d_{\text{Off}} \mid \mathbf{x}_{1:T}, \boldsymbol{\vartheta})$ by using sub-sequences within the trajectories of pixels that were identified 'On' and then altered their state to 'Off', therefore,

$$\Pr(D_{\text{Off}} = d_{\text{Off}} \mid -) \approx \frac{\sum_{s=1}^N \#\{Q_{1:T,s+(d_{\text{Off}}-1)} = 1, Q_{1:T,s:(s+d_{\text{Off}})} = 0, Q_{1:T,s+(d_{\text{Off}}+1)} = 1\}}{\sum_{d_a} \sum_{s=1}^N \#\{Q_{1:T,s+(d_a-1)} = 1, Q_{1:T,s:(s+d_a)} = 0, Q_{1:T,s+(d_a+1)} = 1\}}$$

We consider the maximum duration of either states to be the same as the number of frames in the stack, therefore we can calculate the expected duration of the states

using the average fields in Figure 4.4.5 and Equation (4.2.49) as,

$$\begin{aligned} E(D_{\text{On}}) &= \sum_{d_a=1}^9 d_a \times \text{Pr}(D_{\text{On}} = d_a \mid \mathbf{x}_{401:409}, \boldsymbol{\vartheta}) \approx 1.20, \\ E(D_{\text{Off}}) &= \sum_{d_a=1}^9 d_a \times \text{Pr}(D_{\text{Off}} = d_a \mid \mathbf{x}_{401:409}, \boldsymbol{\vartheta}) \approx 3.16. \end{aligned} \quad (4.4.1)$$

Since the trajectories of pixels are measured on the discrete domain, corresponding to the number of frames within the stack, we round these to $E(D_{\text{On}}) = 1$ and $E(D_{\text{Off}}) = 3$ which describe the expected duration a pixel remains on the 'On' and 'Off' state respectively. We use these expected durations in order to distinguish between different and re-activated molecules within the stack after deciding on the pixels containing the active molecules. Therefore, using the probability maps in Figures 4.4.3 and 4.4.4, we can obtain individual reconstructions of the corresponding underlying structures on each frame using an arbitrary common threshold, presented in Figure 4.4.6. The threshold, here chosen to be 0.50, determines the pixels within each frame that are most probable to contain a molecule, hence can provide a reconstruction of the true positions of the molecules when it is assumed that $r = \theta = 0$. The light blue crosses on the figure represent the true positions of the molecules, where the red coloured circles the drawn ones based on the threshold. A extended discussion on thresholds for the posterior probability maps is presented later in Section 5.3.1.

A final reconstruction of the underlying structure can be obtained by plotting all the individual frame reconstructions, however under the Markov switching model based on SiMPa functions we can use the expected durations of the states to capture and identify reappearances of molecules. The identified pixels containing the active molecules are associated with pixels on the 'On' state, hence using Equations (4.4.1) we can evaluate their corresponding state trajectories. This can provide an overall counting of molecules throughout the stack of frames for the underlying structure.




Figure 4.4.6: Individual frame reconstructions for frames in Figure 4.4.1 using a threshold value of 0.50. The red circles denote the drawn positions by applying the threshold on the posterior probabilities filtered by the average field, whereas the light blue coloured points represent the true positions of the molecules.

In Table 4.4.1 we present the performance of counting if different rounding choices were made. The table consists of four different combinations of the duration of the states, based on alternating rounding of the determined durations in Equation (4.4.1), and declares the percentages of the total correct and wrong predictions. A correct identification of a pixel containing an active molecule is declared as a True Positive (TP), whereas an incorrect identification as a False Positive (FP). Similarly, True Negative (TN) corresponds to correct identification of an empty pixel and False Negative (FN) incorrect declaration of a pixel as empty when an active molecule in fact lies onto.

Counting	TP	FP	TN	FN
$E(D_{\text{ON}}) = 1$ $E(D_{\text{OFF}}) = 3$	100 %	0.3%	99.9%	0%
$E(D_{\text{ON}}) = 2$ $E(D_{\text{OFF}}) = 3$	98.1%	0.1%	99.9%	1.9%
$E(D_{\text{ON}}) = 1$ $E(D_{\text{OFF}}) = 4$	98.1%	0.3%	99.9%	1.9%
$E(D_{\text{ON}}) = 2$ $E(D_{\text{OFF}}) = 4$	96.3%	0.1%	99.9%	3.7%
Real	54	-	1558	-

Table 4.4.1: Counting of molecules using different expected duration $E(D)$ of 'On' and 'Off' states. The table outlines the percentages of correct identified pixels containing the active molecules (True Positive - TP) and the number of correct identified empty pixels (True Negative - TN), along with the incorrect predictions for active pixels (False Positive - FP) and de-active pixels (False Negative - FN). The real number of pixels containing active molecules and empty pixels are also outlined (Real).

Chapter 5

Application and sensitivity analysis

5.1 Introduction

In this chapter we apply the individual frame counting scheme based on the SiMPa functions on synthetic data and compare with one of the most popular off-the-self alternatives. The application consists of a complete analysis on a stack of frames produced by an underlying fixed structure using both our proposed individual frame scheme and ThunderSTORM [Ovesný et al., 2014]. Additionally, we conduct a sensitivity analysis on the performance of our model, on both sets of parameters involved in the localisation/counting schemes and the Markov switching model based on SiMPa functions. The choice not to consider the latter for the application is directly related to computational times and complexity, a matter we discuss on Chapter 6, with the individual frame inference easily set up for parallel computing. Lastly, we apply our individual frame counting scheme on a realistic data set of Tubulins, used for the challenge in 2013 [Sage et al., 2015] to evaluate current super resolution localisation methods.

In Section 5.2, we construct a synthetic data set constituting of a stack of low resolution frames under a predetermined underlying structure. We explain in detail

the parameter selection, followed by an elicitation procedure in order to obtain prior distributions for the single event intensity I and background error precision τ_b . In Section 5.3, we perform a complete analysis of the individual frame counting scheme based on the SiMPa functions using the generated synthetic data set (Section 5.3.1), before comparing and evaluating its performance against ThunderSTORM [Ovesný et al., 2014] (Section 5.3.2). We present visual reconstructions of the underlying reality as well as quantitative evaluation of the output for both methods. We conclude this chapter with a sensitivity analysis in Section 5.4 split into two parts. The first part corresponds to the assessment of our localisation output when different overlapping scenarios are present, as well as investigating the effect of different noise levels (Section 5.4.1). We consider a variety of prior distributions for every parameter in order to assess the robustness of both the individual inference and Markov switching models based on SiMPa functions, forming the second part of the analysis (Section 5.4.2). The chapter concludes with an application on a realistic data set (Section 5.5).

5.2 Synthetic data generation

In this section, we describe the data generation mechanism using the SiMPa functions according to the pseudo-algorithm introduced in Section 4.3, and explain the parameter selection for the synthetic data. We also present a simple mechanism using a number of frames from the stack in order to elicit prior distributions for the single event intensity I and background error precision τ_b .

We decide to construct frames of size $n = 40$ rows and $m = 40$ columns, resulting in a total of $N = 1600$ pixels. The side length of the pixel is arbitrary chosen to be $l = 1$, while each molecule is a point of a circular structure with radius $R = 0$. In order to obtain the positions of the molecules within the image, we create a predetermined structure of two concentric circles, with their radius having a small scale difference, and we name 'Circle within circle' synthetic data. The molecules are placed on the

circumference of each circle, each one constructed as a set of appropriate directions θ from the corresponding centres based on a narrow discretisation of all possible angles. This allows molecules to exist in close distances, with their actual positions described as continuous measurements on the Cartesian coordinate system. The true underlying structure of the 'Circle within circle' synthetic data is displayed in Figure 5.2.1, with the light blue crosses representing the exact positions of the molecules within the image.

In order to generate the low resolution stack of frames using the SiMPa functions, the Cartesian coordinates of every molecule are transformed onto polar coordinates according to the lattice pixels they belong to. Consequently, each molecule is associated with a corresponding position within the image based on a distance r with direction θ with respect to the center of the pixel it lies into, with these centres corresponding to the origin of the SiMPa functions. The starting frame, or frame zero, is arbitrarily chosen to contain 5% of the total number of molecules existing in this structure. Following similar notation to the pseudo-algorithm in Section 4.3, the total number of molecules is $N^{(0)} = 602$ with $N_a^{(0)} = n_{11}^{(0)} = 0.05 \times 602 = 30$ and $N_{da}^{(0)} = 572$. Consulting Dempsey et al. [2011] about different duty cycles of fluorescent dyes, we decide on a transition matrix between 'On', 'Off' and 'Bleached' state to consist of,

$$P = \begin{array}{c} \text{On} \\ \text{Off} \\ \text{Bleached} \end{array} \begin{array}{ccc} \text{On} & \text{Off} & \text{Bleached} \\ \left(\begin{array}{ccc} 0.049 & 0.851 & 0.1 \\ 0.015 & 0.985 & 0 \\ 0 & 0 & 1 \end{array} \right) \end{array}$$

responsible for driving each molecule's state in consecutive frames. The transition matrix P promotes de-active ('Off') molecules to retain their state with a transition probability $p_{\text{Off,Off}} = 0.985$ and active ('On') molecules to alter their state with probability $p_{\text{On,Off}} = 0.851$.

In order to construct the PSF for every molecule using the SiMPa functions, apart from their predetermined position using the 'Circle within circle' structure, denoted by $\{r_q, \theta_q\}$, $\forall q \in \{1, \dots, N^{(0)}\}$ within the corresponding pixels, the single event intensity I , the power of spread parameter c and the background error τ_b need to be selected. Therefore, for the single event intensity I we choose $I = 8$ with the background baseline being $d = 30\%$ of I , hence $I_0 = 2.4$. These choices were chosen arbitrary but in their difference corresponds to clear separation between intensity of signal and background. Since c is bounded (see Section 3.2.4), we consider $c = 0.4$ as an intermediate value and regular background error levels with precision $\tau_b = 10$, or equivalently variance $\sigma_b^2 = 1/\tau_b = 0.1$, corresponding to noise that noticeably alters the intensities of pixels however not on a completely severe manner (see Section 5.4.1 for effect of different noise levels). Adopting these parameter values, a sequence of a total number of $F = 4730$ frames is obtained with the first $F_t, \forall t \in \{1, \dots, 1500\}$ frames having the majority of molecules blinking, whereas the remaining ones are mostly described by sparsity before every molecule reaches the 'Bleached' state. The number of bleached molecules at frame F_{1500} is 545 out of the total $N^{(0)} = 602$ molecules. A few randomly selected low resolution frames from the generated sequence are displayed in Figure 5.2.1. The light blue crosses denote the true positions of the molecules with the square regions representing the spread according to the SiMPa functions.

Recalling the discussion in Section 2.2.3, high density data involve cases where multiple fluorophores are captured on individual frames, translating into light emitters that can potentially coincide with each other, where low density data consist of mostly sparse frames which can achieve temporal separation of fluorophores. The former is usually described by higher signal-to-noise ratio (SNR) while the latter can suffer from potentially challenging lower SNR. According to the data generation mechanism based on the SiMPa functions, different choices on the parameters can result in low resolution frames, promoting overlapping fluorophores or well separated ones with either high or low signal to noise ratio. The latter depends on the




Figure 5.2.1: Circle within circle - synthetic data using the SiMPa functions. On the left top corner of the figure, the true underlying structure of molecules forming a circle within a circle is presented. The following figures are randomly selected frames from the generated stack using pseudo-algorithm in Section (4.3). The square regions on each frame denote the spread according to SiMPa functions while the light blue crosses the true positions of the molecules.

adjustment of the contrast between the single event intensity I and background baseline I_0 in association with appropriate noise levels, with the former corresponding to altering the transition matrix and/or the initial 'On' molecules. For instance, increasing the probability of a molecule remaining on the 'On' state or the probability of transitioning from 'Off' to 'On' can allow an increased number of overlapping cases. In a way, we can arguably construct fairly representative high or low density data by varying the choice of the parameters.

In order to implement the individual frame inference scheme based on SiMPa functions, prior distributions need to be selected accounting for the initial uncertainty on each parameter. Regarding the power of diffraction c , we use a rescaled beta prior with higher density on values away from zero (in which the diffraction decays slower), $\pi(c) \sim \text{RescaledBeta}(4, 1.75)$, whereas for the field parameters we assign independent Gaussian distributions centred at zero with precision 9, $\pi(\beta_0), \pi(\beta_f) \sim \text{N}(0, 9)$. These prior distributions are not elicited but chosen to be sensible, presented on the right panel of the Figure 5.2.3, with values generally between -2 and 2 to serve the purpose regarding β_0 & β_f .

For the single event intensity I and background error precision τ_b , we perform a prior elicitation process to obtain the corresponding hyper-parameters for their Gaussian and Gamma prior distributions respectively. We decide to follow this procedure as the single event intensity and background error are directly related with the measurements on the frames and play a crucial role on the localisation. During a real STORM imaging experiment, prior distributions for the single event intensity I can be arguably recovered or fixed by the user, however we present a mechanism for a complete data set when potentially no prior knowledge is available. For the prior elicitation procedure, we consider the first 100 frames out of the total $F = 4370$ and split each one of them into three separate sets using different quantiles. The lower quantile q_{low} consists of the 98% of the observations forming the hypothetical background intensities, while the middle quantile q_{inter} between 98% and 99% accounts for the single event intensities and the upper quantile q_{upr} above 99% for overlapping

cases. These values can vary depending on the nature of the data and choices might be considered in a sophisticated fashion. In a way, lower values can be considered for data sets consisting of a large amount of overlapping cases or high density data, while potentially higher values for low density data. Gathering the corresponding observations from each partition we obtain Q_{lwr} , Q_{inter} and Q_{upr} , presented in Figure 5.2.2 with the red solid vertical lines being their 10% and 90% quantiles respectively, which we denote with Q_{lwr}^* , Q_{inter}^* and Q_{upr}^* . The reason for considering a within subset of the quantiles is to try and exclude potential extremes on average.




Figure 5.2.2: Quantities used for prior elicitation. Different measurements obtained by using three different quantiles on each one of the first 100 frames of the 'Circle within circle' synthetic data. The top histogram consists of the larger values, hence is assumed to contain the overlapping scenarios, whereas the middle and bottom one consist of the single event case and background baseline respectively. The vertical red solid lines denote the 10% and 90% quantiles of the corresponding sets.

Then,

$$\begin{aligned}\pi(I) &\sim N\left(\mu_I = \text{mean}(Q_{\text{inter}}^*), \tau_I = 1/\text{var}(Q_{\text{inter}}^*)\right) = N(7.32, 2.83), \\ \pi(\tau_b) &\sim \text{Gamma}\left(\text{shape}_{\tau_b} = 2, \text{rate}_{\tau_b} = 4\text{var}(Q_{\text{lwr}}^*)\right) = \text{Gamma}(2, 0.33),\end{aligned}$$

with the prior elicitation procedure summarised in the following Pseudo-algorithm.

Pseudo-algorithm: Prior elicitation for I and τ_b

1. Choose F frames to be used in the elicitation procedure, where \mathbf{y}_f are the observations for frame $f \in F$.
2. Select q_{lwr} , q_{inter} and q_{upr} corresponding to quantiles for the hypothetical background, single and overlapping intensities respectively.
3. For every $f \in F$ obtain the quantities:

$$\left\{ \begin{array}{l} Q_{\text{lwr}_f} = \{\mathbf{y}_f \leq q_{\text{lwr}}\} \\ Q_{\text{inter}_f} = \{\mathbf{y}_f > q_{\text{lwr}} \ \& \ \mathbf{y}_f \leq q_{\text{inter}}\} \\ Q_{\text{upr}_f} = \{\mathbf{y}_f > q_{\text{inter}} \ \& \ \mathbf{y}_f \leq q_{\text{upr}}\} \end{array} \right\} \rightarrow \left\{ \begin{array}{l} Q_{\text{lwr}} = \{Q_{\text{lwr}_f}, \forall f \in F\} \\ Q_{\text{inter}} = \{Q_{\text{inter}_f}, \forall f \in F\} \\ Q_{\text{upr}} = \{Q_{\text{upr}_f}, \forall f \in F\} \end{array} \right\}$$

4. Select common quantiles m_l and m_r to keep:

$$\begin{aligned}Q_{\text{lwr}}^* &= \{Q_{\text{lwr}} \geq m_l \ \& \ Q_{\text{lwr}} \leq m_r\} \\ Q_{\text{inter}}^* &= \{Q_{\text{inter}} \geq m_l \ \& \ Q_{\text{lwr}} \leq m_r\} \\ Q_{\text{upr}}^* &= \{Q_{\text{lwr}} \geq m_l \ \& \ Q_{\text{lwr}} \leq m_r\}\end{aligned}$$

5. Obtain the $\text{mean}(Q_{\text{inter}}^*)$, $\text{var}(Q_{\text{inter}}^*)$ and $\text{var}(Q_{\text{lwr}}^*)$. For the prior distribu-

tion of I , set,

$$\pi(I) \sim N\left(\mu_I = \text{mean}(Q_{\text{inter}}^*), \tau_I = 1/\text{var}(Q_{\text{inter}}^*)\right)$$

whereas for the background error set,

$$\pi(\tau_b) \sim \text{Gamma}\left(\text{shape}_{\tau_b} = 2, \text{rate}_{\tau_b} = 4\text{var}(Q_{\text{lwr}}^*)\right)$$

where we consider four times the variance $\text{var}(Q_{\text{lwr}}^*)$ to allow for higher uncertainty on τ_b and fixed shape at 2 to move away from 0.




Figure 5.2.3: Priors elicited from the 'Circle within circle' synthetic data. On the top left corner of the figure, the prior of the single event intensity I is presented with I_0 being the background intensity, whereas on the bottom left corner the prior of the background precision τ_b is shown. The two remaining chosen priors for the spread c and field control parameters β 's are indicated on the right panel.

The prior distributions elicited and selected are summarised in Figure 5.2.3, with the blue coloured crosses denoting the means of the prior distributions. We adopt these prior distribution for the application on the 'Circle within circle' synthetic data using the individual frame inference based on the SiMPa functions which follows.

5.3 Application and comparisons on synthetic data

5.3.1 Application using individual frame inference based on SiMPa functions

In this section we present the application on the synthetic data of 'Circle within a circle' as described in Section 5.2, using the individual frame inference based on SiMPa functions. From the total number of $F = 4730$ frames in the generated stack we consider the first F_t frames, $t \in \{1, \dots, 100\}$. The main reason behind this choice is that most of the blinking processes occur during the first 1500 frames, with the first 100 frames to be involving most, if not all, of the challenging PSF overlapping scenarios. The remaining frames are generally sparse, where only a few, mostly well separated, molecules are present before all of them are bleached.

In Figure 5.3.1 some of the F_t , $t \in \{1, \dots, 100\}$, frames are displayed, which are part of the application and a complete analysis using the individual frame inference based on SiMPa functions is presented. On each frame, the light blue crosses denote the true positions of the molecules and the surrounding 3×3 pixel regions their intensity diffraction based on the SiMPa functions. The choice of this set of frames is to present a variety of cases, including some challenging scenarios as well less demanding ones. As can be seen in Figure 5.3.1, each one the frames contain a type of an overlapping situation, in which differently positioned molecules within their corresponding pixels have their PSF's overlapping on a number of pixels. On Frame 11, for instance, two molecules share the same pixel resulting in a complete overlapping of their PSFs (middle left side of the frame), with molecules having their PSFs overlapping on six




Figure 5.3.1: Selected frames from the on 'Circle within circle' synthetic data. The frames represent a variety of both overlapping and well isolated PSFs scenarios within the generated stack, and also Frames 3, 22, 30 and 51 contain violations of the main assumption of no more than two PSFs allowed to overlap. The light blue crosses denote the true positions of the molecules with the surrounding 3×3 pixel regions their corresponding intensity diffraction based on the SiMPa functions.

and four pixels also identified (right bottom side of the frame). Additionally, it is apparent that in most of the frames well isolated PSFs can be spotted, where the distance of activate molecules is larger making their diffraction distinguishable from other. All these regions represent scenarios where at most two molecules are extremely close to one another, however this is not necessarily always the case.

High intensity regions can be formed when multiple molecules are simultaneously activated close to each other. Such cases can lead to extremely difficult situations where most of the PSFs in the region share a small number of pixels, making it incredibly challenging to distinguish the molecules positions. One of these cases is illustrated on the top middle part of Frame 3, where four molecules are simultaneously active within a small area producing a high intensity region with multiple overlapping PSFs. As a reminder, our proposed individual frame inference using the SiMPa functions relies on the assumption of having up to two PSFs either partly or completely overlapping on a single region, making the latter a big violation of the assumption. In a similar manner, small violations of the assumption can occur within a frame, for instance when an overlapping case between two PSFs is accompanied with a third PSF interfering with one of the other two PSFs (Frame 30). This can potentially result in missing events or even wrong and less reliable identification of molecules, cases we investigate thoroughly in Section 5.4, however it does not prevent us from conducting inference on such regions. It should be noted that similar high intensity regions can be formed without having our assumption violated, which can result in high precision inference, with an example shown on the left bottom part of Frame 77.

