

Investigating Spatial Organization and Physicochemical Interactions in Biomembranes: Tools and Insights

by

Thomas R. Shaw

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Applied Physics)
in the University of Michigan
2023

Doctoral Committee:

Professor Sarah L. Veatch, Chair
Professor Julie S. Biteen
Professor David K. Lubensky
Associate Professor Kevin Wood
Visiting Professor Haoxing Xu

Thomas R. Shaw

shawtr@umich.edu

ORCID iD: 0000-0002-1619-2401

© Thomas R. Shaw 2023

Dedication

For Sarah, who stands by me in the face of chaos.

Acknowledgements

I thank Dr. Sarah Veatch for guiding me and supporting me throughout this work. Her enthusiastic scientific vision is an inspiration, and her probing questions always lead me toward better understanding. Her encouragement is steadfast, and her advice is always helpful, even when it is occasionally unwelcome. I could not have done this without you.

All of the members of the Veatch lab have contributed greatly to my success and to my happiness throughout this process. To Matt, Marcos, Julia, Kathleen, Abir, Joe, Anna, Frank, Jenny, Andrea, Nat, Yousef, Adam, and Iz – thank you all for all of the things you have taught me, the lab work you have shared and for making the lab always a pleasant place to hang out. Special thanks to Kathleen, with whom I shared many pleasant trips to get liquid nitrogen, and many pots of earl grey tea, beyond your training in fundamental lab techniques.

Thanks to Ben Machta. Your willingness to discuss my sometimes ill-informed questions about critical phenomena and many other topics is always enjoyable, and always leads me to better and clearer ideas.

I would also like to thank Rasmus Waagepetersen and Jesper Møller, who responded to my out-of-the-blue emailed statistical idea with a proposal to collaborate on my first publication. Their engagement with that idea and patience in getting me up to speed on the statistical machinery for analyzing point processes was an incredible learning experience. I hope we can collaborate more in the future.

Many friends have supported me through graduate school. Thanks to Darwin and Cesilie for all the seasons of soccer, picnics and camping, to the trivia crew for distracting evenings every week, and to the virtual bunker, who welcomed me into their transcontinental friendship and made the pandemic much more bearable. My family has also been so important in getting me through. It was especially helpful and fun to overlap in Ann Arbor for 4 years with Clara and Andrew. You eased me into Ann Arbor life and made me so many delicious dinner. Thank you.

Table of Contents

Dedication	ii
Acknowledgements	iii
List of Tables	x
List of Figures	xi
Abstract	xv
Chapter 1 Introduction	1
1.1 Biomembrane structure, physical chemistry, and spatial organization	1
1.2 Single molecule localization microscopy	4
1.3 Detecting spatial organization in SMLM datasets	8
1.4 Overview of the dissertation	12
Chapter 2 Critical Phenomena in Plasma Membrane Organization and Function	15
2.1 Overview	15
2.2 An early history of domains in model and cell membranes.	16
2.3 Liquid-liquid phase separation in model membranes	17
2.3.1 Purified membranes:	17
2.3.2 Isolated membranes:	19
2.4 Phase separated domains are related to but different from ‘raft’ heterogeneity in intact cells.	22
2.5 Criticality and its connection to plasma membrane	24
2.5.1 Evidence for criticality playing a role in cells:	27
2.6 Criticality as it relates to biological function	28
2.6.1 Interactions between proteins:	29

2.6.2 Susceptibility to receptor clustering:	29
2.6.3 Allosteric regulation of single proteins:	31
2.6.4 Tuning binding of allosteric regulators:	31
2.6.5 Criticality coupled to other processes describes a broad array of raft phenomena:	32
2.7 Concluding remarks	33
Chapter 3 Chemical Potential Measurements Constrain Models of Cholesterol-Phosphatidylcholine Interactions	34
3.1 Abstract	34
3.2 Introduction	35
3.3 Materials and Methods	36
Materials	36
3.3.1 Preparing M β CD solutions and saturated solutions of M β CD and M β CD/Chol.	37
3.3.2 Vesicle preparation.....	37
3.3.3 Equilibration and isolation of cholesterol in vesicles and aqueous M β CD solutions. .	38
3.3.4 Measurement of cholesterol and phospholipid concentration	39
3.3.5 Calibrating cholesterol activity in M β CD/Chol mixtures.	41
3.3.6 Evaluating best fit parameters of mean-field models:.....	42
3.3.7 Evaluating approximate phase boundaries and tie lines in mean field models:.....	43
3.3.8 Three-component lattice fluid simulations:.....	44
3.4 Results	45
3.4.1 Measurements of μ_c in binary mixtures with DOPC, POPC, and DPPC.....	45
3.4.2 Measurements of μ_c in ternary mixtures of DOPC, DPPC, and cholesterol.	50
3.4.3 Measurements of μ_c limit the binding affinity of cholesterol-phospholipid complexes in a condensed complex model.....	52
3.4.4 Phase separation and measurements of μ_c are described by models capturing pairwise interactions between components.....	54

3.5 Discussion	57
3.5.1 Experimental trends in μ_c can constrain simple models.....	58
3.5.2 Comparison to past results in model membranes and cells.....	59
3.6 Conclusions	60
3.7 Supplemental Figures and Tables.....	61
3.8 Supplemental Notes.....	73
3.8.1 Supplemental Note 1: computation of chemical potentials in a mole fraction basis....	73
3.8.2 Supplemental Note 2: Regular solution models.....	74
3.8.3 Supplemental Note 3: Mean field regular solution without complexes	75
3.8.4 Supplemental Note 4: mean field regular solution models with complexes	75
3.8.5 Supplemental Note 5: Estimating chemical potentials from lattice model simulations of the regular solution model without complexes.....	79
3.8.6 Supplemental Note 6: Comparison of fitted mean field regular solution model interaction energies to past estimates for a similar mixture.	81
Chapter 4 A Mean Shift Algorithm for Drift Correction in Localization Microscopy.....	82
4.1 Abstract	82
4.2 Introduction	82
4.3 Results	83
4.4 Methods and Materials	91
4.4.1 Simulated Datasets	91
4.4.2 Extracting Close Pairs of Coordinates Between Datasets	91
4.4.3 Determining Shifts between Translated Datasets Using a Mean Shift Algorithm.....	92
4.4.4 Estimates of mean shift error.....	93
4.4.5 Evaluating displacements using nonlinear least squares (NLLS) fitting	94
4.4.6 Correcting continuous drift	95
4.4.7 Evaluating performance of displacement algorithms.....	95

4.4.8 Evaluating the resolution of drift-corrected datasets.....	96
4.4.9 Preparation of cellular samples for imaging.....	96
4.4.10 Single molecule imaging and localization.....	97
4.5 Supplemental Figures.....	98
Chapter 5 Estimating the Localization Spread Function of Static Single Molecule Localization Microscopy Images.....	103
5.1 Abstract	103
5.2 Introduction	104
5.3 Materials and Methods	105
5.3.1 Simulations.....	105
5.3.2 Experimental sample preparation:.....	106
5.3.3 Single molecule imaging and localization.....	107
5.3.4 Evaluation of space-time autocorrelations.	108
5.3.5 Estimation of $g_{LSF}(r, \tau)$ and $\sigma_{xy}(\tau)$	108
5.3.6 Estimating $g_{LSF}(r, \tau)$ by grouping localizations with molecules	109
5.3.7 Measuring distances between distinct molecules on the same ruler.	109
5.3.8 Determining the resolution with Fourier Ring Correlation	109
5.3.9 Determining localization precision using Nearest Neighbor distributions	110
5.4 Results	110
5.4.1 Derivation of the estimated LSF	110
5.4.2 Practical limitations of the LSF estimate.	113
5.4.3 Validation through Simulation.....	115
5.4.4 Estimating the LSF of DNA origami datasets.....	119
5.4.5 Estimating the LSF from datasets of labeled structures in chemically fixed cells.....	122
5.4.6 Comparison to other measures of image resolution	125

5.5 Conclusions	128
5.6 Supplementary Figures.....	130
5.7 Supplementary Note: Derivation of spacetime pair correlation function estimator, and related computations.	141
Chapter 6 Globally Intensity-reweighted Estimators for K - and Pair Correlation Functions	144
6.1 Abstract	144
6.2 Introduction	144
6.3 Preliminaries.....	146
6.4 Global and local intensity-reweighted estimators for K -functions	148
6.4.1 The case of one spatial point process	148
6.4.2 Modifications to account for isotropy	149
6.4.3 Comparison of local and global estimators	150
6.4.4 The case of two spatial point processes.....	151
6.5 Global and local intensity-reweighted estimators for pair correlation functions	152
6.5.1 The case of one spatial point process	152
6.5.2 Modifications to account for isotropy	153
6.5.3 Two point processes	155
6.6 Sources of bias when ρ is estimated.....	156
6.6.1 Bias of local estimators with estimated ρ	156
6.6.2 Bias of global estimators with estimated γ	159
6.7 Computation of γ and γ_{iso}	161
6.8 Simulation study.....	162
6.8.1 Estimation of K - and pair correlation functions.....	164
6.8.2 Estimation of cross K - and cross pair correlation functions	167
6.8.3 Estimation of K -function using a parametric estimate for ρ	171