Depending on the signal-to-noise ratio, the performance of our proposed model can be affected, thus a sensitivity analysis is performed on Section 5.4. For this application we consider the so-called regular noise level as presented in Section 5.2 and apply the individual frame inference based on SiMPa functions using the pseudo-algorithm in Section 3.3.5. We focus on the individual frame counting scheme based on the SiMPa functions, the simplification of our full model which assumes $r = \theta = 0$ for

every molecule within a frame. The main practical reason for this choice is the gain in computational time compared to implementing the localisation scheme. Additionally, our counting scheme based on SiMPa functions aims to tackle one of the main issues in quantitative STORM imaging, constituting of identifying the number of molecules within a cell.

Similar to the implementation presented in Section 3.4, the main output of the procedure for an individual frame involves a posterior probability map of the frame, where higher probabilities indicate the pixels where active molecules more likely lie onto. The posterior probability map is a product of adopting our novel pattern-configuration procedure onto the corresponding full conditional distributions to obtain a probabilistic scheme of drawing positions of molecules (Section 3.3.3). Briefly, for every pixel on every MCMC iteration, the drawn positions of the chosen configuration are recorded, along with the potential positions based on all the possible configurations the pattern can be formed on the corresponding moving window. At the end of the MCMC, the ratio of the drawn and the potential position constitutes the posterior probability map of the frame. Since we apply the individual frame counting scheme, molecules are assumed to have $r = \theta = 0$ within their corresponding pixels therefore no localisation counter is obtained throughout MCMC.

Alongside the posterior probability map, an additional probability map is obtained accounting for double events, corresponding to two simultaneously active molecules on the same pixel. Considering only the posterior probability map, we are able to identify pixels with a single molecule, however this does not provide any information about the number of molecules on the specific pixels. Since we assume that up to two PSFs are allowed to partially or completely overlap, translating into formation of patterns with up to two PSF overlapping on each structure within (e.g. see Figure 3.3.3), we also record the number of times the drawn positions correspond to complete overlap of two molecules. After the MCMC is complete, the ratio of the double events and the drawn positions provides a probability map capable of identifying pixels containing two molecules. In association with the posterior probability map,

pixels which are highly probable of containing a single molecule are accompanied with an additional probability to potentially capture a second molecule as a result of complete overlap of two PSFs.

Finally, as discussed in Section 3.4, we consider two types of filtering on the posterior probability map, and consequently on the double event probability map. According to our proposed model, the fact that the hidden Markov field is the unobserved mechanism declaring the states of the pixels it is safe to claim the molecules will lie somewhere within the determined 'On' pixels of the field. Therefore, we decide to filter the posterior probability maps using either the average field, obtained by taking the ratio of the sum of the fields and the number of iterations, or a binary (0/1) mask created by thresholding the average field. By definition, the former in general decreases the scale of the probabilities as it could be a multiplication with proportions less than one, while the latter keeps the same scale with the drawback of requiring an arbitrary threshold.

For this application, we choose to filter the posterior probability maps based on the average field. This choice does not make a big difference on inference since the results are almost identical to the ones produced using the binary mask, obtained with a threshold of 0.70 on the average field. In a total of 40.000 MCMC iterations after 10.000 iterations as a burn-in period, the posterior probability maps of the first F_t frames were obtained, $t \in \{1, \dots, 100\}$, with the chosen frames (Figure 5.3.1) displayed in Figure 5.3.2. The probability scale bars go from grey (lower probabilities) to white levels (higher probabilities), where large probabilities indicate pixels containing a single molecule. Additionally, the contour bins surrounding the corresponding pixels denote the densities of the regions. The posterior distributions of the parameters I , c , τ_b and $\{\beta_0, \beta_f\}$, the average fields and the traces of the chains for the chosen frames can be found in Appendix A.2.

We note here that due to the nature of our modelling, based on the pattern - configuration concept, there can be cases where a pattern has no compatible configurations.

A compatible configuration is a product of combinations of complete PSFs, as introduced in Section 3.3.3, which can contribute to form a pattern. Since the field, responsible for updating the states of the pixels, does not guarantee correct identification of every pixel in a PSF as 'On' on every iteration, there can be potentially patterns consisting of a large amount of 'On' pixels without compatible configurations. In these situations, our model updates the intensity of the entire moving region as being a product of only 'Off' pixels, since it is incapable of localising any molecules within due to the absence of configurations. This phenomenon can cause a detrimental effect on the parameter updating, especially regarding the background error precision τ_b which could experience high jumps to the lower values, potentially also causing a chain effect to the rest. Such cases, which are sparsely present during our procedure, are not included in the parameter updating.

Each one of the well separated PSFs can be accurately identified, assigning very high probabilities around 0.70-0.80 for the pixels where the molecules lie onto. This can be confirmed in a number of occasions for the chosen frames in Figure 5.3.2. The mass of the regions is also focused on the correctly identified pixels of the molecules as the bins indicate. Similarly, robust identification results are obtained regarding coinciding situations when there is no violation of our up to two PSFs overlapping limit. The scale and structure of overlap is invariant on our identification process which achieves separation of molecules, however compared to isolated molecules the probabilities obey a lower scale around 0.55-0.65. More specifically, PSF overlapping on four and six pixels of two closely located molecules can be spotted on Frame 11, whereas PSFs sharing two and three common pixels are present in Frame 51 and 79 respectively. The former achieves spatial separation of the events with corresponding probabilities around 0.59 and 0.62 for the four and six overlapping pixels cases accordingly. The latter also successfully completes the molecules pixel identification however on a lower probability level of around 0.54 and 0.57 for the two pixel collision on Frame 51, and approximately 0.54 for the three pixel case on Frame 79. For all of these overlapping cases, the probability mass is concentrated on the determined higher probable pixels

accounting for accuracy and precision of identification.

Regarding inference on high intensity regions, identification of pixels containing molecules can be problematic depending on the level of violation of our assumption. In such cases, less accurate estimates are provided corresponding to lower and more spread out probabilities around the region, not generally the case in regions free of violation. Specifically, high intensity regions without a violation, as on Frame 77, still perform properly with most of the density focused on the correct pixels, with slightly varying identification probability values of 0.5, 0.55, 0.57 and 0.56 for each one. Similar types of intensity regions formed by multiple PSFs, but with the difference of violating our assumption on a small degree, have less precise results or sometimes fail to identify every molecule involved. Such cases occur for instance on Frames 22, 30 and 51 (left bottom, left top and left top part of figures respectively), where on the first frame there is a clear and accurate separation of the overlapping molecules, that is 0.56 and 0.59 probabilities with high density, in contrast with the other two producing uncertain results regarding every molecule included. On Frame 30, the area where all three molecules overlap with each other, the pixels with two of the molecules obtain a fine probability with most of the mass (0.57 and 0.58), however the third molecule existing in between cannot be clearly separated with probability ≈ 0.35 . Even more uncertain is the region on Frame 51, resulting in assigning higher probabilities on pixels where there are no molecules, however still managing to capture two out of three pixels with acceptable probabilities.

High intensity regions with big violations of the assumption, as on Frame 3 (top middle part), lead to unreliable results even though there are exceptions where molecules can be accurately identified. In general, regions where most of the violation occurs tend to have a more spread out distribution of probabilities, which can either result in missing a number of molecules or falsely identifying pixels containing active molecules.




Figure 5.3.2: Posterior probability maps filtered by the average field for frames in Figure 5.3.1. Higher probabilities are indicated with white color levels as shown in the light blue probability scale bar on each probability map. The white bins represent the density around the regions.

In Figure 5.3.3, the double event probability maps for the chosés frames are displayed. As already mentioned, this probability map is obtained using the ratio of the drawn positions consisting of two molecules having their PSFs completely overlapping and the corresponding drawn positions for pixels containing single molecules. Similarly to the posterior probability maps, we filter the derived double event probabilities using the average field, relying on the fact that localised molecules must lie within the 'On' pixels of the field. In Figure 5.3.1, we can outline a number of occasions where two molecules have their PSFs completely overlapping, apparent for instance on Frames 11 (middle left part) and 79 (top right corner). As presented in Figure 5.3.3, looking at the double events regions and the corresponding probability levels, a clear identification of such events has been acquired. It should be also noted that, due to the probabilistic nature of our localisation procedure, regions which contain an overlapping scenario can be sometimes declared as ones with completely overlapping PSFs. Such regions can be found in Frames 30 (middle left part) or Frame 51 (right bottom part), with the double event probability map showing an effect although negligible when taking into consideration the corresponding probability levels.

In order to confirm and justify our claim about the different coinciding scenarios we focused on, we additionally obtain individual precision-recall-precision curves for each one of the chosen frames, serving as diagnostic metrics. These curves show the alteration in precision, i.e. positive predictability power, and recall or sensitivity, i.e. true positive rate or positive detection power, when multiple different thresholds are used to capture the true molecule positions. As can be seen in Figure 5.3.4, there is an obvious trade-off between precision and recall when our assumption is violated, whereas the rest of the cases, no matter the type of overlap, have optimal performance. The associated sensitivity-specificity plots, known as ROC curves [Fawcett, 2006], can be found in Appendix A.2, with specificity being the true negative rate accounting for the proportion of correctly identified pixels without a molecule.




Figure 5.3.3: Double event posterior probability maps filtered by the average field for frames in Figure 5.3.1. Higher probabilities are indicated with white color levels as shown in the light blue probability scale bar on each probability map. The white bins represent the density around the regions.




Figure 5.3.4: Precision-Recall curves for 'Circle within circle' frames in Figure 5.3.1. The lines show the performance of our model in terms of balancing the trade-off between precision and recall, with the adapting colour denoting different threshold values.

As already discussed above, the posterior probabilities are affected by the type of overlap present on specific regions within individual frames. Our probabilistic inference based on the pattern-configuration concept provides a reliable measure of uncertainty regarding the location of the molecules on separate frames, however the importance of visualising and exploring molecular behaviour and interaction make reconstructions of structures a necessity. Additionally, processing the derived probability maps is crucial from the quantitative perspective in order to obtain a count on the number of molecules. In a straightforward way, reconstructions of the individual underlying reality can be achieved by applying a threshold to the posterior probability maps on a specific level followed by placing molecules at the center (SiMPa functions origin) of the corresponding identified pixels. This is translated into $r = \theta = 0$ for the localised molecules, accounting for the individual frame counting scheme we perform on this application. The final reconstruction, or super resolution image, can be acquired by plotting all the individual positioned molecules on a single frame.

We focus on two ways for applying a threshold on the posterior probability maps in order to construct a super resolution image. A direct and potentially naive way is to consider fixed common thresholds for the entire frame, with the second way relying on a conventional threshold scheme based on the average field. We discuss both approaches in detail and present reconstructions for both the individual frames and the entire stack of frame F_t , $t = 1, \dots, 100$, accompanied with performance metrics for evaluation.

Fixed Threshold for frames in the stack

In Figure 5.3.5, we present the individual frame reconstructions for the chosen frames in Figure 5.3.1, based on two different fixed threshold values of 0.45 and 0.50. On each individual frame, the light blues crosses indicate the true positions of the molecules with the red empty and yellow filled circles corresponding to the localisations of molecules with $r = \theta = 0$ using the 0.45 and 0.50 thresholds respectively. The values are fixed for every frame and are chosen regardless of the overlapping situation, however the choice comes with a trade-off on identification. The importance and

effect of thresholds when perform a reconstruction can be spotted on a variety of occasions in Figure 5.3.5, with the most immediate and crucial effect being either missing or over-counting molecules.

Missing a molecule during a reconstruction procedure accounts for failing to place a molecule when in reality there exists one. Apart from our identification scheme being imprecise in terms of performance, incapability of identifying pixels where molecules lie onto could be a result of poor judgement on choices for thresholds applied on the posterior probability maps. Additionally, a consequence of different decisions on thresholds and/or our technique's inference accuracy could potentially read in falsely determination of pixels including molecules, or simpler over-counting. Therefore, performing a reliable reconstruction of an underlying structure using the posterior probability maps, sophisticated ways of determining the appropriate threshold levels need to be considered.

Applying the lower threshold level of 0.45, missing molecules is an issue only present in the case of violations on our assumption (Frame 3 and 30) where over-counting is more frequent, mainly on well isolated regions with molecules lying far away from the pixel's origin. On the other hand, adopting the threshold of 0.50 experiences the same limitation within a violation region, but overcomes the drawback of over-counting in most of the cases. A negative consequence of this higher threshold is potentially missing active molecules which were previously identified (Frame 77 - bottom left part or Frame 79 - middle right part), as a result of more uncertain regions due to overlapping PSFs, hence lower probabilities.

Conventional Threshold based on the average field

In Figure 5.3.6, we present the individual frame reconstructions for the chosen frames in Figure 5.3.1, based on a conventional threshold scheme obtained as a product of the average field. On a similar manner to fixed thresholds, the light blue crosses represent the true positions of the molecules, with the orange coloured circles denoting the positions obtained by applying the conventional threshold.

Specifically, we work with the average field for each individual frame obtained at the end our inference in order to decide on the potential complexity of the corresponding regions, hence adjust the level of thresholding. The average fields for the chosen frames can be found in Appendix A.2. In a way, high intensity regions where multiple overlapping or not PSFs potentially contribute to their formation, represent challenging regions of higher complexity therefore a lower threshold can serve the purpose. Molecules on regions where well isolated or fairly distanced PSFs are present can provide higher localisation precision, therefore we can use higher threshold values. With respect to the SiMPa functions, the intensities for a single active molecule spread out on a 3×3 neighbourhood region, hence a complete PSF is associated with nine pixels in total (same size as the moving region ; MR). In an ideal situation, a pixel containing a molecule will have its associated nine neighbouring pixels identified as 'On' by the average field. Taking this into consideration we can learn about the complexity of the outer region of the pixel by examining a larger neighbourhood on the average field, here we consider the 5×5 neighbourhood. For instance, if the larger neighbourhood contains eleven identified 'On' pixels we can assume fairly separated PSFs, or small overlapping levels, thus threshold higher, whereas a case with thirteen 'On' pixels is mostly probably a result of overlapping PSFs, thus we should threshold lower.

Our conventional threshold based on a large neighbourhood of the average field is presented in Table 5.3.1, applied on every individual pixel on each frame F_t , $t = 1, \dots, 100$, using the following notation. Let $\bar{\mathbf{x}}_t = \sum_{q=1}^{IT} x_t^{(q)} / IT$ denote the average field of frame t where IT the number of iterations and $\bar{x}_{s,t}$ a pixel $s \in \{1, \dots, N\}$ on the average field of frame t . If $\sum_s \partial \bar{x}_{s,t}$ denotes the sum of states in the 5×5 neighbourhood of pixel s in the average field t , then,

Convention based on 5×5 neighbourhoods of average field	Threshold
$\sum_s \partial \bar{x}_{s,t} < 12$	0.60
$\sum_s \partial \bar{x}_{s,t} \geq 12$ & $\sum_s \partial \bar{x}_{s,t} < 15$	0.45
$\sum_s \partial \bar{x}_{s,t} \geq 15$ & $\sum_s \partial \bar{x}_{s,t} < 18$	0.40
$\sum_s \partial \bar{x}_{s,t} > 18$	0.35

Table 5.3.1: Conventional threshold scheme based on the average field. The quantity $\sum \partial \bar{x}_{s,t}$ denotes the sum of states in the 5×5 neighbourhood of pixel s in the average field t , associated with conventional conditions to acquire a threshold level.

As can be spotted in Figure 5.3.6, adopting the conventional threshold improves the quantitative ability of our individual counting scheme compared to fixed common thresholds for the entire frame. This is supported and summarised on Table 5.3.2 where all threshold approaches are evaluated based on their identification capabilities. The table contains the following percentages; true positives (TP), correct identification of pixels containing molecules; false positives (FP), false identification of pixels containing molecules; true negatives (TN), correct identification of empty pixels; false negatives, false identification of empty pixels when they contain molecules. These metrics are accompanied with the real number of pixels containing molecules and empty pixels (Real). Additionally, the blue coloured quantities correspond to the performance of counting when we do not take into consideration reappearances of molecules due to blinking, whereas the black colour to the individual frame counts regardless of blinking.




Figure 5.3.5: Individual frame reconstructions for a number of frames included in the application on synthetic data, using two different thresholds for the posterior probability maps filtered by the average field. The red circles denote the drawn positions by thresholding the posterior probabilities filtered by the average field at 0.45, whereas the yellow coloured points at 0.5. The light blue coloured crosses represent the true positions of the molecules.




Figure 5.3.6: Individual frame reconstructions for a number of frames included in the application on synthetic data, using the conventional threshold scheme based on the average field for the posterior probability maps filtered by the average field. The orange coloured circles denote the drawn positions by applying the conventional threshold, with the light blue coloured crosses representing the true positions of the molecules.

Counting	TP	FP	TN	FN				
Thr = 0.45	95.2%	95.9%	1%	1.4%	99.1%	99.1%	4.8%	4.1%
Thr = 0.50	93.3%	93.9%	0.3%	0.6%	99.9%	99.9%	6.7%	6.1%
Thr = 0.55	80.8%	84.7%	0.1%	0.1%	99.9%	99.9%	19.2%	15.3%
Conv Thr	97.1%	98%	0.2%	0.5%	99.9%	99.9%	2.9%	2%
Real	104	98	-	-	1528	1528	-	-

Table 5.3.2: Performance of different threshold approaches in terms of counting for chosen frames in Figure 5.3.1. The table outlines the percentages of correct identified pixels containing the active molecules (TP) and correct identified empty pixels (TN), along with the incorrect positive (FP) and negative (FN) predictions. The black and blue colours denote counting when reappearances of molecules due to blinking are either considered or not respectively. The real number of pixels containing active molecules and empty pixels are also outlined (Real).

As a general comment, applying a fixed common threshold for all frames in the stack, undergoes a trade-off in precise identification and wrong determination of pixels including activated molecules. This effect does not seem to compose an issue if separate thresholds were chosen independently for each frame, based on the spatial association of the molecules within, a claim confirmed by the difference of the threshold scales on the precision-recall curves in Figure 5.3.4. Considering the conventional threshold scheme based on the average field seems to improve the performance of counting, mainly alleviating the over-counting issue when well separated or fairly distanced PSFs are of focus. However, there can still be room for improvement on manipulation of the probability maps obtained via our scheme for reconstructions, potentially by the use of an adaptive threshold based on different criteria, a matter not of our main focus throughout this thesis which discuss in Chapter 6. The detected alteration of the probability levels, and consequently the appropriate threshold scales, when different regions are of main focus as well various noise levels are present, is a case we investigate thoroughly in Section 5.4, where we conduct a complete sensitivity analysis.

Without making use of a tool to identify same molecules undergoing the blinking

process between frames, we present in Figure 5.3.8 the final reconstruction, or super resolution image, of the 'Circle within circle' sequence of frames F_t , $t = 1, \dots, 100$, by combining the Frame-wise localisation. Additionally, in Figure 5.3.7 the reconstruction using only the chosen frames from Figure 5.3.1 is displayed. In both figures, the light blue crosses indicate the true positions of the molecules without taking into consideration potential reappearances due to blinking, with the red, yellow and orange colour points representing the reconstructions based on 0.45, 0.50 and conventional thresholds respectively.

This lack of a mechanism to process the temporal dependency of consecutive frames, when using the individual frame inference based on SiMPa functions, can result in determining the same molecule on multiple occasions while remaining on the active state between frames. Furthermore, the effects identified individually on the chosen frames are consequently apparent and present on the final reconstruction, with the lower threshold leading to over-counting a number of molecules where a higher one is more accurate but fails to capture all of them at times. However, the conventional threshold scheme improves the performance of counting and seems more accurate, a matter supported by Table 5.3.3 which summarises the performance of identification.

Specifically, Table 5.3.3 contains the same metrics as Table 5.3.2, with TP, FP, TN and FN denoting the true positives (correct prediction of pixel containing a molecule), false positives (false prediction of pixel containing a molecule), true negatives (correct prediction of empty pixel) and false negatives (false prediction of empty pixel) respectively. The black and blue colours corresponds to counting when the real reappearances of molecules due to blinking are taken into consideration or not respectively. Focusing on the latter, since we applied the individual frame counting scheme, we can claim that the lower and conventional threshold scheme have similar detection power (here TP=95.7% in both cases), however the amount of incorrect determinations of pixels containing a molecule is significantly higher for the lower threshold (FP=9.5% against 3.1% for conventional). Consequently, this has an immediate effect on the detection of correct empty pixels (TN), whereas the amount




Figure 5.3.7: Final reconstructions of the super resolution image for the chosen frames in Figure 5.3.1, using two fixed thresholds for the posterior probability maps filtered by the average field and the conventional threshold scheme based on the average field. On the top left corner of the figure, the true positions of the molecules are shown (light blue crosses) with the 0.45 (red colour) and 0.5 (yellow colour) fixed threshold final reconstructions to be presented on the right top and left bottom corner respectively. On the bottom right corner of the figure, the final reconstruction using the conventional threshold scheme is displayed.




Figure 5.3.8: Final reconstruction of the super resolution image using the first 100 frames of the 'Circle within circle' stack of frames, using two fixed thresholds for the posterior probability maps filtered by the average field and the conventional threshold scheme based on the average field. On the top left corner of the figure, the true positions of the molecules are shown (light blue crosses) with the 0.45 (red colour) and 0.5 (yellow colour) fixed threshold final reconstructions to be presented on the right top and left bottom corner respectively. On the bottom right corner of the figure, the final reconstruction using the conventional threshold scheme is displayed (orange colour).

of missed molecules (FN) appears to be similar. As we discuss later in Section 5.4.1, the main reason for this incapability is the violation on the assumption of no more than two PSFs overlapping on a single region. Fairly similar behaviours and results are obtained by looking at the black coloured columns, with which we compare against the exact number of molecules due to re-activation between frames, with the false detection of pixels containing active molecules increasing noticeably.

Counting	TP	FP	TN	FN				
Thr = 0.45	95.7%	99.8%	9.5%	35.1%	98.1%	98.1%	3.5%	0.2%
Thr = 0.50	90.8%	98.5%	3%	26%	99.2%	99.2%	9.2%	1.5%
Thr = 0.55	81.5%	96.5%	1.2%	19.2%	99.8%	99.8%	18.5%	3.5%
Conv Thr	95.7%	99%	3.1%	28.9%	99.4%	99.4%	4.3%	1%
Real	892	479	-	-	1472	1472	-	-

Table 5.3.3: Performance of different threshold approaches in terms of counting for frames on application on synthetic data. The table outlines the percentages of correct identified pixels containing the active molecules (TP) and correct identified empty pixels (TN), along with the incorrect positive (FP) and negative (FN) predictions. The black and blue colours denote counting when reappearances of molecules due to blinking are either considered or not respectively. The real number of pixels containing active molecules and empty pixels are also outlined (Real).

5.3.2 Comparison with ThunderSTORM

In this section we apply ThunderSTORM, as introduced by Ovesný et al. [2014], on the 'Circle within circle' synthetic data in order to compare and evaluate the performance of our individual frame inference based on SiMPa functions. As we briefly discussed in Chapter 2, ThunderSTORM is an open source software for processing and visualising images obtained by a super resolution localisation microscopy technique, such as STORM. In order to perform an analysis, a number of steps are considered before the construction of a super resolution image, consisting of four main parts, (i) filtering (ii) identification of approximate molecules positions (iii)

sub-pixel localisation and (iv) post-processing. Each one of these parts can be performed on different ways according to the user's preferences, however default settings are suggested which we adopt and discuss in more detail.

Regarding (i), we use the wavelet filter based on B-spline basis functions. Following Ovesný et al. [2014], an input image V_0 is transformed at different levels j , here $j = 1, 2$, to $F_j = V_{j-1} - V_j$ based on the convolution in $V_j = (V_{j-1} * \mathbf{k}_j) * \mathbf{k}_j^T$, where \mathbf{k}_j the filter's convolution kernel. The V_j 's correspond to a band pass filter allowing only frequencies of a certain range. In this case, the kernels \mathbf{k}_j are obtained based on the B-spline basis functions of order $q = 1, 2, \dots$ [Izeddin et al., 2012], with the user able to adjust the order q and the scaling factor s which is part the kernel. The default settings correspond to the third order basis functions with $s = 2$, which we also consider for the application on the 'Circle within circle'.

The filtered image F , having the same size as the input image V_0 , is used in step (ii) of the analysis in order to determine approximate positions of the molecules, which are then used in order to perform the localisation. In the former, we use the local maxima for positions identification, where each pixel in F is considered as a potential candidate if its corresponding intensity is above a threshold, along with larger intensity value within a specific 4 or 8 pixel neighbourhoods, similar to Equations (3.3.1) and (3.3.2). Both of these features can be selected by the user with the choice of the threshold being of high importance. Following the default setting and what Izeddin et al. [2012] suggested, we keep F_2 to obtain the filtered image and the one times the standard deviation of the intensities F_1 , $s.d.(F_1)$, to derive a threshold level for the approximate positions.

The next step, involves localisation of the molecules based on the determined approximate positions, where the output of the procedure provides a complete set of coordinates for every identified molecule. ThunderSTORM supports a number of localisation procedures, some of which mentioned in Chapter 2, however here we apply two separate approaches, that is 2D Gaussian PSF fitting for single molecules using

weighted least squares and multiple emitter fitting. The former corresponds to fitting a PSF model similar to Equations (2.2.2) and (2.2.4), with the latter following a similar procedure as introduced by Huang et al. [2011a]. This corresponds to a sequential routine of fitting a single PSF model to localise one molecule, followed by fitting a second PSF model to the maximum intensity of the residual image, obtained after subtracting the initial localisation from the input image. This procedure is repeated until the maximum amount of allowed molecules is reached, a number that can be set by the user, followed by statistical tests to obtain the optimal amount (see Ovesný et al. [2014]). Here, we keep the default amount of five molecules. We also note that we set the fitting region to be of size 3×3 pixels similar to the predetermined neighbourhood of the SiMPa functions (Figure 3.2.1).

Applying ThunderSTORM using both single and multiple fitting, we use the estimates of the localised positions of the molecules in order to perform a reconstruction. In agreement with our application using the individual frame counting scheme based on SiMPa functions in Section 5.3.1, we present the individual frame reconstructions for the chosen frames in Figure 5.3.1 as well as the final reconstruction of all frames in the sequence F_t , $t = 1, \dots, 100$. In Figures 5.3.9 and 5.3.10, the individual frame reconstructions for the chosen frames are displayed using ThunderSTORM with single and multiple fitting respectively (yellow coloured circles). Additionally, the true positions of the molecules are marked (light blue crosses) along with the reconstructions using our conventional threshold scheme on the posterior probabilities maps (orange coloured circles), as described on the previous section. In both figures, well isolated regions containing single molecules are accurately separated and identified by both ThunderSTORM and our model, however, ThunderSTORM performance deteriorates in high intensity regions using either single or multiple fitting. These regions can be either regions with adjacent PSFs without overlap or ones where active molecules have their PSFs overlapping. In either case, ThunderSTORM struggles or even fails to detect and spatially separate all of the molecules contributing to the formation of such regions. As we mentioned in Section 5.3.1 and analyse in detail next

in Section 5.4.1, our model mainly struggles to separate PSFs in regions of interest that violate our assumption, of no more than two PSFs overlapping, but performing appropriately in every other case. The final reconstructions of all the frames F_t , $t = 1, \dots, 100$, in the sequence are displayed in Figure 5.3.11, where the individual frame reconstructions are combined and plotted. We note here that this procedure does not involve a mechanism to assess potential reappearances of molecules due to blinking.

In order to evaluate and compare the performance of our model against ThunderSTORM, we consider the localised molecules in the latter to obtain the corresponding pixels where molecules lie onto. In a similar manner as before, a correct identification of a pixel containing an active molecule is declared with TP (True positive) and a false identification with FP (False positive), whereas TN (True negative) and FN (False negative) are the relative correct and false declarations of empty pixel. In Table 5.3.4, we present the comparison of the counting performances, where the blue colours indicate the individual appearances of molecules within frames without considering potential reactivations, a case outlined with the black colour, and 'Thunder', 'Thunder+' and 'SiMPa' corresponding to single, multiple PSF fitting and conventional threshold on posterior probability maps using the SiMPa functions respectively.

Focusing on the final individual frame counting (blue colour), our model outperforms ThunderSTORM, in either localisation procedures of single or multiple PSF fitting. Regarding the positive detection capabilities, there is a quite large difference between 'SiMPa' and 'Thunder', that is TP=95.7% against 70.1%, consequently leading to a similarly large gap on the false detection of pixels containing an active molecule as empty, that is FN=4.3% against 29.9%. ThunderSTORM's counting performance improves when considering multiple emitters fitting, 'Thunder+', increasing the positive predictability power, from 70.1% to 83.5%, however still on noticeably lower levels than 'SiMPa'.




Figure 5.3.9: Individual frame reconstructions for selected frames included in the application on synthetic data, using the conventional threshold scheme based on the average field for the posterior probability maps and the localisations obtained by ThunderSTORM on single molecule fitting. The orange coloured and yellow filled circles denote the drawn positions by applying the conventional threshold and ThunderSTORM respectively. The light blue coloured crosses represent the true positions of the molecules.




Figure 5.3.10: Individual frame reconstructions for selected frames included in the application on synthetic data, using the conventional threshold scheme based on the average field for the posterior probability maps and the localisations obtained by ThunderSTORM on multiple-emitters fitting. The orange coloured and yellow filled circles denote the drawn positions by applying the conventional threshold and ThunderSTORM respectively. The light blue coloured crosses represent the true positions of the molecules.




Figure 5.3.11: Final reconstruction of the super resolution image using the first 100 frames of the 'Circle within circle' stack of frames, using the conventional threshold scheme for the posterior probability maps filtered by the average field and ThunderSTORM based on both single and multiple fitting. On the top left corner of the figure, the true positions of the molecules are shown (light blue crosses) with the conventional threshold scheme (orange colour) displayed on the right top corner. On the bottom right panel of the figure, the final reconstructions using ThunderSTORM (yellow colour) based on single molecule (left bottom panel) and multiple fitting (right bottom panel) are presented.

The percentage of false declaration of empty pixels as one containing a molecule is lower for 'Thunder' compared to 'SiMPa', FP=1.1% against 3.1%, an effect depended on the incapability of the former to capture most of the regions of interest. This improves with multiple fitting ('Thunder+') which allows a larger number of potential regions to be detected and localised, but also increases FP to 5.7% against 3.1% on 'SiMPa'.

Counting	TP	FP	TN	FN				
Thunder	70.1%	92.7%	1.1%	13.4%	99.9%	99.9%	29.9%	7.3%
Thunder+	83.5%	96.5%	5.7%	24.9%	98%	98%	16.5%	3.5%
SiMPa	95.7%	99%	3.1%	28.9%	99.4%	99.4%	4.3%	1%
Real	892	479	-	-	1472	1472	-	-

Table 5.3.4: Performance of ThunderSTORM and individual frame inference based on SiMPa functions for 'Circle within circle' synthetic data, where 'Thunder' denotes single fitting, 'Thunder+' multiple fitting and 'SiMPa' the conventional threshold scheme applied on the individual frame counting scheme based on SiMPa functions. The table outlines the percentages of correct identified pixels containing the active molecules (TP) and the number of correct identified empty pixels (TN), along with the incorrect positive predictions (FP) and negative (FN). The black and blue colours denote counting when reappearances of molecules due to blinking are either considered or not respectively. The real number of pixels containing active molecules and empty pixels are also outlined (Real).

5.4 Sensitivity analysis

5.4.1 Sensitivity analysis for counting and localisation

In this section, we conduct a sensitivity analysis on the counting and localisation procedures with varying levels of error which can affect the accuracy and precision of the posterior probabilities, our key tool in identifying pixels including molecules as well localising them within. This involve a variety of overlapping scenarios including cases of well isolated PSFs, coincidence of two PSFs on a different number of pixels, as well different kind of high intensity regions, with either violations of our assumption

or not.

In order to construct a sequence of frames having different noise levels, we use the structure of the 'Circle within circle' synthetic data generated in Section 5.2, while keeping the exact same positions of molecules amongst the individual frames. Apart from the background levels, all the parameters used in order to create the diffraction of the molecules based on the SiMPa functions remain unchanged, with the single event intensity $I = 8$, background intensity $I_0 = 2.4$ and power of spread $c = 0.4$. After every separate 'Circles within circle' frame is obtained, we assign zero mean Gaussian noise of three different levels, while also keep the free of noise scenario as an additional case. We name as noiseless, small, regular and large noise cases the corresponding scenarios under background noise with precision $\tau_b = 100, 10, 1$, or equivalently variance $\sigma_b^2 = 0.01, 0.1, 1$, for the last three cases respectively.

In Figure 5.4.1, the distribution of signal using the SiMPa functions is presented along with the distribution of the background intensity, when each one of the four noise levels is applied. Here, we refer to signal as the determined SiMPa values according to a distance r and direction θ when $I = 8$ and $c = 0.4$. Initially, the SiMPa functions (Equation 3.2.1) are evaluated over a large set of combinations of r and θ forming the noiseless case in (a), before the obtained quantities become subject to noise in (b), (c) and (d). The distribution of the signal is denoted with the light blue coloured histograms while the background intensity with the light green. The noiseless and small noise scenarios have a clear separation of the signal and background, with the latter causing only a small drift from the true SiMPa values. We note here that the behaviour around 8 corresponds to the central pixels containing the molecules, which under these noise levels are not affected noticeably, hence the peaks, not the case for regular and large noise. A larger variability, responsible also for a small collision between signal and background is apparent on regular noise level, whereas the large level results in a mixture of the two distributions, potentially challenging in distinguishing between signal and background. The regular noise level in (b) was considered to generate the 'Circle within circle' synthetic data for the application in

Section 5.3.




Figure 5.4.1: Different noise levels when generating the 'Circle within circle' synthetic data. The light blue histograms represent the diffraction according to the SiMPa functions for continuous combinations of r and θ with $I = 8$, $c = 0.4$ and different zero mean Gaussian distributed noise. The dark green histograms show the background intensity where $I_0 = I \times d = 2.4$ with $d = 0.3$. (a) No noise (b) $\tau_b = 100$ (c) $\tau_b = 10$ and (d) $\tau_b = 1$.

5.4.1.1 Prior parameters

Before progressing into the sensitivity analysis, we briefly discuss the selection of the corresponding prior distributions for each one of the parameters. In agreement with the prior elicitation for the regular noise case in Section 5.2, we consider the same quantile procedure under similar limits in order to obtain prior distributions for I and τ_b , whereas the ones for c and β 's are unchanged as in Figure 5.2.3 (right column). For the noiseless case, we apply our model by assigning τ_b the elicited prior

for the lowest noise, that is for the small noise case with $\tau_b = 100$. The corresponding elicited prior distributions for I and τ_b , in case of noise free, small and large noise, are displayed in Figure 5.4.2.




Figure 5.4.2: Priors elicited for single event intensity I and precision τ_b when small and large noise levels are present. On the left panel of the figure, the elicited priors under small noise levels are shown, where on the right panel the ones under large noise. The priors from the small noise case are also used for the free of noise scenario. The blue μ marks denote the means of the distributions whereas I_0 the background intensity.

5.4.1.2 Varying noise levels

In order to assess the accuracy of our localisation inference different types of overlapping cases need to be taken into consideration. For that reason, we select four representative frames from the generated stack of 'Circle within circle' data which contain most of the potential structures within a frame, according to our proposed

model based on SiMPa functions. Both for the counting and localisation procedures, these frames are thoroughly investigated and compared under all separate error levels described above, where a detailed analysis is also performed by closely assessing the corresponding overlapping cases. For the localisation algorithm we consider the same discretisation of r and θ as in Section 3.4, with possible directions $\theta = k\frac{\pi}{4}$, $k=0,1,\dots,8$ under the distances $r = 0$ and $r = \frac{r_{\max,\theta}}{2}$, corresponding to the origin of the SiMPa functions and half the distance from the origin on direction θ respectively.

Regular overlapping cases without violation of assumption

Frame 57, presented in the left hand side of Figure 5.4.3 under every different noise level, includes some of the regular PSF coinciding scenarios along with well separated PSFs. Our assumption of up to two PSFs allowed to overlap stands, hence this frame is characterised as a regular overlapping case without violation. In the middle and right hand sides of the figure, the posterior probability maps filtered by the average field are presented under the counting and localisation inference respectively, with the blue boxes indicating the regions of different cases. In general, both procedures perform equally well in all of the noiseless, small and regular noise cases, with counting showing good stability and precision in identification, confirmed by the density having most of mass around the correct pixels containing the molecules. Regarding the large noise case, despite being also visually apparent how it affects the PSFs (left hand side of figure - (d)), both inference procedures seem to perform quite decently for every molecule within the frame. However, there is higher uncertainty on the positions of the molecules regardless the overlapping scenario they belong to, a case apparent by the behaviour and structure of the density bins.




Figure 5.4.3: 'Circle within circle' data - Frame 57 under different noise levels (left) along with posterior probability maps filtered by the average field, using both counting (middle) and localisation (right) algorithms. The light blue crosses on the frames denote the true locations of the molecules, while the blue squares A, B and C regular overlapping scenarios which are investigated thoroughly in Figures 5.4.4, 5.4.5 and 5.4.6 respectively.

For every focused region throughout this analysis, we adopt the same structure for figures consisting of three main parts. The first part, on the left hand side of figures, includes the detailed focused regions under the different error levels, with the true positions of the molecules outlined with the light blue color. The second and third part, on the middle and right hand side of figures respectively, consist of the posterior probability maps filtered by the average field for the counting and localisation procedures respectively, where the actual determined probabilities of the pixels above 0.20 are also outlined with the light blue colour.

Well separated PSF

Region A on Frame 57, presented in Figure 5.4.4, consists of only one well separated PSF. For the small, regular and free of noise cases, both the counting and localisation algorithms perform equally well, assigning probabilities above 65% on the correct pixel containing the molecule. The only difference is the uncertainty around the regions, with the latter being less precise in terms of scales on the neighbouring assigned probabilities. Regarding the large noise levels, both of our algorithms still are able to recognise the correct positions of the molecules, however this well separated PSF was specifically chosen to show an effect which can potentially be present under this error. Since we reconstruct a frame by applying a threshold to the obtained posterior probabilities, the determined level can alter the number and places of molecules. In this case, a threshold of about 0.57 for counting, and 0.46 for localisation, achieve correct identification of the pixel where the molecule lies onto, however additionally wrongly assigns a second one on a neighbouring pixel, causing over-counting or misplacing of a molecule as discussed in Section 5.3.1. For an isolated PSF, this scenario can occur when high level of noise prevents pixels which are part of the PSF to be on average considered as 'On' by the field. Consequently, this results in a constant lower declaration of such pixels being active, resulting on a reduction of precision on the positions inference.

Overlapping PSFs on four pixels

Similarly, precise and quite similar results are obtained for the four overlapping

pixels situation, outlined with Region B on Frame 57, closely focused in Figure 5.4.5. The posterior probabilities under all noisy occasions show capability of accurately spatial separation of the coinciding PSFs, with clearly distinguishable levels from the neighbouring pixels. The uncertainty levels are higher for the localisation algorithm, especially under the large noise scenario, however a threshold ≈ 0.50 achieves optimal placement of molecules in terms of reconstruction.

Overlapping PSFs on six and one pixels

The marked Region C in Frame 57, presented in Figure 5.4.6, consists of two different overlapping cases, each one consisting of two PSFs overlapping on six and one pixels respectively. Based on the counting algorithm, in both cases the molecules are distinguished from their corresponding PSFs with high accuracy. The precision of the results is smaller when it comes to large noise, with probabilities dropping from over 0.55 to ≈ 0.45 . The localisation procedure appears to struggle to identify the correct pixels containing the molecules for the six pixels overlap in the noiseless and small case scenario. The probabilities of the correct pixels are almost identical to two other pixels in the region at around 0.42. This can be an effect of not precise identification on average of the entire region by the field, as discussed above, and/or the complexity levels of the localisation procedure. The latter corresponds to the large number of configurations associated with a realisation existing for any specific pattern formed within a moving region (see Section 3.3.2). In contrast, under regular additive noise, possibly as a result of blurring, the localisation process becomes more adaptive and precise achieving more accurate separation of the two overlapping PSFs on six pixels. For the one common pixel case, the molecules become distinguishable with high certainty, which substantially decreases in scale and becomes less reliable under large error levels.

Small high intensity regions without violation of assumption

Frame 6 in Figure 5.4.7, is a representative regular frame with some overlapping PSF. The main reason we also focus on this one is that Frame 6 falls within the




Figure 5.4.4: Selected region A in Frame 57 along with posterior probability maps filtered by average field, using both counting and localisation algorithms - Well isolated PSF case. The light blue crosses denote the true locations of the molecules, while the actual posterior probabilities higher than 30% are outlined both for the counting and localisation procedures.




Figure 5.4.5: Selected region B in Frame 57 along with posterior probability maps filtered by average field, using both counting and localisation algorithms - Two PSFs overlapping on four pixels. The light blue crosses denote the true locations of the molecules, while the actual posterior probabilities higher than 30% are outlined both for the counting and localisation procedures.




Figure 5.4.6: Selected region C in Frame 57 along with posterior probability maps filtered by average field, using both counting and localisation algorithms - Two cases of two PSFs overlapping on six and one pixel. The light blue crosses denote the true locations of the molecules, while the actual posterior probabilities higher than 30% are outlined both for the counting and localisation procedures.

category we call small high-intensity region without violation, where more than two PSFs, that is three or four, form a region of interest of higher total intensity than ordinary. In agreement with the posterior probability maps for Frame 57, both procedures of either counting or localising provide reliable identification for pixels containing molecules, with the latter being generally less precise. Regarding the large noise case for both procedures, an issue can be recognised on the top right corner of the figure, where the overlapping of two PSFs on two pixels is highly affected by the error levels. Even though separation is achieved with less certainty, these noise scales could lead to the field being more variable, from incapability of recognising complete regions of interest on a stable basis, followed by larger uncertainty on inference.