6.9 Extensions	175
6.10 Conclusion.....	175
Chapter 7 Spatial Pair Correlation Analysis of Localization Microscopy Data under Spatially Nonuniform Sampling	176
7.1 Abstract	176
7.2 Introduction	177
7.3 Results and Discussion.....	179
7.3.1 Naive correlation function estimate yields artifactual correlations.....	179
7.3.2 Local and global density corrections for cross correlation functions.....	180
7.3.3 Variance of global estimator is lower than that of local estimator.....	181
7.3.4 Estimating density corrections when density is unknown.....	182
7.3.5 Pairwise edge-correction for the normalization factor gamma reduces systematic errors	184
7.3.6 Short range density variation induces bias in kernel based gamma estimates, especially for large sigma.....	185
7.3.7 Density corrected estimators underestimate the strength of correlations, especially for small sigma.....	185
7.3.8 Application to an SMLM dataset	186
7.4 Conclusions	187
7.5 Materials and Methods	188
7.5.1 Simulated SMLM datasets	188
7.5.2 Computations for correlation functions.....	188
7.5.3 Sample preparation.....	189
7.5.4 Imaging and localization	189
Chapter 8 Conclusions	199
Bibliography	203

List of Tables

Table 3.1 Fitted parameter values and confidence intervals.....	72
Table 5.1 Summary of values obtained for several resolution measures for the datasets shown in figures	130
Table 6.1. Mean (\pm st. dev.) of CVL and LCV bandwidths, for each type of spatial point process we considered. The expected number of points for each listed process is 400.	166
Table 6.2 RIMSE $\times 10^2$ of local and global K -function estimators with CVL and LCV bandwidths.	168
Table 6.3 Mean (\pm st. dev.) of CVL and LCV selected bandwidths for the simulated two point process cases. Expected number of points is 400 for each listed process.....	171
Table 6.4 Root integrated mean squared errors $\times 10^2$ of local and global K_{12} -function estimators with CVL and LCV bandwidths.	172
Table 6.5 Root integrated mean squared errors $\times 10^2$ of local and global K -function estimators with parametric intensity estimator, applied to point processes with intensity function “waves” or “deep waves”.	174

List of Figures

Figure 2.1 Phase diagram of lipid mixtures	18
Figure 2.2 Phase diagrams of ternary lipid mixtures exhibit the same overall topology.....	20
Figure 2.3 Critical systems exhibit universal features that extend beyond the 2 phase region	25
Figure 2.4 Four Functional Mechanisms Primarily Driven by L_o/L_d Partitioning in a Near-Critical Membrane	30
Figure 3.1 Measurement of cholesterol activity in membranes referenced to cholesterol crystals.	47
Figure 3.2 Cholesterol activity (A), chemical potential (B), and excess Gibbs free energy per molecule (C) in binary PC/Chol membranes.....	48
Figure 3.3 μ_c in ternary DOPC/DPPC/cholesterol membranes.	51
Figure 3.4 Fits of μ_c to two models with condensed complexes.	53
Figure 3.5 Fits of μ_c to mean field (A) and lattice (B) models of pairwise interactions between components.	55
Supplemental Figure 3.6 Saturated solutions of cholesterol in 5 mg/ml M β CD equilibrate within 5 min and remain stable over extended times.....	61
Supplemental Figure 3.7 NBD-PE is not detected in the M β CD containing aqueous phase.	62
Supplemental Figure 3.8 Experimental determination cholesterol content of vesicles equilibrated with M β CD solutions.....	63
Supplemental Figure 3.9 Equilibration of systems of DPPC and cholesterol.	64
Supplemental Figure 3.10 Characterization of saturated cholesterol solutions in 5 mg/ml M β CD.	65
Supplemental Figure 3.11 Calibration of phospholipid concentration using NBD-PE.....	66
Supplemental Figure 3.12 Calibration relating % saturation of cholesterol in 5 mg/ml M β CD solutions to cholesterol activity.	67
Supplemental Figure 3.13 Comparison of μ_c measurements from the present manuscript to past work from Tsamaloukas et al (35).....	68

Supplemental Figure 3.14 Measurements of cholesterol chemical potential from Figure 3.3 replotted as chemical activity (A-C) and excess Gibbs free energy per molecule (D-F).	69
Supplemental Figure 3.15 Fraction of cholesterol that remains free.	70
Supplemental Figure 3.16 Fitting μ_C to two regular solutions models with complexes and two repulsive interactions.	71
Figure 4.1 Demonstration of the mean shift algorithm.	84
Figure 4.2 Evaluating the mean shift (MS) algorithm on simulated data translated in 2 dimensions, compared to nonlinear least squares (NLLS) fitting approach.	85
Figure 4.3 Demonstration of mean shift (MS) drift correction of a 2D SMLM dataset of antibody-labelled Nup210 in nuclear pore complexes, within the nuclear envelope of primary mouse neurons.	88
Figure 4.4 Demonstration of mean shift (MS) drift correction of a 3D SMLM dataset of B cell receptors at the ventral plasma membrane of CH27 B cells.	90
Figure 4.5 Evaluating the mean shift (MS) and NLLS algorithms with the start-point at the origin.	98
Figure 4.6 Drift estimates between overlapping temporal bins are prone to bias.	99
Figure 4.7 Drift correction diagnostics for the nuclear pore complex dataset of Figure 4.3.	100
Figure 4.8 Evaluating the mean shift (MS) algorithm on 3D simulated data, compared to the NLLS approach.	101
Figure 4.9 2D projections of 3D data degrade mean shift (MS) shift estimation performance. ..	101
Figure 4.10 Drift correction diagnostics for the 3D B cell dataset of Figure 4.4.	102
Figure 5.1 Validation of approach through simulation.	117
Figure 5.2 Validation of approach through simulation with drift and drift correction.	118
Figure 5.3 Experimental observations of DNA origami rulers labeled with AlexaFluor647.	120
Figure 5.4 Experimental observations of DNA origami rulers imaged with DNA PAINT, using an Atto655 imaging strand.	121
Figure 5.5 Experimental observations of Nuclear Pore Complexes (NPCs) within primary mouse neurons, antibody-labeled with AlexaFluor647.	123
Figure 5.6 Experimental observations of clathrin coated pits within CH27 B cells, imaged using a nanobody-coupled Atto655 DNA-PAINT scheme.	124

Figure 5.7 Experimental observations of F-actin on the ventral surface of a CH27 B cell using phalloidin-AlexaFluor647.....	125
Figure 5.8 Experimental observations of membrane anchor peptide Src15-mEos3.2 on the ventral surface of a CH27 B cell.....	126
Supplemental Figure 5.9 Plots of $g(r < 25 \text{ nm}, \tau)$ for the experimental samples shown the main text.....	131
Supplemental Figure 5.10 Subtracting Gaussian shapes with different width leads to distortion in $\Delta g(r) = g(r, \tau) - g(r, \tau_{max})$ when gr, τ and $g(r, \tau_{max})$ have similar amplitudes but different widths.....	131
Supplemental Figure 5.11 Simulation with drift, drift correction and incoherent single molecule motions.....	132
Supplemental Figure 5.12 10 μm by 10 μm region showing simulated localizations from Figs 1-2. The full simulated area was 40 μm by 40 μm . Scale bar is 1 μm	133
Supplemental Figure 5.13 Localization spread functions (LSFs) estimated as $\Delta g(r, \tau)$ (points) are well described by Gaussian functions (Eqn. 8 of main text; solid lines). Fits are used to extract $\sigma_{xy}(\tau)$ reported in the main text.	134
Supplemental Figure 5.14 10 μm by 10 μm region showing DNA origami rulers analyzed in Figure 5.3. The full imaged area was 40 μm by 40 μm . Scale bar is 1 μm	135
Supplemental Figure 5.15 10 μm by 10 μm region showing DNA origami rulers analyzed in Figure 5.4. The full imaged area was 40 μm by 40 μm . Scale bar is 1 μm	136
Supplemental Figure 5.16 Resolution calculation with Fourier Ring Correlation (FRC).	137
Supplemental Figure 5.17 Comparison of LSF widths (left) and distributions of CRLB localization errors returned from fits (right) for datasets fit without background subtraction.	138
Supplemental Figure 5.18 Localization precision determined by the Nearest Neighbor analysis (NeNA).	139
Supplemental Figure 5.19 Estimates of the autocorrelation of the 2-dimensional LSF, $g_{LSF}(r, \tau)$, for the Nuclear Pore Complex data of Figure 5.5.	140
Supplemental Figure 5.20 Estimates of the weighted average of the autocorrelation of the 2-dimensional LSF g_{LSF} , for the indicated datasets.....	140
Figure 6.1 Plots of the “hole”, “waves” and “LGF” thinning profiles.	164
Figure 6.2 Averages of estimates of $L(r) - r$ obtained from simulations in case of the “waves” intensity function with 400 simulated points on average.....	165