This small high-intensity region in Frame 6, presented in Figure 5.4.8, is formed by two molecules having their PSFs sharing six pixels with an additional PSF of a third active molecule right next to the other ones. Based on the posterior probabilities of the counting procedure, all of the small, regular and free of noise cases provide rather accurate identifications of pixels containing the molecules, achieving clear spatial separation of the corresponding PSFs. In contrast, uncertainty is much higher on the larger error case, which apart from being responsible for the lower scale of precision, also seems to prevent the complete separation of PSFs. The correct pixels are still distinguishable with probabilities of 0.49, 0.49 and 0.41 based on the counting algorithm, although the last probability is fairly close to the surrounding ones. This is not exactly the case for the localisation algorithm. Less accurate results are even apparent on the noiseless case, even though two of the three PSFs are positively resolved under any noise type. The third molecule's position probabilities are 0.43 for the noiseless and small noise cases and ≈ 0.38 for regular and large cases, preventing clear distinction within the region. Additionally, the pixel of the PSF on the right middle end is consistently wrongly identified, with 0.46, 0.46, 0.46 and 0.45 for the different noise levels respectively, a case that could lead to over-counting and/or misplacing due to a regular threshold ≈ 0.40 or missing molecules ≈ 0.47 .




Figure 5.4.7: 'Circle within circle' data - Frame 6 under different noise levels (left) along with posterior probability maps filtered by the average field, using both counting (middle) and localisation (right) algorithms. The light blue crosses on the frames denote the true locations of the molecules, while the blue square A a small high intensity region without violation, investigated thoroughly in Figure 5.4.8




Figure 5.4.8: Selected region A in Frame 6 along with posterior probability maps filtered by average field, using both counting and localisation algorithms - Three PSFs forming a small high intensity region without violation. The light blue crosses denote the true locations of the molecules, while the actual posterior probabilities higher than 30% are outlined both for the counting and localisation procedures.

Large high intensity regions without violation of assumption

Frame 62, presented in Figure 5.4.9, consists of a main region from the case we call large high-intensity region under no violation. In such a case, multiple PSFs which are allowed pairwise overlapping form a large region of high total intensity. As a general comment, inference in these types of regions can be affected by different error levels, while at the same time the determined inference algorithm can arguably provide altered results as we discuss next. Similarly to Frames 6 and 57, under no or small noise, both counting and localising procedures identify the correct pixels containing the molecules, with the former being noticeably more precise. Furthermore, equally accurate identification is derived under regular noise scale during counting, however this is not the case for the localisation procedure. Lastly, both procedures can be either less precise or even have essential issues regarding large scaled noise.

Specifically, in Figure 5.4.10, we outline the big high-intensity Region A under no violation in more detail. As already spotted from the larger scale picture in Figure 5.4.9, the small, regular and free of noise cases allow the counting procedure to obtain an accurate and reliable identification outcome, whereas the large noisy version decreases the precision for a number of pixels where molecules lie into. Additionally, it is essentially affected by the error in terms of missing some of the molecules, if a fairly low threshold ≈ 0.40 was considered for placing them. Regarding the localisation procedure, things get complicated earlier on, where even though under small and no noise levels separation and identification of the correct pixels is mostly achieved, the levels of uncertainty are quite high throughout the neighbourhood. This could lead to one, or more, of the described issues of over-counting, misplacing and/or missing molecules, a statement quite obvious when evaluating the probabilities within the region. As a matter of fact, these potential issues have a knock on effect when moving to cases of bigger noise levels, corresponding to either regular or large noise scenarios. In such high intensity regions, the vast amount of possible configurations associated with realisations can potentially prevent the localisation scheme of performing decently (see Section 3.3.2).




Figure 5.4.9: 'Circle within circle' data - Frame 62 under different noise levels (left) along with posterior probability maps filtered by the average field, using both counting (middle) and localisation (right) algorithms. The light blue crosses on the frames denote the true locations of the molecules, while the blue square A a big high intensity region without violation, investigated thoroughly in Figure 5.4.10




Figure 5.4.10: Selected region A in Frame 6 along with posterior probability maps filtered by average field, using both counting and localisation algorithms - Three PSFs forming a large high intensity region without violation. The light blue crosses denote the true locations of the molecules, while the actual posterior probabilities higher than 30% are outlined both for the counting and localisation procedures.

High intensity regions under violation of assumption

In Frame 23, presented in Figure 5.4.11, the posterior probability maps using both the counting and localisation algorithms are shown, along with two separate cases which belong to the category that violates our main assumption of having at most two PSFs overlapping in a region. In general, as the probability heat-maps outline, the overall performance of our inference schemes are not affected by the corresponding violations regardless of the noise levels. On the contrary, the corresponding regions with violations tend to provide alternating identification inference under the various error levels. Specifically, two different regions are of main focus, with Region A falling within the category of big high-intensity region under violation, in which multiple PSFs (in this case four) coincide on several occasions. In a similar manner, Region B belongs to a closely related category of small high-intensity regions under violation, accounting for existence of commonly shared pixels between three PSFs.

Regarding the small high-intensity Region B, displayed in Figure 5.4.12, the capability of the counting algorithm to identify the pixels containing the molecules appears to be invariant of the violation. In terms of reconstructing the region, for every noisy scenario a regular threshold value of ≈ 0.45 accounts for optimal performance, however a smaller one at ≈ 0.40 results in wrongly placing molecules in some pixels. This is not exactly the case based on the localisation algorithm. For the small, regular and free of noise scenarios, a similar issue can be outlined accounting for imprecision on identification of the correct pixels regardless of the threshold level for reconstruction. A similar value of ≈ 0.45 achieves separation for two of the molecules within the region however falsely includes a third one (right bottom corner). A lower threshold of ≈ 0.40 can resolve every molecule included, on the drawback of still falsely placing a fourth one on the same location. In the most noisy version, uncertainty overwhelms the entire localisation inference, assigning probabilities on a larger neighbourhood, however the position inference is quite similar to the regular noise case.




Figure 5.4.11: 'Circle within circle' data - Frame 23 under different noise levels (left) along with posterior probability maps filtered by the average field, using both counting (middle) and localisation (right) algorithms. The light blue crosses on the frames denote the true locations of the molecules, while the blue squares A and B a small and big high intensity region under violation, investigated thoroughly in Figures 5.4.13 and 5.4.12 respectively.




Figure 5.4.12: Selected region B in Frame 23 along with posterior probability maps filtered by average field, using both counting and localisation algorithms - Three PSFs forming a small high intensity region under violation. The light blue crosses denote the true locations of the molecules, while the actual posterior probabilities higher than 30% are outlined both for the counting and localisation procedures.

Large high intensity regions under violation of assumption

Focusing at Region A in Figure 5.4.13, of big high-intensity under violation, both algorithms regardless of the noise case seem to face difficulties on providing accurate and precise inference for pixels containing molecules. Based on the counting procedure, accuracy in recognition of the molecules within the region is fairly similar across small, regular and free of error cases, with most uncertain inference around the area where three molecules overlap with each other (middle area). In terms of placing molecules according to a probabilities threshold, values ≈ 0.40 lead to almost ideal results in small and free of noise cases, subject to a single missed molecule, an issue not present on the regular noise results, however might be a result of randomness. In a decision of lowering the threshold below 0.40, all correct pixels can be identified under the disadvantage of falsely placing more molecules within the region. Large scale noise under this scenario seems to severely affect the identification capabilities of the algorithm, producing unstable results which can lead to both missing molecules and wrong placement of molecules. This is also the case for the localisation procedure with the inference being even more variable. Uncertainty appears to overpower the predictability of the procedure even under small, regular and free of noise versions although on a less severe scale. Correct recognition of the molecules pixels can be obtained under various thresholds, however no matter the determined level, misplacements and/or missed molecules outcome is inevitable.

The precision-recall curves for the frames investigated are presented in Figures 5.4.14 and 5.4.15 for the small/free of noise and regular/large noise cases respectively, confirming the effect large noise has in the inference procedures, even though both still perform quite well with fairly similar trade-offs under close thresholds. This negative impact can be identified in most of the frames we investigated, with lower noise levels appearing to have a small influence on identification of pixels containing the active molecules. The change in accuracy and precision can additionally be outlined when challenging overlapping and/or not regions are present, again mainly problematic under large noise. The sensitivity-specificity curves can be found in Appendix A.3.




Figure 5.4.13: Selected region A in Frame 6 along with posterior probability maps filtered by average field, using both counting and localisation algorithms - Five PSFs forming a large high intensity region under violation. The light blue crosses denote the true locations of the molecules, while the actual posterior probabilities higher than 30% are outlined both for the counting and localisation procedures.




Figure 5.4.14: Precision-Recall curves under small and free of noise levels for the frames used in sensitivity analysis, including both the counting and localisation procedures. The line types denote the corresponding noise cases whereas the adaptive color the trade-off in precision-recall for different thresholds on the posterior probabilities.




Figure 5.4.15: Precision-Recall curves under regular and large noise levels for the frames used in sensitivity analysis, including both the counting and localisation procedures. The line types denote the corresponding noise cases whereas the adaptive color the trade-off in precision-recall for different thresholds on the posterior probabilities.

High intensity regions under extreme violations of assumption

Lastly, we investigate a special case corresponding to extreme scenarios, where a patch is produced from multiple PSF overlapping in a fairly small region under violation of our assumption on a large scale. In such cases, procedural issues can potentially arise, mainly because of the incapability of our algorithms to estimate the corresponding intensities of the focused spot leading to inaccurate estimation of parameters and pixels containing the molecules. On the top panel (a) of Figure 5.4.16, we present Frame 1 and 10 from the 'Circle within circle' stack reading as representative frames of extreme violations of the assumption. Frame 1, consists of a number of overlapping PSFs with a number of violations of the assumption, along with an extreme violation on a high-intensity region formed by four molecules active really close to each other (bottom left part of figure). On a similar manner, Frame 10 consists of a few PSFs satisfying the assumption whereas a high intensity region formed by six simultaneously active molecules is present on the right middle panel of the figure.

These type of extreme violations can potentially cause our procedures to fail on identifying every region of interest (ROI), as the field is unable to determine the pixels that belong to a PSF on a stable basis, as a result of a negative chain effect on estimation of the parameters. The average fields for these two frames are displayed on part (b) of Figure 5.4.16 where this instability of acquiring the majority of true 'On' pixels is apparent. As a consequence of this, the posterior probability maps (part (c)), which depend on the states of the fields on every iteration of the MCMC, struggle to provide accurate positions of molecules and are highly uncertain, not only on the regions with violations but generally on the frame. In reality, when high density data are analysed, the assumption of up to two PSFs overlapping can be relaxed in order to allow for more intense overlapping, however with a cost on computing times and possibly on the stability of the probabilistic scheme, as we discuss in Chapter 6.




Figure 5.4.16: Position inference using the counting algorithm on Frames 1 and 10 from 'Circles within circle' data used in application. (a) Raw intensities of Frames 1 and 10, with the light blue crosses denoting the true positions of the molecules (b) Average fields of Frames 1 and 10, with the scale bar denoting the 'On' probabilities for each pixel. (c) Posterior probability heat-maps filtered by the average field.

In more detail, during the update of the parameters the total intensity of a moving region based on the SiMPa functions has an upper bound, according to the potential maximum intensity using the most up to date parameters and appropriate configurations with respect to the pattern. As a matter of fact, on an extreme overlapping situation, even if our probabilistic scheme accomplishes to regularly draw configurations accounting for maximum corresponding total intensity, a big distance between the observed and estimated intensity of the moving region still exists. This has a major impact on the parameter estimation, starting on the estimation of the background error precision τ_b , followed by a chain effect on I and c , or the other way around. This distance can prevent τ_b 's estimation to move away from values approaching zero or result in a bimodal distribution with higher values. The latter is the case for Frame 10 while the former can be spotted on Frame 1, both displayed on the middle part of Figure 5.4.17. Additionally, in an attempt to reach the true intensity of the moving region, the spread parameter c goes down to zero which allows brightest neighbourhoods, or slower diffraction, and the single event intensity I to higher values, resulting in higher intensity values within the moving region. The prior and posterior distributions for I and c for Frames 1 and 10 are displayed on the top and bottom panel of Figure 5.4.17.

As briefly mentioned above, this chain effect prevents the field to provide sensible regions of interest and the update is fairly unstable. On average, mainly pixels of higher intensity tend to be identified as 'On' with a number of pixels that are in reality parts of a PSF to be potentially switched 'Off'. The instability of the field is more obvious on Frame 1, where both τ_b and c go to zero, and less on Frame 10 (part (b) of Figure 5.4.16). In either cases, the outcome of identifying pixels containing molecules is fairly similar, led by high levels of uncertainty on the posterior probabilities, a fact which can cause either over-counting and/or missing active molecules.




Figure 5.4.17: Prior and Posterior distributions of I (top panel of figure), τ_b (middle panel of figure) and c (bottom panel of figure) for Frames 1 and 10 from 'Circles within circle' data used in application. The histograms along with the blue solid line represent the posterior samples with the red solid lines being the corresponding prior distributions.

General conclusions on localisation

On average, the counting scheme is consistent on identifying the correct pixels containing active molecules, when no violations of the assumption exist, regardless of the noise levels. This identification corresponds to precisely placing molecules at the center of pixels, or the origin of the SiMPa functions with $r = \theta = 0$. Under scenarios that violate the assumption of up to two PSFs allowed to overlap, the counting algorithm still performs quite well in relation with the scale of violation in terms of the complexity levels produced in the focused region, however there is a possibility

of over-counting and/or missing events. Despite the accuracy of the probabilistic outcome of our model, levels of threshold can be crucial when reconstructing the specimen under investigation, or placing molecules, leading to the need of acquiring thresholds using sophisticated ways, a fact we further discuss in Chapter 6.

In a similar manner, the localisation scheme also performs well on average regardless of the overlapping case when no violation occurs. However, localisation appears much more sensitive to the scale of noise, with large levels noticeably affect the identification capability of the procedure. It is highly complex as a process, mainly due to the probabilistic scheme relying on a large, though finite, number of configurations associated with a realisation for most of the potential patterns (see Section 3.3.3). As a result, this can bound the precision on identification of pixels containing molecules when a moving region focuses on high intensity regions. In cases of violations of the assumption, the probabilistic identification appears equally unstable, similarly with respect to the level of violation, specifically under large noise levels.

Furthermore, the localisation procedure does not provide clearly distinguishable positions of molecules within their corresponding identified pixels, hence allows room for improvement (discussion on Chapter 6). The majority of probability sets for positions of molecules within the identified pixels according to the discretisation of r and θ , appears to be uniformly distributed, especially for the regular and large noise levels. The corresponding realisation probability heat-maps are omitted from this section but can be found in Appendix A.3.

Regardless of the chosen algorithm, frames containing extreme cases of overlapping molecules, in terms of being subject to large scale violations of the assumption, potential issues in stable identification of regions of interest (ROI) can be addressed. The main reason is the incapability of our procedures to estimate the corresponding intensities of the extreme regions leading to inaccurate parameter estimation. The parameters τ_b and c tend to move down to low values (close to zero) with I shifting to higher levels. In these cases, the field is noticeably affected by being generally

unstable and often failing to identify correctly the 'On' pixels, resulting in less reliable localisation inference. Relaxing the assumption of no more than two PSFs overlapping is straightforward, a case we discuss in detail in Chapter 6.

Lastly, a similar issue can potentially be the case when excessive error levels exist within frames which contain many overlapping patches. This can result in pixels altering entirely their intensity levels, especially the ones which are parts of a PSF and do not contain an active molecule, leading to an unstable field hence less precise localisation. In such cases, more informative prior distributions could improve the performance and enhance the capabilities of the procedures however do not necessarily guarantee alleviation of the issues. On a general comment, the counting algorithm shows more stability even under more challenging situations and is less computationally intensive than the localisation procedure.

5.4.2 Sensitivity analysis on parameters

In this section we perform a sensitivity analysis on the parameters needed to perform our proposed methodology, when either using the individual frame inference scheme or the Markov switching model based on SiMPa functions. Using the individual frame counting scheme we investigate the effect of the prior distributions on the single event intensity I , background error precision τ_b , power of spread parameter c and the field control parameters β_0 and β_f . Regarding the Markov switching model we only consider different prior distributions for the transition matrix ξ since the rest of the parameters remain the same. Lastly, we stress out the importance of having a well established estimate for the proportion of the background intensity d since based on our proposed model the background intensity $I_0 = d \times I$, with d being a constant, a matter that can be relaxed as we discuss in Chapter 6.

Regarding the individual frame scheme part, where we perform an analysis on the effect of different prior parameter configurations on inference, we consider a single frame from the generated 'Circle within circle' synthetic data used for the appli-

cation in Section 5.3. This frame, Frame 338, is selected as a representative one constituting of a number of well isolated and overlapping PSFs when the assumption of no more than two PSFs overlapping is not violated, and is displayed in Figure 5.2.1. Additionally, we revisit the subset from the stack of the 'Circles within circles' considered during the implementation of the Markov switching model based on the SiMPa functions in Section 4.4, displayed in Figure 4.4.1. This subset constitutes the baseline to the analysis of the effect of different prior distributions on the transition matrix ξ .

In order to conduct the analysis on the parameters we consider two types of priors apart from the elicited (for I and τ_b) and chosen (c and β_0 & β_f) ones in Figure 5.2.3. The one set of prior parameters has an identical or similar location as the elicited or chosen but assigned a fairly larger variance, while the second set retains the elicited variance however shifting the location away either to the right or left. We name the former set 'Larger variance case' whereas the latter 'Location shifted case'. Regarding the prior distributions for β_0 & β_f , we decide not to investigate on a location shifted case, since values on the scale between -2 and 2 are reasonable parameter values which serve the purpose. In Figure 5.4.18 we present the corresponding prior distributions for I , τ_b , c and β_0 & β_f which we consider throughout this analysis. The red solid lines denote either the elicited or chosen prior distributions, whereas the blue and light green solid lines correspond to the larger variance and location shifted cases respectively. We focus on the effect these prior distributions have on the corresponding posterior distributions and consequently on localisation, by changing one prior at a time while keeping the rest on their pre-determined structure, as denoted with the red solid lines in Figure 5.4.18.




Figure 5.4.18: Prior distributions for parameters I , τ_b , c and β_0 & β_f used for sensitivity analysis. The red solid lines correspond to either the elicited (I and τ_b) or chosen (c and β 's), whereas the blue and light green solid lines to the larger variance and location shifted cases respectively.

In Figure 5.4.19, we present the posterior distributions of every parameter under the different corresponding prior distributions. Each sub-figure has been obtained from an individual implementation of our counting algorithm based on SiMPa functions on Frame 338, on a run of 30000 MCMC iterations after a burn-in period of 10000 iterations. The left column of the figure corresponds to the obtained posterior distributions of the separate runs using the elicited or chosen prior distributions, while the

middle and right columns consist of the posterior distributions under larger variance and location shifted cases respectively. The red solid lines denote the prior distributions of each parameter, while the histograms associated with the blue solid density lines the corresponding posteriors.

In order to investigate how the posterior distributions are affected by the prior choice, hence potentially the localisation inference, we compare the sub-figures in Figure 5.4.19 row-wise which can show and outline shifts in the posteriors. As can be easily identified, the prior choices do not alter the behaviour of the model therefore it is safe to say that we obtain robust estimates regardless of the initial choices. As a matter of fact in Bayesian analysis, if the prior distributions are completely wrong, mainly referring to extreme cases with point mass distributions centred far from truth, this can potentially highly damage our model inference. This is a situation not very common in super-resolution imaging where large amounts of data can be obtained, however we suggest the use of non-informative prior distributions in clueless cases. One should be careful on the prior determination of the single event intensity I , and consequently on the background proportion d , if he was to proceed on a prior elicitation process (as for instance in Section 5.2) to acquire hyper-parameters. We next investigate the effect on choices of d which we consider as a constant throughout this thesis. We note that the rest of the posterior distributions for each one of the cases investigated and presented throughout this section, as well as the posterior probability heat-maps filtered by the average field, can be found in Appendix A.3, omitted from the main body due to their robustness and almost perfect similarity.

We consider two separate values for the background proportion d , one higher than the true value with $d = 0.45$ and one lower with $d = 0.20$, while keeping the original elicited and chosen prior for each one of the parameters (Figure 5.2.3). In both cases, the parameter estimation of I , τ_b and c is affected, shifting the location of the posterior distributions either to the left or right. Regarding $d = 0.45$, corresponding to the left panel of the Figure 5.4.20, the posterior of the precision τ_b is robust (similar to Figure 5.4.19), however the posterior distributions of I and c are shifted to the




Figure 5.4.19: Prior/Posterior distributions obtained using different priors for the parameters I , τ_b , c and β_0 & β_f , part of the sensitivity analysis on the parameters. Each separate prior used belongs to either the elicited location but larger variance case or location shifted on same elicited variance. The histograms represent the posterior samples whereas the red solid lines the prior distributions. The light blue histogram on the bottom part relates to β_f while the regular one to β_0 .

left assigning lower values, as a result of a knock on effect. Since I , which corresponds to the level of the intensities with respect to the SiMPa functions, is dropping, in order for the model to approach the performance of the SiMPas it acquires tighter diffraction within the neighbourhoods, hence c is going down. In a similar manner, a background proportion of $d = 0.20$ (right panel of Figure 5.4.20), causes an overestimation of I shifting its posterior distribution to the right, which consequently drags the posterior of c to the right. Large values of c boost the SiMPa functions which translates into faster intensities diffraction within a neighbourhood, a consequence on our model attempting to approach the observed intensities. Additionally, in this case our model struggles to accurately estimate the background noise precision τ_b , shifting the posterior distribution to the left as displayed on the middle right panel of Figure 5.4.20. We note that the posterior distributions of β_0 & β_f remain unaffected in both cases, hence not presented in the section.