Figure 6.3 Averages of estimates of $L(r) - r$ obtained from simulations in case of the waves intensity function with 400 simulated points on average.....	167
Figure 6.4 Averages and 95 % pointwise probability intervals for estimates of $L(r) - r$ in case of the waves intensity function with 400 simulated points on average.....	168
Figure 6.5 Averages of estimates of $g_1(r)$ obtained from simulations in case of the “waves” intensity function with 400 simulated points on average.....	169
Figure 6.6 Averages of estimates of cross- $L(r) - r$ in case of the “waves” intensity function with 400 simulated points on average.....	171
Figure 6.7 Averages and 95 % pointwise probability intervals for estimates of $L_{12}(r) - r$ in case of the “waves” intensity function with 400 simulated points on average.....	172
Figure 6.8 Averages of estimates of $c(r)$ in case of the waves intensity function with 400 simulated points on average.	173
Figure 6.9 Averages and 95 % pointwise probability intervals for estimates of $L(r) - r$ in case of the “waves” (top row) or “deep waves” (bottom row) intensity function with 400 simulated points on average.	174
Figure 7.1 Schematic of a cross-correlation measurement without density correction, for a simulated SMLM dataset.	191
Figure 7.2 Two density correction approaches for the cross-correlation measurements of Figure 7.1.....	192
Figure 7.3 Global density correction achieves lower variance than local correction in nonuniform labeling conditions.	193
Figure 7.4 Density corrections applied to simulated SMLM dataset with uniform labeling but non-uniform detection of localizations.	194
Figure 7.5 Comparing pairwise and pointwise edge corrections for the kernel-based γ estimators.	195
Figure 7.6 High-amplitude, short-range density variations lead to bias in estimated normalization factor γ_σ	196
Figure 7.7 Strong correlations lead to bias in gamma estimates.....	197
Figure 7.7 Density corrected cross-correlations of phalloidin-AlexaFluor647 and cholera toxin B subunit-AlexaFluor532 labeling a CH27 B cell.....	198

Abstract

The lipid membranes of cells are complex structural and functional landscapes. Beyond being a selective barrier separating the cell from its surroundings, the membrane serves also as a two-dimensional solvent dictating the thermodynamic environment in which membrane protein biochemistry takes place, and as a platform that facilitates and responds to the organization of membrane proteins into functional domains. Membranes including vesicles derived from eukaryotic plasma membranes also exhibit liquid-liquid phase coexistence. This dissertation aims to link the biochemical and organizational properties of membranes to their phase behavior.

The membrane's role as a thermodynamic platform is addressed in a chapter on the availability of cholesterol, specifically its chemical potential (Chapter 3). This work consists of measurements of the chemical potential of cholesterol in a family of synthetic lipid membrane compositions. This chemical potential describes the availability of cholesterol, and is a primary determinant of the occupancy of protein binding sites for cholesterol. The synthetic membranes used in this study are similar to mammalian plasma membranes in phase behavior and cholesterol concentration. The measurements show a close connection between the role of cholesterol in phase separation of these membranes and its availability. This finding suggests that treatments that modify the phase behavior of the membrane, of which many are known, may act through their effect on the availability of cholesterol. In addition, this study provides a framework for how to approach other questions about the biochemistry of cholesterol.

The remaining chapters describe methods that will enable more precise and robust measurement and analysis of the organization of membrane proteins using single molecule localization microscopy (SMLM). SMLM techniques produce location information of target molecules with precisions on the order of 10 nm, and so have been invaluable for characterizing protein organization in membranes. The methods contributions include direct improvements to the precision of these datasets through improved sample drift correction (Chapter 4), a novel method for characterizing SMLM measurement precision (Chapter 5), and a method for correcting spatially non-uniform labeling or detection artifacts in measurements of colocalization (Chapters 6 and 7). These methods extend the usefulness of SMLM so that it can detect more detailed and

subtler structure in the organization of proteins on membranes. In particular, they will enable future experiments to measure the role of membrane phase behavior in biological systems where it has been too subtle to detect using past methods.

Overall, the developments described in this dissertation strengthen the connections between membrane phase behavior and biological function, by linking phase behavior to a new biochemical property of the membrane, and by enabling future investigations into how it organizes membrane proteins.

Chapter 1 Introduction

This dissertation consists of several developments that relate to the role of biomembrane phase behavior in biological function. This chapter gives background and context for the mechanisms and experimental methods that are discussed in later chapters. Section 1.1 describes membranes and their phase transitions, as well as some of the ways that membrane physical and chemical properties can serve biological function. Section 1.2 introduces single molecule localization microscopy, a high resolution microscopy technique that is useful for measuring nanometer scale organization of membrane proteins. Section 1.3 defines pair correlation functions, which give a rigorous way to statistically quantify colocalization, including in circumstances where organization is subtle. Finally, section 1.4 gives an overview of the work that is presented in the rest of the dissertation and its significance.

1.1 Biomembrane structure, physical chemistry, and spatial organization

The core structure of all biological membranes is the lipid bilayer (1). Lipids are amphiphilic molecules, meaning that they have one hydrophilic end, known as the “head group”, that interacts favorably with water, and a long hydrophobic region that has unfavorable interactions with water. This hydrophobic group is usually made up of a pair of acyl chain “tails”. These molecules form stable bilayer membranes: the hydrophobic tails of the lipids align next to each other, preventing unfavorable contacts with the aqueous phase, and two such leaflets assemble with their hydrophobic faces apposed to each other. The bilayer is thus composed of a hydrophobic core sandwiched between two hydrophilic head group regions that each face the aqueous phase.

Lipid bilayers exhibit several kinds of phase transitions. Chapter 2 contains a more thorough review of what is known about the phase behavior of lipid membranes, with special emphasis on critical phenomena, their observation in cell-derived vesicles, and potential mechanisms by which the physical properties of the critical point could affect biological function. By way of introduction, a few details are highlighted below.

First, at sufficiently high temperatures most bilayers form liquid phases, meaning that lipids can readily diffuse in the plane of the membrane. This is the case for most biological

membranes, and many biological processes depend on the fact that their membranes are fluid so that membrane proteins can move in the membrane to encounter signaling partners or substrates. Decreasing temperature induces most fluid membranes to transition to a solid phase, with lipids essentially fixed in place next to each other (2). The temperature at which this phase transition occurs is referred to as the membrane's melting temperature T_m .

Some mixtures of lipids can support two liquid phases that coexist (3–11). In model membrane systems this is often observed by mixing three lipids: a high- T_m lipid, a low- T_m lipid, and cholesterol. When these are mixed within a certain range of ratios, the membrane phase-separates into a so-called liquid-ordered (l_o) phase rich in high- T_m lipid and slightly enriched in cholesterol, and a liquid-disordered (l_d) phase rich in low- T_m lipid and slightly depleted of cholesterol (12–14). The presence of coexisting phases can be directly visualized using fluorescent lipophilic dyes, many of which partition strongly into the l_d phase (6). Coexisting phases have also been observed by atomic force microscopy on supported lipid bilayers (15), where the contrast between the phases is due to the different thicknesses and elastic properties of the two phases – in most cases the l_o phase is thicker and more rigid. Both of these are liquid phases, although the l_o phase is more viscous than the l_d phase (16).

A common theme has been the importance of cholesterol, whose rigid ring chemical structure constrains the conformational flexibility of nearby acyl chains, promoting extended conformations in saturated, high- T_m lipids (17, 18). The resulting l_o phase is more tightly packed, allowing for more favorable van der Waals interactions between the acyl chains. Low- T_m lipids containing double bonds in their acyl chains are excluded from this more packed environment because the kink of the double bonds makes extended conformations energetically unfavorable.

Liquid-liquid phase separation has also been observed directly in vesicles derived from the plasma membranes of mammalian cells (16, 19, 20). These vesicles phase separate at temperatures somewhat below the growth temperature of cells (21), suggesting that plasma membranes are unlikely to ever experience phase separation per se. However, the phase behavior can still play a role at growth temperature due to the existence of a critical phase transition in this system (20). Near a critical point, many thermodynamic quantities diverge as power laws as a function of the distance from the critical point (22). As a result, these quantities can be anomalously large in the vicinity of the critical point, even if the system never crosses the phase transition. One such property is the susceptibility, meaning the magnitude of response that is induced by a given

external applied force. In the context of the membrane, this external force could be the clustering of proteins that prefer the l_o (or l_d) phase, and the response would be a concomitant local change of membrane composition to more closely resemble the composition of the l_o (or l_d) phase. This local compositional heterogeneity, in turn, can recruit other proteins that favor the same phase. Recent measurements have validated this concept by directly comparing the partitioning of membrane anchored peptides in phase-separated plasma membrane vesicles, to their sorting around clustered B-cell receptors, which prefer the l_o phase (23).

The chemical potentials of lipids in the membrane can also directly inform biological function. In many membrane proteins including diverse ion channels, lipids such as cholesterol and phosphatidylinositol phosphates are thought to bind to specific binding sites (24–32). The chemical potential μ_L of each lipid L describes its availability to participate in such interactions. Formally, μ_L is the Gibbs free energy increment associated with adding one molecule of species L to the system. The chemical potential must be equal in all phases of an equilibrated system, since unequal μ_L would imply that the overall Gibbs free energy could be decreased by moving a molecule of L from a phase with high μ_L to a phase with low μ_L . A similar argument implies that a binding site for L is occupied with increasing probability as μ_L increases.

Because μ_L is a derivative of the Gibbs free energy, it is determined by both enthalpic and entropic contributions to the free energy. In the ideal case of a dilute mixture (low concentration of L), the enthalpic contribution of each molecule is identical – each molecule interacts exclusively with the other components of the mixture, to a good approximation. Therefore the concentration dependence of μ_L is entirely determined by the mixing entropy, giving simply $\mu_L = \mu_0 + k_B T \log x_L$, where x_L is the mole fraction of the lipid L . Several groups have published measurements of cholesterol partitioning between vesicles of different phospholipid compositions in this dilute regime (33–35).