These types of behaviours of the parameters can also potentially affect the localisation process. In Figure 5.4.21, we present the posterior probability heat-maps filter by the average field for Frame 338, when $d = 0.45$ (left panel of the figure) and $d = 0.20$ (right panel of the figure). Higher probabilities denote pixels which contain a single active molecule, with their levels shown on the probability scale bar on the bottom right corners, and the bins representing the corresponding densities around regions. As can be identified, we can argue that the localisation is robust and remains unaltered regardless of the choice of d , however the uncertainty levels on pixels identification are higher. This is more obvious when $d = 0.45$, where in most neighbourhoods the density is more spread out than clearly focused on the correct pixels containing the molecules, also supported by the scale bar declaring lower probability levels. Applying a reconstruction setup based on the probability heat-maps could lead to either over-counting or missing active molecules, similar to what we addressed and discussed in Section 5.3.1. We discuss about relaxing the fact that d is a constant in Chapter 6.

Lastly, we revisit the subset from the stack of frames used to implement the Markov

switching model based on SiMPa functions in Section 4.4, displayed in Figure 4.4.1.




Figure 5.4.20: Prior and Posterior distributions of I , τ_b and c for Frame 338 from 'Circle within circle' data (displayed in Figure 5.2.1) using different proportions of the background intensity, that is $d = 0.45$ and $d = 0.20$ on the left and right panel respectively. The histograms along with the blue solid line represent the posterior samples with the red solid lines being the corresponding prior distributions.

This subset consists of nine frames, that is F_t , $t = 401, \dots, 409$ from the 'Circle within circle' synthetic data set, and we use in order to investigate the effect of different prior distributions on the transition matrix ξ . As a reminder, ξ contains the probabilities of a pixel interchanging between 'On'(1) and 'Off'(0) states, with ξ_{11} and ξ_{00} denoting the probabilities of a pixel remaining on the 'On' and 'Off' states respectively. Similar to Section 4.4, we consider different independent Beta prior distributions for ξ_{11} and ξ_{00} , displayed in Figure 5.4.22. The red solid lines

correspond to the prior distributions used on the implementation, concentrated on lower and higher probabilities for ξ_{11} and ξ_{00} respectively, hence promoting 'On' pixels to alter and 'Off' pixels to retain their states. In contrast, the blue solid lines represent the opposite scenario where higher probabilities are considered for 'On' pixels to retain and 'Off' pixels to alter their states. The prior distributions denoted with the light green solid lines represent a non-informative case where neither state is promoted.

In Figure 5.4.23, we present the posterior distributions of ξ_{11} and ξ_{00} , associated with their corresponding prior distributions, on a run of 25.000 MCMC iterations after a burn-in period of 10.000 iterations of the Markov switching model based on the SiMPa functions. Supported by the posterior distribution on the figure, we can safely argue that the prior choices for the transition matrix ξ do not affect or alter the performance of our model significantly, since the results are robust and consistent. The posterior distributions of the I , τ_b , c and $\{\beta_{0_t}, \beta_{f_t}\}$, $t = 401, \dots, 409$ are almost identical to the ones obtained during the implementation in Section 4.4, displayed in Figure 4.4.2 and in Appendix A.1 for $\{\beta_{0_t}, \beta_{f_t}\}$, as well as the localisation inference presented in Figures 4.4.3 and 4.4.4.

General conclusions on parameter sensitivities

Both the individual frame inference scheme and the Markov switching model based on the SiMPa functions are consistent and robust to the selection of prior distributions of the parameters I , τ_b , c , $\{\beta_0, \beta_f\}$'s and ξ when d is well defined. Since the single event intensity I and background proportion d are connected within our models, one should be careful if an elicitation procedure is responsible for their determination. The fact that d is a constant, can affect the localisation as a result of a knock on effect on the parameter estimation, when values far away from the truth are selected. This corresponds to regions in the posterior probability heat-maps having higher levels of uncertainty, a matter that can cause over-counting and/or missing active molecules when a reconstruction setup is applied. Again, we outline the fact that both I and d are observable quantities which can be well-established during a STORM experiment.




Figure 5.4.21: Posterior probability maps filtered by the average field for Frame 338 from 'Circle within circle' (displayed in Figure 5.2.1), using different proportions of the background intensity d . On the left panel of the figure, d is higher than the true value with $d = 0.45$, while on the right panel lower with $d = 0.20$.




Figure 5.4.22: Prior distributions for the parameters ξ_{11} and ξ_{00} of the transition matrix ξ used for sensitivity analysis. The red and blue solid lines correspond to more concentrated prior distributions around specific probabilities, whereas the light green solid lines denote less informative prior distributions.




Figure 5.4.23: Prior and Posterior distributions for the parameters ξ_{11} and ξ_{00} of the transition matrix ξ by applying the Markov switching model based on SiMPa functions on a subset of frames from 'Circle within circle' data (displayed in Figure 4.4.1). The histograms along with the blue solid line represent the posterior samples with the red solid lines being the corresponding prior distributions.

5.5 Application on realistic data

In this section we apply our individual frame inference scheme on realistic data used for the challenge in 2013, in order to evaluate and assess current single molecule localisation methods [Sage et al., 2015]. The chosen dataset represents Tubulins with various diameters, consisting of a long sequence of 10000 frames with low density (LDLS), where mostly well separated PSFs of active molecules are present on individual frames. For this implementation, we consider a subset of the total $F = 10000$

frames in the middle part of the sequence consisting of 100 frames, that is F_t for $t \in \{502, 522, 542, \dots, 2502\}$. Our choice is based on the idea that for such a long sequence, big steps (here we consider a step of 20 frames) will most probably allow to focus on different molecules undergo blinking and not ones that retain their 'On' state over consecutive frames.

In Figure 5.5.1 some of the F_t , $t \in \{502, 522, 542, \dots, 2502\}$, frames are displayed, which are part of the application and in agreement with Section 5.3, a complete analysis using the individual frame inference based on SiMPa functions is presented. Every frame in the sequence is represented by a $n = 128 \times 128 = m$ lattices of pixels consisting of intensities coming from either only background or parts of one or more PSFs. In order to have similar intensities to our implementations throughout this thesis we rescale the values by dividing them by 100.

Based on a visual and preliminary investigation of the frames, we can potentially argue that the majority of active molecules are captured by their well separated PSFs, with a few probably overlapping in some cases, for instance in the middle part of Frame 1142 or the left middle area of Frame 1222. Additionally, regardless of the noise levels which we discuss next, different large regions within the frames appear to contain higher intensities without implying the existence of active molecules. This can be potentially a result of non-uniform background or out of focus light, as we discussed in Chapter 2, and could complicate our individual frame inference when we identify the regions of interest based on the MRF. Lastly, the production of such long sequences of low density, translates into longer acquisition times hence the probability of capturing active molecules during their state transition is increased, a fact that can lead frames to contain events of lower intensities. These matters need to be taken into consideration and are discussed next where we describe the prior setting.




Figure 5.5.1: Selected frames from a subset of the realistic LDLS Tubulins data used for application. The frames are arbitrarily chosen and appear to contain both well isolated and overlapping PSFs from active molecules. The intensity levels are described by the grey-scale colours with white representing higher values.

Prior setting

In order to obtain prior distributions for the single event intensity I and background noise precision τ_b , we follow a similar procedure as the elicitation procedure described in Section 5.2, whereas we keep the same rescaled beta prior with higher density on values away from zero for the power of diffraction parameter c , i.e. $\pi(c) \sim \text{RescaledBeta}(4, 1.75)$ and independent Gaussian distributions for the field parameters β_0 and β_f , i.e. $\pi(\beta_0), \pi(\beta_f) \sim \text{N}(0, 9)$. The choice for c promotes large values, hence faster decay of the diffraction based on the SiMPa functions, with values between -2 and 2 to serve the purpose regarding β_0 & β_f .




Figure 5.5.2: Quantities used for prior elicitation - realistic LDLS Tubulins data. Different measurements obtained by using three different quantiles on each one of the frames in the subset of the realistic LDLS Tubulins used for the application. The top histogram consists of the larger values, hence is assumed to contain the overlapping scenarios, whereas the middle and bottom one consist of the single event case and background baseline respectively. The vertical red solid lines denote the 10% and 90% quantiles of the corresponding sets.

For elicitation, we use the sequence F_t , $t \in \{502, 522, 542, \dots, 2502\}$, and the lower quantile q_{lwr} corresponds to the 99.5% of the observations, the middle quantile q_{inter} between 99.5% and 99.925% and q_{upr} above 99.925%, accounting for the hypothetical background, single event and overlapping intensities respectively. As we discussed in Section 5.2, we consider higher values for these quantiles since we deal with a long sequence of low density data. In Figure 5.5.2, the obtained sets of observations Q_{lwr}^* , Q_{inter}^* and Q_{upr}^* are presented with the red solid vertical lines being their 10% and 90% quantiles respectively. Based on these sets, for the single event intensity I we obtain $\pi(I) \sim N(8.6, 1.91)$ with $d = 0.25$, and for the background error precision τ_b , $\pi(\tau_b) \sim \text{Gamma}(2, 0.94)$, with the prior distributions displayed in Figure 5.5.3 in the red coloured solid lines. Also, the light green solid line denotes a very precise prior distribution for τ_b , which is the one we consider for this application and discuss next.




Figure 5.5.3: Priors elicited or chosen from the realistic LDLS Tubulins data. On the top left corner of the figure, the prior of the single event intensity I is presented with I_0 being the background intensity, whereas on the bottom left corner the prior of the background precision τ_b is shown, as selected to be a very precise prior centred on a value above 1 (here at μ_{τ_b}). The two chosen prior distributions for the spread c and field control parameters β 's are indicated on the right panel.

Application

The LDLS Tubulins data fall within the category we considered as 'Large noise' case during the sensitivity analysis in Section 5.4.1, a fact supported by the large variability in the hypothetical background intensities in Figure 5.5.2. If we take into account the mean, $\mu_I = 8.6$, of the hypothetical single event intensity I , and the background intensity with a mean around 2.2 varying between 0 and 5, we can argue we have a case fairly similar to the diffraction of background levels, when $I = 8$ and the proportion $d = 0.3$, hence the background intensity is $I_0 = 2.4$, under a large scale noise. This can be found at the bottom right panel of Figure 5.4.1. As thoroughly discussed in Section 5.4.1, large noise levels can noticeably affect and alter the capabilities of our localisation algorithms, since the intensity levels of the pixels can be altered substantially, mostly complicating cases with overlapping PSFs.

Even though there is higher uncertainty on inference, we showed that, regardless of the overlapping scenarios, both our counting and localisation procedures are capable of constantly identifying all regions of interest when complete events are present. Complete events correspond to an active molecule diffraction described as a result of the single event intensity I , thus single emitters produce similar signal on individual frames. This is not always the case in the LDLS sequence. If we assume that $\mu_I = 8.6$ (as elicited), there are a number of occasions where captured PSFs have a diffraction of lower or weaker signal, with two potential reasons being that the noise levels have crucially altered the diffraction or incomplete events are present possibly during a state transition. The latter is directly related with the blinking rate of the respective fluorophore and the camera's frame rate, as we discussed in Chapter 2. Our proposed model does not include a mechanism to identify incomplete events, however a potential extension is discussed in Chapter 6, where we consider a local parameter q to account for the proportion of captured events in terms of the single event intensity I .

In preliminary runs of individual frames from the LDLS Tubulins, using the prior setting with the red solid lines in Figure 5.5.3, we discovered that the combination

of large noise and potentially incomplete events within frames can lead to an incapability of our model to identify every region of interest, potentially resulting in only capturing the high intensity regions as ones with active molecules. Specifically, the posterior distribution of the background error precision τ_b is centred on a small value close to zero, less than 0.5, (translating into large variance), making either lower intensity pixels most probable to be identified as 'Off', or highly improbable to cycle back to 'On' from the 'Off' state. Also, apart from missing regions of interest, another effect could be to create unnecessary large 'On' patches or islands around the determined regions, that can make inference less accurate.

These should not be the cases if the posterior of τ_b was not stuck on such small values, also a possible consequence of high complexity within a frame, even when incomplete events are present. We note here that, regardless of the noise levels, if an incomplete event has a diffraction with intensities closer to background, there is a high probability to be identified as 'Off' on average through our inference. The posterior probability maps and the average fields for the chosen frames in Figure 5.5.1, when the elicited prior distribution for τ_b is considered, can be found in Appendix A.4. For this application we consider a very precise prior distribution for τ_b on a value away from zero, here this is the elicited mean μ_{τ_b} , displayed on the left bottom panel of Figure 5.5.3 with the light green colour, however any very precise prior centred on values above 1 can serve the purpose. Such an adjustment forces τ_b to move away from small values close to zero, accomplishing correct identification of every region of interest even when incomplete events exist, however we should note here that our model is still functional and provides inference on the determined regions (Appendix A.4). As already mentioned, we discuss potential extensions and improvements of our proposed model that can potentially alleviate these issues in Chapter 6.

Similar to the application on the 'Circle within circle' synthetic data in Section 5.3.1, we apply the individual frame counting scheme in a total of 25000 MCMC iterations after 5000 iterations as a burn-in period. We obtain the posterior probability maps filtered by the average field for the F_t frames, $t \in \{502, 522, \dots, 2502\}$, with the chosen

frames (Figure 5.5.1) displayed in Figure 5.5.4, indicating the probabilities of pixels containing a single molecule.

Alongside the posterior probability map, the double event posterior probability maps (also filtered by the average field) are obtained, accounting for two simultaneously active molecules on the same pixel. For the chosen frames (Figure 5.5.1) these maps are displayed in Figure 5.5.5. The posterior distributions of the parameters I , c , τ_b and $\{\beta_0, \beta_f\}$, and the average fields for the chosen frames are omitted from the main body and can be found in Appendix A.4.

The posterior probability maps associated with the double events heatmap constitute the positional inference to identify pixels containing active molecules. In Figure 5.5.4, the capability of our model to determine such pixels is apparent, with regions that are potentially a result of isolated molecules to be resolved with high probabilities (around 0.75) compared to more complex ones with overlapping molecules (around 0.50). The probability levels can drop even lower in cases where either large patches are created by the field (see Appendix A.4 for the average fields) or more complicated overlapping situations occur, probably ones violating the assumption of no more than two PSFs allowed to overlap.

Additionally, since the LDLS Tubulins are subject to large noise levels, we expect to have a more variable field hence larger uncertainty on the localisation inference, a matter we thoroughly investigated in Section 5.4.1. Briefly, large noise could substantially alter the intensity levels of pixels with the consequence of preventing our model to identify complete PSFs, thus decreasing the precision in identification of pixels containing active molecules, more apparent in overlapping PSFs (see Figure 5.4.6 or 5.4.10) and even sometimes in well isolated PSFs (Figure 5.4.4). This is an effect we identify in the results of this application, more clear during the reconstruction steps that follow, with potential consequences being over-counting and/or missing molecules.




Figure 5.5.4: Posterior probability maps filtered by the average field for frames in Figure 5.5.1. Higher probabilities are indicated with white color levels as shown in the light blue probability scale bar on each probability map. The white bins represent the density around the regions.




Figure 5.5.5: Double event posterior probability maps filtered by the average field for frames in Figure 5.5.1. Higher probabilities are indicated with white color levels as shown in the light blue probability scale bar on each probability map. The white bins represent the density around the regions.

In agreement with Section 5.3.1, construction of a super resolution image can be performed by applying a threshold on the posterior probabilities, determining the pixels containing a single active molecule, followed by a threshold on the double event probability map to further identify if the drawn positions contain two active molecules. The corresponding localisation of molecules will be fixed at the center of the pixels (origin of the SiMPa functions), i.e. $r = \theta = 0$, as we applied the individual frame counting scheme. We consider the same two schemes to threshold our probability maps, with the first being common thresholds for all frames in the sequence and the second a conventional scheme that takes into account a larger neighbourhood within their average fields to determine regions with higher complexity.

For the fixed threshold scheme, we keep the same levels, that is 0.45 (lower) and 0.50 (intermediate), as in Section 5.3.1, which are intermediate values and can potentially resolve both well separated and overlapping PSFs. According to the sensitivity analysis on localisation in Section 5.4.1, regarding the overlapping, or not, scenarios, as well as the stability of the field, we expect such values to lead both in over-counting and missing active molecules in some of the cases. In Figure 5.5.6, we present the frame-wise reconstructions for the chosen frames in Figure 5.5.1, using the common thresholds of 0.45 and 0.50, where on each individual frame the red empty and yellow filled circles correspond to the localisations of molecules with $r = \theta = 0$ respectively.

We focus on Frames 902 and 1222, which appear to contain different intensity regions with different overlapping scenarios, in order to discuss the performance of the fixed thresholds. Regarding potential overlapping regions, the patch on the middle left area of Frame 902 is identified as being a product of two simultaneously active molecules in close proximity for the lower threshold, not the case for the intermediate one. On the left area of Frame 1222 for instance, there seems to be four different overlapping regions, however under both these thresholds only the top one determines molecule pixels. These thresholds appear to be quite high for such regions hence failing to place molecules, a consequence outlined in Section 5.4.1 when large noise was investigated.




Figure 5.5.6: Individual frame reconstructions for a number of frames included in the application on realistic LDLS Tubulins data, using two different thresholds for the posterior probability maps filtered by the average field. The red circles denote the drawn positions by thresholding the posterior probabilities filtered by the average field at 0.45, whereas the yellow coloured points at 0.5.

The right top corner and the bottom areas of Frames 902 and 1222 respectively, appear to consist of a number of potential well isolated PSFs. It is quite clear that in almost every case, both the low and intermediate thresholds lead to possible over-counting, a matter that was also the case in the application in Section 5.3.1 for fixed thresholds. In an attempt to have a common threshold capturing active molecules on entire frames, isolated PSFs which can be accurately resolved by higher thresholds, lead to over-counting since there is higher uncertainty under large noise (details in Section 5.4.1).

In Figure 5.5.7, we present the individual frame reconstructions for the chosen frames in Figure 5.5.1, using a conventional threshold scheme based on the average field, with the drawn positions indicated with the orange coloured circles. In Table 5.5.1 we present the conventional threshold scheme, where some adjustments have been made compared to the one used in the application on the 'Circle within circle' synthetic data (Table 5.3.1 in Section 5.3), and we discuss next.

As a reminder, the idea behind the conventional threshold scheme is to provide a broad classification on the complexity of regions of interest, in terms of indicating potential overlapping or not scenarios by focusing on the states of pixels in a larger neighbourhood. Since the LDLS Tubilins sequence is subject to large noise levels, we alter the structure of the scheme as the fields during the MCMC are generally less stable (average fields for chosen frames can be found in Appendix A.4). Similarly as before, $\bar{\mathbf{x}}_t = \sum_{q=1}^{IT} x_t^{(q)} / IT$ denotes the average field of frame t where IT the number of MCMC iterations and $\bar{x}_{s,t}$ a pixel $s \in \{1, \dots, N\}$ on the average field of frame t . Then, $\sum_s \partial \bar{x}_{s,t}$ denotes the sum of states in the 5×5 neighbourhood of pixel s in the average field t . Compared to the conventional threshold scheme for the 'Circle within circle', the limits for $\sum_s \partial \bar{x}_{s,t}$ have been decreased to account for the instability of the fields, along with the threshold values due to the higher uncertainty in the regions of interest under large noise levels. Additionally, the first three rows in Table have been added to deal with incomplete PSFs, here we assume a complete PSF is described

by a 3×3 pixel region, a case more frequent under large noise (for instance Figure 5.4.4).

Applying the conventional threshold scheme, the reconstruction in Figure 5.5.7 appears to improve in terms of not placing some of the potential unnecessary molecules especially in regions where most probably single molecules exist, while also placing molecules in more challenging regions which were before omitted due to the larger thresholds. However, this convention is not ideal since it could still incorrectly misplace or include molecules in the reconstruction, with different trade-offs in precision when altering either the limits of $\sum_s \partial \bar{x}_{s,t}$ or the threshold levels. Obtaining an optimal threshold is beyond the scope of this thesis, however we discuss the possibility of an adaptive threshold in Chapter 6.

Convention based on 5×5 neighbourhoods of average field	Threshold
$\sum_s \partial \bar{x}_{s,t} \leq 4$	0.95
$\sum_s \partial \bar{x}_{s,t} \geq 5$ & $\sum_s \partial \bar{x}_{s,t} < 7$	0.75
$\sum_s \partial \bar{x}_{s,t} = 7$	0.60
$\sum_s \partial \bar{x}_{s,t} \geq 8$ & $\sum_s \partial \bar{x}_{s,t} < 10$	0.55
$\sum_s \partial \bar{x}_{s,t} \geq 10$ & $\sum_s \partial \bar{x}_{s,t} < 13$	0.45
$\sum_s \partial \bar{x}_{s,t} \geq 13$ & $\sum_s \partial \bar{x}_{s,t} < 16$	0.40
$\sum_s \partial \bar{x}_{s,t} \geq 16$	0.35

Table 5.5.1: Conventional threshold scheme based on the average field for application on realistic LDLS Tubulins. The quantity $\sum_s \partial \bar{x}_{s,t}$ denotes the sum of states in the 5×5 neighbourhood of pixel s in the average field t , associated with conventional conditions to acquire a threshold level.

Lastly, in Figures 5.5.8 and 5.5.9 we present the final super resolution images by combining the individual frame reconstructions for the chosen frames in Figure 5.5.1 and the entire sequence of frame F_t , $t \in \{502, 522, \dots, 2502\}$, respectively, where we also considered a common threshold of 0.55.




Figure 5.5.7: Individual frame reconstructions for a number of frames included in the application on realistic LDLS Tubulins data, using the conventional threshold scheme based on the average field (Table 5.3.1) for the posterior probability maps filtered by the average field. The orange coloured circles denote the drawn positions by applying the conventional threshold scheme.




Figure 5.5.8: Final reconstructions of the super resolution image for the chosen frames in Figure 5.5.1, using three fixed thresholds for the posterior probability maps filtered by the average field and the conventional threshold scheme based on the average field (Table 5.3.1). On the top panel of the figure, the 0.45 (left - red colour) and 0.5 (right - yellow colour) fixed threshold reconstructions are displayed, with the 0.55 (green colour) and the conventional threshold (orange colour) shown on the bottom left and right corners respectively.




Figure 5.5.9: Final reconstruction of the super resolution image using the subset of the realistic LDLS Tubulins data, using three fixed thresholds for the posterior probability maps filtered by the average field and the conventional threshold scheme based on the average field (Table 5.3.1). On the top panel of the figure, the 0.45 (left - red colour) and 0.5 (right - yellow colour) fixed threshold reconstructions are displayed, with the 0.55 (green colour) and the conventional threshold (orange colour) shown on the bottom left and right corners respectively.

ThunderSTORM/SiMPa

Finally, we use ThunderSTORM [Ovesný et al., 2014] on the subset of frames using single PSF fitting, similar to Section 5.3.2, in order to qualitatively compare the results against the individual counting scheme based on the SiMPa functions. In Figure 5.5.10, we present the final reconstruction by combining the individual localisations using the conventional threshold scheme on the posterior probabilities (left panel - orange coloured circles) along with the final reconstruction using ThunderSTORM (right panel - yellow coloured circles).




Figure 5.5.10: Final reconstruction of the super resolution image for the subset of the realistic LDLS Tubulins data using SiMPa and ThunderSTORM. For the SiMPa, the conventional threshold scheme (Table 5.3.1) is applied on the posterior probabilities filtered by the average fields, displayed on the left panel (orange coloured circles), while the reconstruction using ThunderSTORM on single PSF fitting are displayed on the right panel (yellow coloured circles).

Generally, the overall structure of the underlying Tubulins is similarly reconstructed using both SiMPas and ThunderSTORM, with a few differences easily spotted in some occasions. Based on the SiMPa functions, the structure appears to consist of thicker lines compared to ThunderSTORM in which they are sharper, potentially a result of the simplification in our probabilistic counting scheme which assumes that every molecule lies at the center of pixels ($r = \theta = 0$).

Regarding molecule identification, two main differences can be spotted on the final reconstructions. Firstly, on the top panel of the images, the two main formed lines appear to extend a lot more according to SiMPas, with ThunderSTORM appearing to fail on identifying these regions as regions of interest. SiMPa finds several active molecules, whereas ThunderSTORM completely misses a few potentially active molecules in the corresponding areas. In contrast, ThunderSTORM seems to localising a few molecules at the right hand side of the second line from the bottom, with the SiMPas leaving a small gap. Secondly, as we thoroughly discussed in Section 5.5, even though the conventional threshold scheme appears to improve on both unnecessary placement of molecules (possible over-counting) and missing molecules, the large noise levels often reduce the accuracy of our schemes due to larger uncertainty (see Section 5.4.1). As a result, SiMPa can possibly over-count in some regions that appear to be a product of single active molecules, but seems much more powerful regarding potential overlapping PSFs, a matter that can be a consequence for the unidentified areas with ThunderSTORM. This can be supported by our comparisons in the 'Circle within circle' synthetic data in Section 5.3.2.

Some of the above differences between SiMPa and ThunderSTORM regarding the reconstructions, can be seen in the individual reconstructions in Figure 5.5.11 for the chosen frames in the LDLS Tubulins data (from Figure 5.5.1), using both the conventional threshold scheme and ThunderSTORM with single molecule fitting. The former is denoted by the orange whereas the latter with the yellow coloured circles.




Figure 5.5.11: Individual frame reconstructions for a number of frames included in the application on realistic LDLS Tubulins data, using the conventional threshold scheme based on the average field (Table 5.3.1) for the posterior probability maps filtered by the average field (left panel - orange colour) and ThunderSTORM using single PSF fitting (right panel - yellow colour).

5.6 Application on SuReSim data and comparisons

In this section we first apply our individual frame counting scheme on realistic data generated by SuReSim [Venkataramani et al., 2016], a software which simulates 3-D structures for localisation microscopy based on an underlying set truth, and compare the results with ThunderSTORM [Ovesný et al., 2014]. For the former, the procedures (prior elicitation/inference) and output based on the SiMPa functions are in complete agreement with Sections 5.3 and 5.5, where for the latter, ThunderSTORM is adjusted to the nature of the simulated data and camera’s setup accordingly. The evaluation and comparison of the results is performed both qualitatively and quantitatively.

SuReSim data generation

The generated dataset represents a 3-D structure of Microtubules where the ground truth of molecules’ positions is known. Most of the default settings for the Microtubules structure were maintained, where only a few of them were altered to meet the purpose of our proposed model based on SiMPa functions, i.e. high density data where PSF overlap is highly possible. Therefore, the number of frames in the generated stack was reduced to 5000 (instead of 10000) and the option to ensure single PSF’s was deselected. Even though we apply our individual frame inference scheme, hence we do not have a mechanism to account for blinking events over consecutive frames, we selected to allow potential contribution of blinking molecules to multiple frames, however, increased the minimum photon count of blinking. Relaxing the latter corresponds to having proportions of the single event intensity I for different molecules and is discussed in Chapter 6. Finally, the window size for PSF rendering was selected to be of 3×3 pixels, similar to the neighbourhood region of the SiMPa functions, again a choice that can be relaxed and has been introduced in Section 3.2.2 and discussed in Chapter 6.

In general, the simulated SuReSim Microtubules stack of $F = 5000$ frames consists of high density data, where both well separated and overlapping PSF’s are present,




Figure 5.6.1: Selected frames from SuReSim data. The frames are arbitrarily chosen and appear to contain both well isolated and overlapping PSFs from active molecules. The intensity levels are described by the grey-scale colours with white representing higher values. The light blue crosses correspond to the true positions of the molecules.

either representing entire events of higher intensities or reappearances of blinking events of potential lower intensities. Every frame in the sequence is represented by a $n = 34 \times 40 = m$ lattice of pixels consisting of intensities coming from either only background or parts of one or more PSFs (pixel to nm ratio for the stack is set to 133 by default). In order to have similar intensities to our implementations throughout this thesis we rescale the values by dividing them by 100. For this implementation, we consider a subset of the sequence of $F = 5000$ frames, that is F_t for $t \in \{2, 52, 102, \dots, 5002\}$ (every 50th frame starting at frame 2), some of which are displayed in Figure 5.6.1, for which the output from the individual SiMPa frame inference will be presented.

Based on a visual and preliminary investigation of the frames we can identify active molecules both captured by their well separated PSFs and overlapping ones, with the potential difference in intensity being apparent for different molecules. This translates into PSFs having either brighter or relatively darker diffraction, with a number of them being almost visually inseparable from background. Since the true position of molecules is known, we can also see that the number of molecules that overlap per region exceed two, one of the main assumptions for the current SiMPa model. An extension is straightforward and discussed in Chapter 6. These matters can potentially affect both localisation and parameter estimation, therefore will be taken into consideration and described next during the application.

Prior setting

Following a similar procedure as described in Section 5.2, we elicit prior distributions for the single event intensity I and background noise precision τ_b , whereas we again keep the rescaled beta prior for the power of diffraction c , i.e. $\pi(c) \sim \text{RescaledBeta}(4, 1.75)$, and independent Gaussian distributions for the field parameters β_0 and β_f , i.e. $\pi(\beta_0), \pi(\beta_f) \sim \text{N}(0, 9)$. As before, the choice for c promotes large values, translating into faster decay of the diffraction, and values between -2 and 2 to serve the purpose regarding β_0 & β_f about the behaviour of the fields.




Figure 5.6.2: Quantities used for prior elicitation - SuReSim data. Different measurements obtained by using three different quantiles on each one of the frames in the subset of the SuReSim data used for the application. The top histogram consists of the larger values, hence is assumed to contain the overlapping scenarios, whereas the middle and bottom one consist of the single event case and background baseline respectively. The vertical red solid lines denote the 10% and 90% quantiles of the corresponding sets.

For elicitation, we use the sequence F_t , $t \in \{1, 31, 51, \dots, 2501\}$ and based on a preliminary analysis, the lower quantile q_{lwr} corresponds to the 75% of the observations, the middle quantile q_{inter} between 75% and 99.995% and the upper q_{upr} above 99.95%, accounting for the hypothetical background, single event and overlapping intensities respectively. As we discussed in Section 5.2, these quantiles are considered to deal with such high density data. In Figure 5.5.2, the obtained sets of observations Q_{lwr}^* , Q_{inter}^* and Q_{upr}^* are presented with the red solid vertical lines being their 10% and 90% quantiles respectively. Based on these sets, for the single event intensity I we obtain $\pi(I) \sim N(11, 5.4)$ with $d = 0.27$, and for the background error precision τ_b ,

$\pi(\tau_b) \sim \text{Gamma}(2, 2)$, with the prior distributions displayed in Figure 5.5.3 in the red coloured solid lines. Also, the light green solid line denotes a very precise prior distribution for τ_b , which is the one we consider for this application and discussed.




Figure 5.6.3: Priors elicited or chosen from SuReSim data. On the top left corner of the figure, the prior of the single event intensity I is presented with I_0 being the background intensity, whereas on the bottom left corner the prior of the background precision τ_b is shown, as selected to be a very precise prior centred on a value above 1 (here at μ_{τ_b}). The two chosen prior distributions for the spread c and field control parameters β 's are indicated on the right panel.

Application

The microtubules data fall within the category we considered as 'Regular noise' case during the sensitivity analysis in Section 5.4.1, a fact supported by the variability in the hypothetical background and single event intensities in Figure 5.5.2 which also

present similar behaviour with the Quantiles for 'Circle within circle' data in Section 5.2.2.

As we have showed in sensitivity analysis in Section 5.4, regardless of the noise levels and overlapping scenario that satisfies the assumption of up to two PSFs overlapping per region, both our counting and localisation procedures are capable of consistently identifying all regions of interest when complete events are present. As a reminder, complete events correspond to an active molecule diffraction described as a result of the single event intensity I , thus single emitters produce similar signal on individual frames. For the microtubules data this is not always the case due to how the data are generated, a fact also visually identified on frames in Figure 5.6.1. Alongside overlapping events, blinking events that can reappear in multiple frames might produce weaker signal from incomplete events.

Therefore, similar to the tubulins application in Section 5.5, and based on preliminary runs of individual frames from the microtubules data, using the prior setting with the red solid lines in Figure 5.6.3, we again discovered that a combination of extreme overlapping and/or potentially incomplete events within frames can lead to an incapability of our model to identify every region of interest. The impact of this can be potentially only capturing the high intensity regions as ones with active molecules. As with tubulins in Section 5.5, the posterior distribution of τ_b is centred around a small value close to zero, translating into large variability of intensities, making either lower intensity pixels most probable to be identified as 'Off' (empty).

As previously, a posterior distribution of τ_b that is not stuck on such small values, can alleviate this issue, hence for this application we consider a very precise prior distribution for τ_b on a value away from zero and above one, displayed on the left bottom panel of Figure 5.6.3 with the light green colour. We note here that incomplete events with an intensities diffraction closer to background, can also be potentially identified as 'Off' on average through our inference.




Figure 5.6.4: Posterior probability maps filtered by the average field for frames in Figure 5.6.1. Higher probabilities are indicated with white color levels as shown in the light blue probability scale bar on each probability map. The white bins represent the density around the regions.




Figure 5.6.5: Double event posterior probability maps filtered by the average field for frames in Figure 5.6.1. Higher probabilities are indicated with white color levels as shown in the light blue probability scale bar on each probability map. The white bins represent the density around the regions.

Similar to both the applications on the 'Circle within circle' synthetic and Tubulins realistic data in Sections 5.3.1 and 5.5 respectively, we apply the individual frame counting scheme in a total of 25000 MCMC iterations after 5000 iterations as a burn-in period. The resulting posterior probability maps filtered by the average field for frames F_t , $t \in \{2, 52, \dots, 2502\}$, with the chosen frames (Figure 5.5.1) displayed in Figure 5.6.4, indicating the probabilities of pixels containing a single molecule. Alongside those, in Figure 5.6.5 the double event posterior probability maps (also filtered by the average field) are obtained, accounting for two simultaneously active molecules on the same pixel. The posterior distributions of the parameters I , c , τ_b and $\{\beta_0, \beta_f\}$ and the average fields for the chosen frames in Figure 5.6.1 can be found in Appendix A.5.

Using the main positional inference tools, that is the posterior probability maps associated with the double events heat-map, we are able to identify pixels containing active molecules. Looking at Figure 5.6.4, the capability of our model to determine such pixels is apparent, with regions that are potentially a result of isolated molecules to be resolved with high probabilities (≈ 0.7) compared to more complex ones with overlapping molecules (≈ 0.45). The probability levels can drop even lower (≈ 0.3 - 0.35) in cases where either large patches are created by the field (see Appendix A.5 for the average fields) or more complicated overlapping situations occur, consisting mostly of violations of the assumption of no more than two PSFs allowed to overlap in a region.

As a result of this variability of the posterior probabilities, especially for complex regions and those that violate the assumption of no more than two PSFs overlapping, identifying molecules to reconstruct a super resolution image based on a thresholding scheme can have an effect of overcounting or missing molecules. Such processing of the posterior probability maps has an important role from a quantitative perspective. Of course, as previously has been discussed, incapability of localising molecules can be a result of creation of large islands in the field from many closely PSFs either overlapping or not. As thoroughly described in Sections 5.3 and 5.5, we consider

two different threshold approaches, these are (i) *fixed thresholds for all frames* and (ii) *a conventional threshold scheme based on average field*, each one having pros and cons. Generally, we prefer the use of the latter since it can reduce the number of false positives (incorrect identification of pixels with active molecules), however, can sometimes lead to missing true positives (correct pixels containing active molecules), as we will see later in this section. According to our individual frame counting scheme, both reconstruction approaches for individual frames consist of placing molecules at the center of the identified pixels (origin of SiMPa), that is when $r = \theta = 0$, followed by stacking up all reconstructions on a single frame.

Mainly based on preliminary analysis of the results from SiMPa, we consider the fixed threshold levels for all frames to be 0.40 and 0.45. These are both 0.05 lower than the ones applied to 'Circle within circle' data in Section 5.3 that have similar noise levels but differ in number of potentially overlapping PSFs and reappearances of blinking events. This is done to allow potential identification of active pixels for more complex regions, but at the same time increases the possibility of overcounting. For the same reasons, for the conventional threshold scheme, we maintain the same conditions for the average field as in Table 5.3.1, while we increase each one by 0.05 with the last one starting at 0.40. The conventional threshold scheme for this application on microtubules data can be seen in Table 5.6.1. As a reminder, this scheme classifies the regions of the average field based on the state of a larger neighbourhood of 5×5 pixels. Both of these approaches can be improved, mainly by having a dynamic threshold, as discussed in Chapter 6.

Convention based on 5×5 neighbourhoods of average field	Threshold
$\sum_s \partial \bar{x}_{s,t} < 12$	0.65
$\sum_s \partial \bar{x}_{s,t} \geq 12$ & $\sum \partial \bar{x}_{s,t} < 15$	0.5
$\sum_s \partial \bar{x}_{s,t} \geq 15$ & $\sum \partial \bar{x}_{s,t} < 18$	0.45
$\sum_s \partial \bar{x}_{s,t} > 18$	0.4

Table 5.6.1: Conventional threshold scheme based on the average field for SuReSim data. The quantity $\sum \partial \bar{x}_{s,t}$ denotes the sum of states in the 5×5 neighbourhood of pixel s in the average field t , associated with conventional conditions to acquire a threshold level.

In Figure 5.6.6, the individual frame reconstructions based on the conventional threshold scheme for the chosen frames in Figure 5.6.6 are displayed. From a first glance, we can argue that well isolated complete events or ones that had their PSFs overlapping while satisfying the assumption can be accurately identified, a case not necessarily true for incomplete events or scenarios of violation in overlap. The final reconstructions of the super resolution image using both fixed and conventional thresholds is presented in Figure 5.6.7 along with the ground truth. An overall performance of these schemes for all frames F_t , $t \in \{1, 31, 51, \dots, 2501\}$, in the application can be found later in Table 5.6.2, where we compare and discuss results with ThunderSTORM [Ovesný et al., 2014].

Similar to our other two applications, we present in Figure 5.6.8, the individual precision-recall-precision curves for each one of the chosen frames, serving as diagnostic metrics. These curves correspond to the positive predictability power (precision), against positive detection power (recall), over multiple different thresholds levels in order to capture the true molecule positions. As already mentioned, it is apparent that the threshold levels should not be globally defined but need to follow some sophisticated rules in order to be efficient (discussion in Chapter 6). The associated sensitivity-specificity plots can be found in Appendix A.5, with specificity being the proportion of correctly identified inactive pixels.




Figure 5.6.6: Individual frame reconstructions for frames in Figure 5.6.1 included in the application on SuReSim data, using the conventional threshold scheme based on the average field for the posterior probability maps filtered by the average field. The orange coloured circles denote the drawn positions by applying the conventional threshold, with the light blue coloured crosses representing the true positions of the molecules.




Figure 5.6.7: Final reconstruction of the super resolution image using the subset of frames of the 'Circle within circle' stack of frames, using two fixed thresholds for the posterior probability maps filtered by the average field and the conventional threshold scheme based on the average field. On the top left corner of the figure, the true positions of the molecules are shown (light blue crosses) with the 0.4 (red colour) and 0.45 (yellow colour) fixed threshold final reconstructions to be presented on the right top and left bottom corner respectively. On the bottom right corner of the figure, the final reconstruction using the conventional threshold scheme is displayed (orange colour).




Figure 5.6.8: Precision-Recall curves for SuReSim data frames in Figure 5.6.1. The lines show the performance of our model in terms of balancing the trade-off between precision and recall, with the adapting colour denoting different threshold values. Frames 1102 and 1852 initiate the curve at 0 because the higher threshold does not correspond to a TP, but a FP, not the case for the rest of the frames.

ThunderSTORM

Lastly, we again use ThunderSTORM [Ovesný et al., 2014] on the same subset of frames F_t , $t \in \{2, 52, \dots, 2502\}$, using single and multiple PSF fitting, as described in Section 5.3.2, in order to both quantitatively and qualitatively compare the results against the individual frame counting scheme based on the SiMPa functions. We note here that the ThunderSTORM algorithm was preliminary implemented a number of times to acceptably adjust the parameters in a way to perform at full capacity. The camera setup was identically determined to be the one used during microtubules data generation, while the fitting window was set to be of size 5×5 pixels where ThunderSTORM performs better, even though the window size for PSF rendering during data generation was 3×3 to be in agreement with SiMPa's neighbourhood.

In Figure 5.6.9, we present the final reconstruction by combining the individual localisations using the conventional threshold scheme on the posterior probabilities (top panel - orange coloured circles) along with the final reconstructions using ThunderSTORM with single PSF fitting (bottom left panel - yellow coloured circles) and multiple PSF fitting (bottom right panel - yellow coloured circles). Additionally, the ground truth is displayed with the light blue crosses on the top left panel of the figure. From a qualitative point of view, both SiMPa and ThunderSTORM seem to identify the general structure, with ThunderSTORM producing sharper areas with molecules appearing to be in closer distance with each other. For the single PSF fitting the reconstruction clearly appears to be visually poor, especially regarding the number of identified molecules. SiMPa produces thicker regions of molecules, a result of the simplification during counting of having $r = \theta = 0$.

From a quantitative perspective, we present in Table 5.6.2 the same metrics used throughout this thesis for comparisons. As a reminder, the abbreviation TP, FP, TN and FN denote the true positives (correct prediction of pixel containing a molecule), false positives (false prediction of pixel containing a molecule), true negatives (correct prediction of empty pixel) and false negatives (false prediction of empty pixel) respectively. Two different colours, black and blue, have been used, with the former

corresponding to counting when the real reappearances of molecules due to blinking are taken into account or not respectively. Focusing on the latter, since the individual frame counting scheme has been applied, we can claim that the SiMPa based model has a better detection power (TP) which can almost reach $\approx 16\%$ improvement compared to ThunderSTORM, while keeping the same percentage of incorrect pixels (FP) and correct identification of empty ones (TN). This can be increased by either adjusting the conventional threshold scheme in a different way or in that case choosing the common threshold of 0.45 for all frames, however with the cost of increasing the FP. As a general comment, SiMPa seems to be noticeably better than ThunderSTORM, even though SiMPa requires minimal information about the design and details of the experiment compared to ThunderSTORM. Additionally, extensions and improvement of the current form of SiMPas are straightforward to implement and can potentially allow for tackling scenarios and cases that are now problematic, such as incomplete events, extreme overlapping etc (discussed in Chapter 6).