However, when L is not dilute, the enthalpic contributions can vary substantially and lead to non-idealities of the chemical potential. Cholesterol, for example, is present in cell plasma membranes at concentrations of 30 – 40 mole%. At such high mole fractions, the dilute solution model is no longer a good approximation: contacts between pairs of cholesterol molecules become important. Furthermore, cholesterol is known to induce structural changes in the membrane, for example making it thicker (15). These concentration dependent effects are expected to lead to non-

idealities in the concentration-dependence of the chemical potential, but past work has not explored this regime experimentally.

1.2 Single molecule localization microscopy

As discussed above, the organization of biological membranes by phase-like lipid interactions does not lead to macroscopic phase separation, and indeed the relevant length-scales are frequently less than 100 nm (36). Therefore, detecting and characterizing this organization requires measurements that can achieve precisions much better than 100 nm. Single molecule localization microscopy (SMLM) is well suited to this case.

SMLM refers to a family of microscopy techniques that determine positional information about a molecule of interest with precision on the order of 10 nm using a standard far-field fluorescence microscope. Far-field microscopes are diffraction-limited, so that a point source in the sample gives rise to an extended spot of finite width in the image produced by such a microscope. The intensity profile of this spot is known as the point-spread function (PSF) of the microscope. When the microscope is properly focused and in the absence of aberrations, the width of the PSF is given by the Abbé diffraction limit: $\lambda / 2NA$, where λ is the wavelength of light that is being detected and NA is the numerical aperture of the microscope, a dimensionless quantity that describes the angular extent of the cone of light that can be collected by the microscope. In state of the art microscopes NA can be as high as 1.5, so that the diffraction limit is about one third of the wavelength of the light that is being collected. Thus for a red-emitting fluorophore ($\lambda \approx 630$ nm), the diffraction limit is still more than 200 nm. As a result, a sample with many point sources each spaced by much less than 200 nm results in an image where the PSFs all overlap, so that features smaller than about 200 nm are not readily distinguishable from a single point source – information about the nanometer-scale arrangement of the molecules is lost.

The fundamental insight of SMLM techniques is that sparsely distributed point sources do not have this problem: the location of a single fluorescent molecule can be determined very precisely from its image by finding the center of the PSF. A useful simplified picture is that the image of a single fluorophore is a collection of the locations where N photons originating from that fluorophore were detected. Each detection is distributed according to a probability density function with some standard deviation σ on the order of the diffraction limit. By collecting N photons we can improve the precision of the estimate of the center of the distribution to σ / \sqrt{N} .

So for $\sigma \approx 200$ nm and $N \approx 400$ we can obtain a localization precision of approximately 10 nm. To apply this insight to samples that are dense, we must apply some trick so that we only see fluorescence from a sparse subset of the molecules at any given time. In addition, that subset must change over time, so that any given molecule is visible at some point during the observation time with high probability.

Since the inception of SMLM nearly 20 years ago (37–39), several such tricks have been devised and implemented. Many of these involve choosing fluorophores that can be manipulated using buffer chemistry and/or illumination conditions so that only a small fraction of the molecules in the sample are fluorescent at any given time, while most molecules occupy a non-fluorescent “dark” state. These include organic dyes with long-lived triplet excited states (40, 41) or that reversibly isomerize between a fluorescent and non-fluorescent form (42). Other dyes and fluorescent proteins rely on light-induced irreversible transitions from a non-fluorescent to a fluorescent form, followed by photobleaching (43–46).

Another class of SMLM strategies known as PAINT (Points Accumulation In Nanoscale Topography) depends on weak transient binding of a bright soluble fluorophore to the (non-fluorescent) target molecule, so that only a few of the molecules of interest have a bound fluorophore at any given time (47). The off-rate of binding is engineered to be somewhat longer than the integration time of the camera, so that bound fluorophores can be readily distinguished from the diffuse background fluorescence produced by the fast-diffusing soluble population. In particular, DNA-PAINT (48) takes advantage of the well-understood energetics and cheap availability of DNA oligomers by conjugating complementary single-stranded DNA oligomers to the target molecule and the fluorophore (known as the docking strand and the imager strand, respectively). The off-rate of this DNA-DNA binding interaction can be straightforwardly tuned by varying the length of the base-pairing sequence in the DNA oligomers.

Because SMLM observes a small fraction of the target molecules in each image, many images must be collected to fully sample the locations of all of the target molecules, over an observation time measured in minutes (49). In living cells, most structures are dynamic on these timescales, so molecules that are imaged at the end of the observation time represent different structures than existed at the beginning of the observation time. It is therefore common to use cells that are chemically fixed e.g. by paraformaldehyde to crosslink the cell’s structures and prevent them from moving during imaging. Then the SMLM dataset is effectively a snapshot of the cell’s

organization at the time of fixation. Most of the work in this thesis is primarily applicable to fixed-cell SMLM data. Note that live cell SMLM data can still be informative by applying different analytical techniques than would be useful in fixed-cell data. These techniques include single-particle tracking for diffusion analysis (50–52), as well as spatiotemporal correlation analysis, both of which take into account the time delays between observations in addition to their spatial properties (53).