Counting	TP	FP	TN	FN				
SiMPa - ConvT	57.2%	58.6%	32.8%	33.5%	90.3%	90.3%	42.8%	41.4%
SiMPa - T0.40	64.8%	66.2%	52.9%	54.2%	86.9%	86.9%	35.2%	33.8%
SiMPa - T0.45	50.3%	51.7%	26.8%	27.2%	91.1%	91.1%	49.7%	48.3%
Thunder	28.7%	29.5%	13.9%	14.2%	94.6%	94.6%	71.3%	70.5%
Thunder+	41.9%	42.8%	29.2%	29.9%	90.6%	90.6%	58.1%	57.2%
Real	1754	1700	-	-	925	925	-	-

Table 5.6.2: Performance of ThunderSTORM and individual frame inference based on SiMPa functions for SuReSim data. where 'Thunder' denotes single fitting, 'Thunder+' multiple fitting, 'SiMPa - Conv' the conventional threshold scheme applied on the individual frame counting scheme based on SiMPa functions and 'SiMPa - T0.40' & 'SiMPa - T0.45' fixed thresholds of levels 0.4 & 0.45 respectively. The table outlines the percentages of correct identified pixels containing the active molecules (TP) and the number of correct identified empty pixels (TN), along with the incorrect positive predictions (FP) and negative (FN). The black and blue colours denote counting when reappearances of molecules due to blinking are either considered or not respectively. The real number of pixels containing active molecules and empty pixels are also outlined (Real).




Figure 5.6.9: Final reconstructions of the super resolution image on SuReSim data, using SiMPa and ThunderSTORM. The conventional threshold scheme for the posterior probability maps filtered by the average field has been applied for SiMPa, whereas ThunderSTORM is considered on both single and multiple fitting. On the top left corner of the figure, the true positions of the molecules are shown (light blue crosses) with the conventional threshold scheme (orange colour) displayed on the right top corner. On the bottom right panel of the figure, the final reconstructions using ThunderSTORM (yellow colour) based on single molecule (left bottom panel) and multiple fitting (right bottom panel) are presented.

Chapter 6

Conclusion and Discussion

In this thesis, we propose a novel approach to process SMLM images, based on a novel set of structural functions (SiMPa), that model the intensity diffraction along a predetermined neighbourhood with respect to the position of a molecule within a pixel. We introduce a new concept within a Bayesian framework, based on the pattern-configuration-realisation scheme that allows for probabilistic inference on the positions of the molecules. Its main feature includes a moving region which scans frames associated with the corresponding structural pattern of a hidden mechanism determining the unobserved state of the pixels. Setting an allowance of overlapping PSFs, every pattern has a finite number of PSFs that can be formed from, hence for every moving region potential positions can be drawn in a probabilistic manner. Cycles of this procedure can result in localisations of active molecules not only when well isolated PSFs are focused, but also when overlapping situations occur from molecules in close proximity.

The SiMPa functions based model, associated with the pattern - configuration - realisation scheme, was extended in order to allow processing on complete stacks of frames as obtained by a SMLM imaging experiment, while accounting for time dynamics. A stack of frames, considered as a state space model, progresses via a

Markov switching model while taking into account the corresponding fluorophore transition properties as well as the spatial dependencies. A forward filtering backwards sampling (FFBS) algorithm was considered for updating the unobserved time related states of the pixels, whereas the molecules spatial separation was achieved using the SiMPa functions within the proposed probabilistic scheme.

Considering a data generation mechanism based on the SiMPa functions, we obtained a synthetic data - stack of frames in which we applied our individual inference scheme, presenting a complete inference and features in order to evaluate and reconstruct super resolution images. We showed that our model is robust on any changes of the parameter prior distributions, and additionally retain its powerful capability to resolve overlapping PSFs, or more generally high intensity regions containing multiple PSFs, especially when the allowance of overlapping PSFs is met. We outlined the fact that large noise levels, that can alter substantially the intensity of pixels, can have a negative effect on localisation adding more uncertainty to the results. We compared our model against a popular off-the-shelf alternative (ThunderSTORM) using the synthetic data set, in which our model appears to be superior from a quantitative perspective, while offering similar visual reconstructions.

Lastly, we applied our individual frame inference scheme on a realistic data set of Tubulins, consisting of a long sequence of low density data, and on simulated SuReSim dataset of high density data to evaluate and compare our proposed methodology. Next, we discuss potential extensions and future work that can either improve and enhance our current procedures or enable them to deal with super resolution data sets of different properties, i.e. high density data where multiple PSFs could overlap with each other.

Computing times - Implementation on GPUs

Our probabilistic inference schemes, either counting or localisation, appear to be quite powerful in resolving overlapping PSFs and high intensity regions, however require a number of iterations to achieve this since the schemes are embedded within

a Bayesian framework. In this thesis, we mainly focused on the counting algorithm which is less computationally intensive. Furthermore, for the Markov switching model based on SiMPa functions we only performed an implementation along with a parameter sensitivity analysis on a small subset of a stack, since it requires even more computational power and memory.

For every model, the codes have been written in R using vectorization, with the frame-wise counting and localisation algorithms applied in parallel using the University of Sheffield cluster. For individual frames, the computational time is independent of its dimensions, however is directly dependent on the complexity of the regions and consequently the number of active molecules, leading to a necessary trade-off in order to resolve overlapping PSFs. On a laptop with i7-6500U CPU @ 2.50GHz & 16G RAM, computation of 30.000 MCMC iterations need,

Complexity	Counting	Localisation
Multiple High intensity regions ; ~ 20 molecules	1.5 ⁻ days	2 ⁺ days
Multiple Overlap (>5 cases) ; ~ 20 molecules	1.5 ⁻ days	2 ⁺ days
Regular Overlap (<5 cases) ; ~ 20 molecules	1.25 ⁻ days	2 ⁻ days

Table 6.0.1: Computational times for individual frame inference based on SiMPa functions, on a total of 30000 MCMC iterations. For the localisation algorithm the realisations come from a discretisation, with possible directions $\theta = k\frac{\pi}{4}$, $k=0,1,\dots,8$ under the distances $r = 0$ and $r = \frac{r_{\max,\theta}}{2}$ (2 distances over 8 angles).

We plan on optimising our algorithms in R, and then create an R package implementing our localisation schemes. A straightforward way to parallelise the code is distribute the localisation or counting schemes into separate CPUs, since our corresponding functions use as input the number of 'On' pixels. Here, we considered a 3×3 moving region, thus the possible 'On' pixels are 1-9, allowing for splitting the code into 9 CPUs. One of the main future plans is to transfer the procedures into GPUs, so computational speed will increase rapidly, and potentially create an

add-on for ImageJ as a user friendly interface.

Improvement of localisation scheme

In order to increase the accuracy and precision on localisation of molecules, with respect to a discretisation of r and θ , and also obtain a faster and less complex localisation scheme, a two stage procedure could be considered. Specifically, our probabilistic counting scheme relies on the pattern - configuration concept where every molecule is assumed to be fixed at the center of pixels, or the origin of the SiMPa functions with $r = \theta = 0$, while the localisation scheme additionally allows all possible realisations (in the 1st quadrant) on each configuration. For that reason, inference with the latter relies on a large number of probabilities for each moving window, hence making it less stable and more computational expensive (Section 3.3.3). Therefore, with a potential two stage procedure within our inference, one could have a mechanism to obtain the pixels containing the active molecules before proceeding to localisation, instead of applying a simultaneous step that takes into account all possible configurations associated with the corresponding realisations.

Relaxation of assumption & Extended SiMPa (SiMPa_X)

The main assumption we made throughout this thesis, reads in allowing no more than two PSFs overlapping on a single region. Specifically, let a moving region landing in a central pixel where the active molecule lies into. When considering the SiMPa functions that correspond to a 3×3 neighbourhood of pixels, we assume that this particular moving region cannot have an extra two PSF contributing but one up to one extra one. We showed that, even under violations of this assumption inference on such region is obtained, however there might be cases with inaccurate identification and/or missing molecules. The main reasons for this allowance, is firstly related to the nature of SMLM imaging, where mostly temporally separated molecules are captured, and secondly, reduces the complexity and computational times of our model. As we discussed before, a proper implementation on GPUs can allow a straightforward extension of this assumption, probably a matter of high importance since high density data are in the forefront of interest. Regarding data

sets that appear to contain molecules diffusing in a larger neighbourhood, we plan on applying the extended SiMPa model (SiMPa_X), also needing more computational power.

Background proportion d (or background intensity I_0)

Throughout this thesis, we assumed the background, I_0 , and single event, I , intensities are related, based on the argument that the single event intensity is a quantity added on top of some background levels, hence I_0 can be considered proportional to I . The proportionality relationship was described by the constant d , which when multiplied with I gives the background intensity. Two main extensions can be potentially considered, one accounts for allowing d to be a parameter, maybe a-priori following a Beta distribution, and another assuming no relationship between I and I_0 . Both are straightforward to implement, with the former relaxing the need to have an accurate estimate of the background intensity proportion, even though it is an observable quantity during the imaging experiment.

Proportion of event parameter q

In order to capture incomplete events, that is active molecules with lower emission intensity than ordinary, an additional local parameter q can be introduced, accounting for the proportion of an event captured on the frame in terms of the total intensity levels. This parameter can be potentially embedded in both the individual frame and Markov switching models. In the former, this can help identifying regions of interest that could be otherwise missed, especially in cases where the background I_0 and single event intensity I are closely valued, or under large noise cases. Apart from this, in the latter, this can also be used to determine active molecules retaining their 'On' state between frames, according to the relation and synchronization between frame (camera settings) and blinking (duty cycle of molecule) rates. This parameter would be multiplied by the mean $\mu(x_{ij})$ ('On' state only) in Equation (3.3.6), where throughout this thesis we can say we assumed $q = 1$ hence focusing on complete

events, therefore

$$\mu(x_{ij}) = \begin{cases} q I (1 - \alpha_O(r, \theta, c)), & \text{if } x_{ij} = 1 \text{ ('On')} \\ I d, & \text{if } x_{ij} = -1 \text{ ('Off')} \end{cases}.$$

Background error precision τ_b

We can allow $\tau_{\text{MR}}(x_{ij})$ to depend on the 'On' and 'Off' pixels, instead of assuming that each pixel regardless of its state has a common background error τ_b , generated by a zero mean Gaussian distribution with precision τ_b , hence for an 3×3 MR - $\tau_{\text{MR}}(x_{ij}) = \tau_b/N_p$ (Equation (3.3.7)). For instance, let τ_{Off} be the precision of an 'Off' pixel (or variance V_{Off}) and τ_{On} the precision of an 'On' pixel (or variance V_{On}), then if n_{Off} and n_{On} are the number of 'Off' and 'On' pixels in the MR respectively,

$$\tau_{\text{MR}}(x_{ij}) = \frac{(\tau_{\text{On}}/n_{\text{On}})(\tau_{\text{Off}}/n_{\text{Off}})}{(\tau_{\text{On}}/n_{\text{On}}) + (\tau_{\text{Off}}/n_{\text{Off}})},$$

obtained by $\tau_{\text{MR}}(x_{ij}) = 1/V_{\text{MR}}(x_{ij}) = 1/(n_{\text{Off}}V_{\text{Off}} + n_{\text{On}}V_{\text{On}})$. Such an extension can potentially allow our models to have a more stable performance on localisation, especially cases of large noise levels.

Adaptive threshold scheme for reconstructions

Finally, in order to reconstruct super resolution images using the posterior probability maps, we considered either different common thresholds for every frame regardless of the complexity of individual regions within, or a conventional threshold scheme based on a mechanism using the average field on each frame to distinguish between high and regular intensities regions. Even though these schemes provide insightful reconstructions, there is definitely room for improvement. For future work, we will try to construct an adaptive threshold scheme that can take into account not only neighbourhoods in the field, but also potentially the probability densities or the total intensity of the regions with respect to the single event intensity I .

Appendix A

Appendix

This Appendix contains a number of figures that are omitted from the main body of the thesis.

Part [A.1](#) corresponds to the implementation of the Markov switching model based on SiMPa in Section [4.4](#).

Part [A.2](#) corresponds to the application of the individual frame inference based on the SiMPa functions to the synthetic data set of a 'Circle within circle' in Section [5.3.1](#).

Part [A.3](#) corresponds to the sensitivity analysis in Section [5.4](#).

Part [A.4](#) corresponds to the application of the individual frame inference based on the SiMPa functions to realistic data in Section [5.5](#).

Part [A.5](#) corresponds to the application of the individual frame inference based on the SiMPa functions to SureSim data in Section [5.5](#).

A.1 Figures from implementation of MSM SiMPa model (Sec 4.4)




Figure A.1.1: Prior and posterior distributions for the frame wise parameters β_0 's for frames in Figure 4.4.1. The red solid lines denote the prior distributions for each β_{0_t} , $t = 401, \dots, 409$ while the histograms associated with the blue solid lines denote the corresponding posterior distributions.




Figure A.1.2: Prior and posterior distributions for the frame wise parameters β_f 's for frames in Figure 4.4.1. The red solid lines denote the prior distributions for each β_{f_t} , $t = 401, \dots, 409$ while the histograms associated with the blue solid lines denote the corresponding posterior distributions.

A.2 Figures - Application on synthetic data (Sec 5.3)




Figure A.2.1: Average fields for 'Circle within circle' frames in Figure 5.3.1. The larger proportions a pixel is identified as 'On' are declared with the light blue color levels as shown in the scale bar on each frame. Each figure represents the average state of each frame after the end of the MCMC.




Figure A.2.2: Sensitivity-Specificity (ROC) curves for 'Circle within circle' frames in Figure 5.3.1. The lines show the performance of our model in terms of balancing the trade-off between sensitivity and specificity, with the adapting colour denoting different threshold values.




Figure A.2.3: Posterior distributions of single event intensity I for frames in Figure 5.3.1 presented in the application on 'Circle within circle' synthetic data. The histograms represent the drawn posterior samples whereas the red solid lines the prior distributions.




Figure A.2.4: Trace plots of the drawn posterior sample of the single event intensity I for frames in Figure 5.3.1 presented in the application on 'Circle within circle' synthetic data.




Figure A.2.5: Posterior distributions of background error τ_b for frames in Figure 5.3.1 presented in the application on 'Circle within circle' synthetic data. The histograms represent the drawn posterior samples whereas the red solid lines the prior distributions.




Figure A.2.6: Trace plots of the drawn posterior sample of the background error precision τ_b for frames in Figure 5.3.1 presented in the application on 'Circle within circle' synthetic data.




Figure A.2.7: Posterior distributions of power of spread parameter c for frames in Figure 5.3.1 presented in the application on 'Circle within circle' synthetic data. The histograms represent the drawn posterior samples whereas the red solid lines the prior distributions.




Figure A.2.8: Trace plots of the drawn posterior sample of the power of spread parameter c for frames in Figure 5.3.1 presented in the application on 'Circle within circle' synthetic data.




Figure A.2.9: Posterior distributions of field control parameters β_0 and β_F for frames in Figure 5.3.1 presented in the application on 'Circle within circle' synthetic data. The histograms represent the drawn posterior samples whereas the red solid lines the prior distributions.




Figure A.2.10: Trace plots of the drawn posterior sample of the field control parameter β_0 for frames in Figure 5.3.1 presented in the application on 'Circle within circle' synthetic data.




Figure A.2.11: Trace plots of the drawn posterior sample of the field control parameter β_f for frames in Figure 5.3.1 presented in the application on 'Circle within circle' synthetic data.

A.3 Figures - Sensitivity analysis (Sec 5.4)




Figure A.3.1: Prior/Posterior distributions obtained for the parameters τ_b , c and β_0 & β_f for Frame 338, using different prior distributions for I . The left column contains the posterior distributions when the elicited prior is used for I , whereas the middle and right columns the cases of larger variance and wrong location respectively. The histograms represent the posterior samples whereas the red solid lines the prior distributions. The light blue histogram on the bottom part relates to β_f while the regular one to β_0 .




Figure A.3.2: Prior/Posterior distributions obtained for the parameters I , c and β_0 & β_f for Frame 338, using different prior distributions for τ_b . The left column contains the posterior distributions when the elicited prior is used for τ_b , whereas the middle and right columns the cases of larger variance and wrong location respectively. The histograms represent the posterior samples whereas the red solid lines the prior distributions. The light blue histogram on the bottom part relates to β_f while the regular one to β_0 .




Figure A.3.3: Prior/Posterior distributions obtained for the parameters I , τ_b and β_0 & β_f for Frame 338, using different prior distributions for c . The left column contains the posterior distributions when the elicited prior is used for c , whereas the middle and right columns the cases of larger variance and wrong location respectively. The histograms represent the posterior samples whereas the red solid lines the prior distributions. The light blue histogram on the bottom part relates to β_f while the regular one to β_0 .




Figure A.3.4: Prior/Posterior distributions obtained for the parameters I , τ_b and c for Frame 338, using different prior distributions for β_0 and β_f . The left column contains the posterior distributions when the elicited prior is used for β_0 and β_f , whereas the right column the case of larger variance. The histograms represent the posterior samples whereas the red solid lines the prior distributions.




Figure A.3.5: Posterior probability maps filtered by the average field for Frame 338 used on parameter sensitivity. Higher probabilities are indicated with white color levels as shown in the light blue probability scale bar on each probability map. The white bins represent the density around the regions.




Figure A.3.6: Precision-Recall curves under regular and large noise levels for the frames used in sensitivity analysis, including both the counting and localisation procedures. The line types denote the corresponding noise cases whereas the adaptive color the trade-off in precision-recall for different thresholds on the posterior probabilities.




Figure A.3.7: Sensitivity-Specificity (ROC) curves under small and free of noise levels for the frames used in sensitivity analysis, including both the counting and localisation procedures. The line types denote the corresponding noise cases whereas the adaptive color the trade-off in sensitivity-specificity for different thresholds.




Figure A.3.8: Frame 57 - Posterior probability map filtered by the average field for specified locations within the pixels, according to our localisation algorithm based on SiMPa functions. Every pixel has been expanded into a 3×3 super-pixel region, each one denoting a combination of the applied discretisation of r and θ .




Figure A.3.9: Frame 6 - Posterior probability map filtered by the average field for specified locations within the pixels, according to our localisation algorithm based on SiMPa functions. Every pixel has been expanded into a 3×3 super-pixel region, each one denoting a combination of the applied discretisation of r and θ .




Figure A.3.10: Frame 62 - Posterior probability map filtered by the average field for specified locations within the pixels, according to our localisation algorithm based on SiMPa functions. Every pixel has been expanded into a 3×3 super-pixel region, each one denoting a combination of the applied discretisation of r and θ .




Figure A.3.11: Frame 23 - Posterior probability map filtered by the average field for specified locations within the pixels, according to our localisation algorithm based on SiMPa functions. Every pixel has been expanded into a 3×3 super-pixel region, each one denoting a combination of the applied discretisation of r and θ .

A.4 Figures - Application on realistic data (Sec 5.5)




Figure A.4.1: Posterior distributions of single event intensity I for frames in Figure 5.5.1 presented in the application on realistic LDLS Tubulins data. The histograms represent the drawn posterior samples whereas the red solid lines the prior distributions.




Figure A.4.2: Posterior distributions of background error precision τ_b for frames in Figure 5.5.1 presented in the application on realistic LDLS Tubulins data. The histograms represent the drawn posterior samples whereas the red solid lines the prior distributions.




Figure A.4.3: Posterior distributions of power of spread parameter c for frames in Figure 5.5.1 presented in the application on realistic LDLS Tubulins data. The histograms represent the drawn posterior samples whereas the red solid lines the prior distributions.




Figure A.4.4: Posterior distributions of field control parameters β_0 and β_F for frames in Figure 5.5.1 presented in the application on realistic LDLS Tubulins data. The histograms represent the drawn posterior samples whereas the red solid lines the prior distributions.




Figure A.4.5: Average fields for frames in Figure 5.5.1. The larger proportions a pixel is identified as 'On' are declared with the light blue color levels as shown in the scale bar on each frame. Each figure represents the average state of each frame after the end of the MCMC.




Figure A.4.6: Posterior probability maps filtered by the average field for frames in Figure 5.5.1, when the elicited prior distribution for τ_b is used. Higher probabilities are indicated with white color levels as shown in the light blue probability scale bar on each probability map. The white bins represent the density around the regions.




Figure A.4.7: Average fields for frames in Figure 5.5.1, when the elicited prior distribution for τ_b is used. The larger proportions a pixel is identified as 'On' are declared with the light blue color levels as shown in the scale bar on each frame. Each figure represents the average state of each frame after the end of the MCMC.

A.5 Figures - Application on SuReSim data (Sec 5.6)




Figure A.5.1: Posterior distributions of single event intensity I for frames in Figure 5.6.1 presented in the application on SuReSim data. The histograms represent the drawn posterior samples whereas the red solid lines the prior distributions.




Figure A.5.2: Posterior distributions of background error precision τ_b for frames in Figure 5.6.1 presented in the application on SuReSim data. The histograms represent the drawn posterior samples whereas the red solid lines the prior distributions.




Figure A.5.3: Posterior distributions of power of spread parameter c for frames in Figure 5.6.1 presented in the application on SuReSim data. The histograms represent the drawn posterior samples whereas the red solid lines the prior distributions.




Figure A.5.4: Posterior distributions of field control parameters β_0 and β_F for frames in Figure 5.6.1 presented in the application on SuReSim data. The histograms represent the drawn posterior samples whereas the red solid lines the prior distributions.




Figure A.5.5: Average fields for frames in Figure 5.6.1. The larger proportions a pixel is identified as 'On' are declared with the light blue color levels as shown in the scale bar on each frame. Each figure represents the average state of each frame after the end of the MCMC.




Figure A.5.6: Sensitivity-Specificity (ROC) curves for SuReSim data frames in Figure 5.6.1. The lines show the performance of our model in terms of balancing the trade-off between sensitivity and specificity, with the adapting colour denoting different threshold values.

Bibliography

- E. Abbe. A contribution to the theory of the microscope and the nature of microscopic vision. *Proc. Bristol Nat. Soc.*, 1:200–261, 1874.
- A. V. Abraham, S. Ram, J. Chao, E. Ward, and R. J. Ober. Quantitative study of single molecule location estimation techniques. *Optics Express*, 17(26):23352–23373, 2009.
- W. Ambrose, T. Basché, and W. Moerner. Detection and spectroscopy of single pentacene molecules in ap-terphenyl crystal by means of fluorescence excitation. *The Journal of chemical physics*, 95(10):7150–7163, 1991.
- S. B. Andersson. Localization of a fluorescent source without numerical fitting. *Opt. Express*, 16(23):18714–18724, Nov 2008. doi: 10.1364/OE.16.018714. URL <http://www.opticsexpress.org/abstract.cfm?URI=oe-16-23-18714>.
- I. M. Antolovic, S. Burri, C. Bruschini, R. A. Hoebe, and E. Charbon. Spad imagers for super resolution localization microscopy enable analysis of fast fluorophore blinking. *Scientific reports*, 7:44108, 2017.
- H. Babcock, Y. M. Sigal, and X. Zhuang. A high-density 3d localization algorithm for stochastic optical reconstruction microscopy. *Optical Nanoscopy*, 1(1):6, 2012.
- D. Baddeley, M. B. Cannell, and C. Soeller. Three-dimensional sub-100 nm super-resolution imaging of biological samples using a phase ramp in the objective pupil. *Nano Research*, 4(6):589–598, 2011.
- J. Besag. Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society: Series B (Methodological)*, 36(2):192–225, 1974.
- E. Betzig, G. H. Patterson, R. Sougrat, O. W. Lindwasser, S. Olenych, J. S. Bonifacino, M. W. Davidson, J. Lippincott-Schwartz, and H. F. Hess. Imaging intracel-

- lular fluorescent proteins at nanometer resolution. *Science*, 313(5793):1642–1645, 2006.
- M. Billio, A. Monfort, and C. P. Robert. Bayesian estimation of switching arma models. *Journal of econometrics*, 93(2):229–255, 1999.
- N. Brede and M. Lakadamyali. Graspj: an open source, real-time analysis package for super-resolution imaging. *Optical Nanoscopy*, 1(1):11, 2012.
- S. Cox. Super-resolution imaging in live cells. *Developmental biology*, 401(1):175–181, 2015.
- S. Cox, E. Rosten, J. Monypenny, T. Jovanovic-Talisman, D. T. Burnette, J. Lippincott-Schwartz, G. E. Jones, and R. Heintzmann. Bayesian localization microscopy reveals nanoscale podosome dynamics. *Nature methods*, 9(2):195, 2012.
- H. Cramér. *Mathematical methods of statistics (PMS-9)*, volume 9. Princeton university press, 2016.
- J. C. Croney, D. M. Jameson, and R. P. Learmonth. Fluorescence spectroscopy in biochemistry: teaching basic principles with visual demonstrations. *Biochemistry and Molecular Biology Education*, 29(2):60–65, 2001.
- S. Culley, D. Albrecht, C. Jacobs, P. M. Pereira, C. Leterrier, J. Mercer, and R. Henriques. Quantitative mapping and minimization of super-resolution optical imaging artifacts. *Nature methods*, 15(4):263, 2018.
- G. T. Dempsey, J. C. Vaughan, K. H. Chen, M. Bates, and X. Zhuang. Evaluation of fluorophores for optimal performance in localization-based super-resolution imaging. *Nature methods*, 8(12):1027, 2011.
- H. Deschout, A. Shivanandan, P. Annibale, M. Scarselli, and A. Radenovic. Progress in quantitative single-molecule localization microscopy. *Histochemistry and cell biology*, 142(1):5–17, 2014.
- A. Doucet, A. Logothetis, and V. Krishnamurthy. Stochastic sampling algorithms for state estimation of jump markov linear systems. *IEEE Transactions on Automatic Control*, 45(2):188–202, 2000.
- C. Eggeling, J. Widengren, R. Rigler, and C. Seidel. Photobleaching of fluorescent dyes under conditions used for single-molecule detection: Evidence of two-step photolysis. *Analytical chemistry*, 70(13):2651–2659, 1998.

- T. Fawcett. An introduction to roc analysis. *Pattern recognition letters*, 27(8):861–874, 2006.
- T. H. Fereja, A. Hymete, and T. Gunasekaran. A recent review on chemiluminescence reaction, principle and application on pharmaceutical analysis. *Isrn Spectroscopy*, 2013, 2013.
- W. Fong, S. J. Godsill, A. Doucet, and M. West. Monte carlo smoothing with application to audio signal enhancement. *IEEE transactions on signal processing*, 50(2):438–449, 2002.
- S. Frühwirth-Schnatter. *Finite mixture and Markov switching models*. Springer Science & Business Media, 2006.
- S. Geman and D. Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. In *Readings in computer vision*, pages 564–584. Elsevier, 1987.
- T. J. Gould, V. V. Verkhusha, and S. T. Hess. Imaging biological structures with fluorescence photoactivation localization microscopy. *Nature protocols*, 4(3):291, 2009.
- M. G. Gustafsson. Nonlinear structured-illumination microscopy: wide-field fluorescence imaging with theoretically unlimited resolution. *Proceedings of the National Academy of Sciences of the United States of America*, 102(37):13081–13086, 2005.
- M. G. Gustafsson, L. Shao, P. M. Carlton, C. R. Wang, I. N. Golubovskaya, W. Z. Cande, D. A. Agard, and J. W. Sedat. Three-dimensional resolution doubling in wide-field fluorescence microscopy by structured illumination. *Biophysical journal*, 94(12):4957–4970, 2008.
- J. D. Hamilton. A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica: Journal of the Econometric Society*, pages 357–384, 1989.
- J. M. Hammersley and P. Clifford. Markov fields on finite graphs and lattices. *Unpublished manuscript*, 46, 1971.
- R. Heintzmann, T. M. Jovin, and C. Cremer. Saturated patterned excitation microscopy—a concept for optical resolution improvement. *JOSA A*, 19(8):1599–1609, 2002.
- S. W. Hell and J. Wichmann. Breaking the diffraction resolution limit by stimulated

- emission: stimulated-emission-depletion fluorescence microscopy. *Optics Letters*, 19(11):780–782, 1994.
- R. Henriques, M. Lelek, E. Fornasiero, F. Valtorta, C. Zimmer, and M. Mhlanga. Quickpalm: 3d real-time photoactivation nanoscopy image processing in imagej. *PLoS Medicine*, 7(5):339–340, 5 2010. ISSN 1549-1277. doi: 10.1038/nmeth0510-339.
- S. T. Hess, T. P. Girirajan, and M. D. Mason. Ultra-high resolution imaging by fluorescence photoactivation localization microscopy. *Biophysical journal*, 91(11): 4258–4272, 2006.
- K. L. Hey, H. Momiji, K. Featherstone, J. R. Davis, M. R. White, D. A. Rand, and B. Finkenstädt. A stochastic transcriptional switch model for single cell imaging data. *Biostatistics*, 16(4):655–669, 2015.
- S. J. Holden, S. Uphoff, and A. N. Kapanidis. Daostorm: an algorithm for high-density super-resolution microscopy. *Nature methods*, 8(4):279, 2011.
- B. Huang, M. Bates, and X. Zhuang. Super-resolution fluorescence microscopy. *Annual review of biochemistry*, 78:993–1016, 2009.
- F. Huang, S. L. Schwartz, J. M. Byars, and K. A. Lidke. Simultaneous multiple-emitter fitting for single molecule super-resolution imaging. *Biomedical optics express*, 2(5):1377–1393, 2011a.
- F. Huang, T. M. Hartwich, F. E. Rivera-Molina, Y. Lin, W. C. Duim, J. J. Long, P. D. Uchil, J. R. Myers, M. A. Baird, W. Mothes, et al. Video-rate nanoscopy using sCMOS camera-specific single-molecule localization algorithms. *Nature methods*, 10(7):653, 2013.
- Z.-L. Huang, H. Zhu, F. Long, H. Ma, L. Qin, Y. Liu, J. Ding, Z. Zhang, Q. Luo, and S. Zeng. Localization-based super-resolution microscopy with an sCMOS camera. *Optics Express*, 19(20):19156–19168, 2011b.
- F. Humblot and A. Mohammad-Djafari. Super-resolution using hidden markov model and bayesian detection estimation framework. *EURASIP Journal on Advances in Signal Processing*, 2006(1):036971, 2006.
- G. Hummer, F. Fricke, and M. Heilemann. Model-independent counting of molecules in single-molecule localization microscopy. *Molecular biology of the cell*, 27(22): 3637–3644, 2016.

- I. Izeddin, J. Boulanger, V. Racine, C. Specht, A. Kechkar, D. Nair, A. Triller, D. Choquet, M. Dahan, and J. Sibarita. Wavelet analysis for single molecule localization microscopy. *Optics express*, 20(3):2081–2095, 2012.
- A. Jablonski. Efficiency of anti-stokes fluorescence in dyes. *Nature*, 131(3319):839–840, 1933.
- S. Karlin and H. E. Taylor. *A second course in stochastic processes*. Elsevier, 1981.
- S. M. Kay. *Fundamentals of statistical signal processing*. Prentice Hall PTR, 1993.
- A. Kechkar, D. Nair, M. Heilemann, D. Choquet, and J.-B. Sibarita. Real-time analysis and visualization for single-molecule based super-resolution microscopy. *PLoS One*, 8(4):e62918, 2013.
- H. Kirshner, F. Aguet, D. Sage, and M. Unser. 3-d psf fitting for fluorescence microscopy: implementation and localization application. *Journal of microscopy*, 249(1):13–25, 2013.
- U. Köthe, F. Herrmannsdörfer, I. Kats, and F. A. Hamprecht. Simplestorm: a fast, self-calibrating reconstruction algorithm for localization microscopy. *Histochemistry and cell biology*, 141(6):613–627, 2014.
- J. R. Lakowicz. *Principles of fluorescence spectroscopy*. Springer Science & Business Media, 2013.
- S. Z. Li. *Markov random field modeling in image analysis*. Springer Science & Business Media, 2009.
- J. W. Lichtman and J.-A. Conchello. Fluorescence microscopy. *Nature methods*, 2(12):910–919, 2005.
- H. Ma, F. Long, S. Zeng, and Z.-L. Huang. Fast and precise algorithm based on maximum radial symmetry for single molecule localization. *Opt. Lett.*, 37(13):2481–2483, Jul 2012. doi: 10.1364/OL.37.002481. URL <http://ol.osa.org/abstract.cfm?URI=ol-37-13-2481>.
- P. Meyer and J. Dworkin. Applications of fluorescence microscopy to single bacterial cells. *Research in microbiology*, 158(3):187–194, 2007.
- J. Min, C. Vonesch, H. Kirshner, L. Carlini, N. Olivier, S. Holden, S. Manley, J. C. Ye, and M. Unser. Falcon: fast and unbiased reconstruction of high-density super-resolution microscopy data. *Scientific Reports*, 4:4577, 2014.

- E. A. Mukamel and M. J. Schnitzer. Unified resolution bounds for conventional and stochastic localization fluorescence microscopy. *Physical review letters*, 109(16):168102, 2012.
- E. A. Mukamel, H. Babcock, and X. Zhuang. Statistical deconvolution for super-resolution fluorescence microscopy. *Biophysical Journal*, 102(10):2391 – 2400, 2012. ISSN 0006-3495. doi: <https://doi.org/10.1016/j.bpj.2012.03.070>. URL <http://www.sciencedirect.com/science/article/pii/S0006349512004122>.
- R. P. Nieuwenhuizen, K. A. Lidke, M. Bates, D. L. Puig, D. Grünwald, S. Stallinga, and B. Rieger. Measuring image resolution in optical nanoscopy. *Nature methods*, 10(6):557, 2013.
- D. Nino, N. Rafiei, Y. Wang, A. Zilman, and J. N. Milstein. Molecular counting with localization microscopy: a bayesian estimate based on fluorophore statistics. *Biophysical journal*, 112(9):1777–1785, 2017.
- H. Nyquist. Certain topics in telegraph transmission theory. *Transactions of the American Institute of Electrical Engineers*, 47(2):617–644, 1928.
- R. J. Ober, S. Ram, and E. S. Ward. Localization accuracy in single-molecule microscopy. *Biophysical journal*, 86(2):1185–1200, 2004.
- M. Ovesný, P. Křížek, J. Borkovec, Z. Švindrych, and G. M. Hagen. Thunderstorm: a comprehensive imagej plug-in for palm and storm data analysis and super-resolution imaging. *Bioinformatics*, 30(16):2389–2390, 2014.
- A. R. Pagan and G. W. Schwert. Alternative models for conditional stock volatility. *Journal of econometrics*, 45(1-2):267–290, 1990.
- R. Parthasarathy. Rapid, accurate particle tracking by calculation of radial symmetry centers. *Nature Methods*, 9:724–726, 2012.
- T. Quan, S. Zeng, and Z. Huang. Localization capability and limitation of electron-multiplying charge-coupled, scientific complementary metal-oxide semiconductor, and charge-coupled devices for superresolution imaging. *Journal of biomedical optics*, 15(6):066005, 2010.
- T. Quan, H. Zhu, X. Liu, Y. Liu, J. Ding, S. Zeng, and Z.-L. Huang. High-density localization of active molecules using structured sparse model and bayesian information criterion. *Optics Express*, 19(18):16963–16974, 2011.
- L. R. Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.

- C. R. Rao. Information and the accuracy attainable in the estimation of statistical parameters. In *Breakthroughs in statistics*, pages 235–247. Springer, 1992.
- G. C. Rollins, J. Y. Shin, C. Bustamante, and S. Pressé. Stochastic approach to the molecular counting problem in superresolution microscopy. *Proceedings of the National Academy of Sciences*, 112(2):E110–E118, 2015.
- R. Rottenfusser, E. E. Wilson, and M. W. Davidson. The point spread function. URL <https://www.zeiss.com/microscopy/int/solutions/reference/basic-microscopy/the-point-spread-function.html>.
- M. J. Rust, M. Bates, and X. Zhuang. Sub-diffraction-limit imaging by stochastic optical reconstruction microscopy (storm). *Nature methods*, 3(10):793, 2006.
- D. Sage, H. Kirshner, T. Pengo, N. Stuurman, J. Min, S. Manley, and M. Unser. Quantitative evaluation of software packages for single-molecule localization microscopy. *Nature methods*, 12(8):717, 2015.
- M. Sauer, J. Hofkens, and J. Enderlein. *Handbook of fluorescence spectroscopy and imaging: from ensemble to single molecules*. John Wiley & Sons, 2010.
- M. C. Seiler and F. A. Seiler. Numerical recipes in c: The art of scientific computing. *Risk Analysis*, 9(3):415–416, 1989. doi: 10.1111/j.1539-6924.1989.tb01007.x. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1539-6924.1989.tb01007.x>.
- C. E. Shannon. Communication in the presence of noise. *Proceedings of the IRE*, 37(1):10–21, 1949.
- A. Shivanandan, H. Deschout, M. Scarselli, and A. Radenovic. Challenges in quantitative single molecule localization microscopy. *FEBS letters*, 588(19):3595–3602, 2014.
- H. Shroff, C. G. Galbraith, J. A. Galbraith, and E. Betzig. Live-cell photoactivated localization microscopy of nanoscale adhesion dynamics. *Nature methods*, 5(5):417, 2008.
- A. Small and S. Stahlheber. Fluorophore localization algorithms for super-resolution microscopy. *Nature methods*, 11(3):267, 2014.
- A. R. Small and R. Parthasarathy. Superresolution localization methods. *Annual Review of Physical Chemistry*, 65(1):107–125, 2014. doi: 10.1146/annurev-physchem-040513-103735. URL <https://doi.org/10.1146/annurev-physchem-040513-103735>. PMID: 24274701.

- C. S. Smith, N. Joseph, B. Rieger, and K. A. Lidke. Fast, single-molecule localization that achieves theoretically minimum uncertainty. *Nature methods*, 7(5):373, 2010a.
- J. F. Smith, A. Pillai, K. Chen, and B. Horwitz. Identification and validation of effective connectivity networks in functional magnetic resonance imaging using switching linear dynamic systems. *Neuroimage*, 52(3):1027–1040, 2010b.
- S. Stallinga and B. Rieger. Accuracy of the gaussian point spread function model in 2d localization microscopy. *Opt. Express*, 18(24):24461–24476, Nov 2010. doi: 10.1364/OE.18.024461. URL <http://www.opticsexpress.org/abstract.cfm?URI=oe-18-24-24461>.
- R. Starr, S. Stahlheber, and A. Small. Fast maximum likelihood algorithm for localization of fluorescent molecules. *Optics Letters*, 37(3):413–415, 2012.
- G. G. Stokes. Xxx. on the change of refrangibility of light. *Philosophical transactions of the Royal Society of London*, (142):463–562, 1852.
- T. Takeshima, T. Takahashi, J. Yamashita, Y. Okada, and S. Watanabe. A multi-emitter fitting algorithm for potential live cell super-resolution imaging over a wide range of molecular densities. *Journal of Microscopy*, 271(3):266–281, 2018.
- Y. Tang, L. Dai, X. Zhang, J. Li, J. Hendriks, X. Fan, N. Gruteser, A. Meisenberg, A. Baumann, A. Katranidis, et al. Snsml, a real-time single molecule identification and localization algorithm for super-resolution fluorescence microscopy. *Scientific Reports*, 5:11073, 2015.
- Y. Tang, J. Hendriks, T. Gensch, L. Dai, and J. Li. Automatic bayesian single molecule identification for localization microscopy. *Scientific Reports*, 6:33521, 2016.
- M. A. Thompson, J. S. Biteen, S. J. Lord, N. R. Conley, and W. Moerner. Molecules and methods for super-resolution imaging. In *Methods in enzymology*, volume 475, pages 27–59. Elsevier, 2010.
- R. E. Thompson, D. R. Larson, and W. W. Webb. Precise nanometer localization analysis for individual fluorescent probes. *Biophysical journal*, 82(5):2775–2783, 2002.
- N. J. Turro. *Modern molecular photochemistry*. University science books, 1991.
- B. Valeur. Molecular fluorescence. *Digital Encyclopedia of Applied Physics*, pages 477–531, 2003.

- V. Venkataramani, F. Herrmannsdörfer, M. Heilemann, and T. Kuner. Suresim: simulating localization microscopy experiments from ground truth models. *Nature methods*, 13(4):319, 2016.
- Y. Wang, T. Quan, S. Zeng, and Z.-L. Huang. Palmer: a method capable of parallel localization of multiple emitters for high-density localization microscopy. *Opt. Express*, 20(14):16039–16049, Jul 2012. doi: 10.1364/OE.20.016039. URL <http://www.opticsexpress.org/abstract.cfm?URI=oe-20-14-16039>.
- S. Wolter, U. Endesfelder, S. van de Linde, M. Heilemann, and M. Sauer. Measuring localization performance of super-resolution algorithms on very active samples. *Optics Express*, 19(8):7020–7033, 2011.
- S. Wolter, A. Löschberger, T. Holm, S. Aufmkolk, M.-C. Dabauvalle, S. Van De Linde, and M. Sauer. rapidstorm: accurate, fast open-source software for localization microscopy. *Nature methods*, 9(11):1040, 2012.
- B. Yu, D. Chen, J. Qu, and H. Niu. Fast fourier domain localization algorithm of a single molecule with nanometer precision. *Opt. Lett.*, 36(22):4317–4319, Nov 2011. doi: 10.1364/OL.36.004317. URL <http://ol.osa.org/abstract.cfm?URI=ol-36-22-4317>.
- B. Zhang, J. Zerubia, and J.-C. Olivo-Marin. Gaussian approximations of fluorescence microscope point-spread function models. *Appl. Opt.*, 46(10):1819–1829, Apr 2007. doi: 10.1364/AO.46.001819. URL <http://ao.osa.org/abstract.cfm?URI=ao-46-10-1819>.
- L. Zhu, W. Zhang, D. Elnatan, and B. Huang. Faster storm using compressed sensing. *Nature methods*, 9(7):721, 2012.