The various strategies discussed above for achieving sparse stochastic fluorescence from a target molecule of interest all lead to relatively similar datasets coming off the microscope: a large number of images (typically several thousands) that each consist of some bright PSFs in an otherwise dark image. As a result, the process of analyzing raw SMLM image data can typically be shared among the modalities. The standard analytical “pipeline” can be broken into several independent steps, which we refer to as segmentation, fitting, culling, and drift correction, described in the next few paragraphs. When two colors are used, there is an additional step after fitting to align the two color channels with subpixel accuracy.

Segmentation refers to an initial coarse peak-finding step to determine which regions of the image contain candidate fluorophores (54). The raw image is convolved with a filter that is most sensitive to objects that are the size of the PSF. A noise floor is determined from the variance of the same image filtered at a higher frequency. The filtered image is then compared to a user-defined threshold signal to noise ratio, and regions that exceed the threshold are considered to be candidate fluorophore PSFs. For the datasets analyzed for this thesis, a wavelet filter was applied following the method of (55), and a threshold signal to noise ratio of 2 was standard.

The fitting step considers each candidate PSF separately, fitting it to a statistical model to obtain an estimate of the center location of the PSF, which is therefore also an estimate of the true location of the fluorophore in the sample (56). The fitting model can be thought of as specifying a PSF model $I_{PSF}(i, j | x_0, y_0)$ giving the expected number of photons absorbed by pixel (i, j) given a fluorophore at location (x_0, y_0) , and additionally a noise model $P_{noise}(n | m)$, the probability that the camera will record n digital counts for a pixel that absorbs m photons. I_{PSF} essentially models the shape of the PSF, and is often taken to be a 2-dimensional Gaussian function, integrated over the area of each pixel. While this model is approximate, it has been found to be adequate in practice, with negligible loss of precision compared to more accurate models of the PSF, in realistic conditions (57). The noise model frequently considers the Poisson-distributed photon counting

noise, and often treats other sources of camera noise as negligible. Given this model of the single molecule fluorescence data, the best fit location (x_0, y_0) can be determined by maximizing the likelihood function implied by the model, with respect to the data given by the camera’s digital output for each pixel. In the present work fitting is typically performed on a square subregion 7 px in each dimension using GPU-based fitting code adapted from (58).

The culling step considers the parameter estimates from the fit as well as the goodness-of-fit criteria that can be determined from the fitting step, and rejects fits that are deemed unlikely to be a true single-molecule PSF, either because two molecules were present or some camera noise was mistaken for a fluorophore (54).

Finally, although microscope stages are engineered to be quite stable, many of them still exhibit some stage drift that can amount to 10 – 100 nm over the course of a 10 – 30 minute SMLM image acquisition time (59, 60). In a single frame, localization precisions can easily be smaller than 10 nm, so uncorrected drift represents a substantial loss of precision when comparing localizations that are distant in time. Drift is modeled as a time-dependent rigid translation $\mathbf{u}_d(t)$ of the sample with respect to the microscope reference frame, so that a fluorophore at location \mathbf{x} in the sample reference frame is at location $\mathbf{x} + \mathbf{u}_d(t)$ in the microscope reference frame at time t (61). Thus, given an estimate of the drift trajectory $\mathbf{u}_d(t)$, the drift correction is as simple as subtracting $\mathbf{u}_d(t)$ from the observed location of each localization from time t . Estimating the drift correction is less straightforward. Various successful methods have been proposed (61–65). A common scheme is to divide the localizations into disjoint temporal windows centered on t_1 and t_2 , find the spatial displacement that maximizes the similarity of the two subsets, and take that displacement as an estimate of $\mathbf{u}_d(t_2) - \mathbf{u}_d(t_1)$.

It is possible to apply SMLM to two molecules of interest at once by labeling them with different fluorophores whose emission spectra can be separated by a dichroic beam splitter in the emission path of the microscope, so that emission from the first fluorophore – the first “color channel” – is imaged onto half of the camera’s field of view, and the other color channel is imaged onto the other half of the camera’s field of view. The segmentation and fitting process can then be carried out independently for each color channel, with resulting localizations assigned to the type of molecule being imaged in that color channel. To obtain precise information about the arrangement of one type of molecule with respect to the other, a precise measurement of the spatial transformation from one color channel to the other is required. A sparse sample of nano-scale

beads that are decorated with multiple fluorophores whose emission spectra fall in both color channels can be used to do so. Each such bead produces a PSF-like image in both color channels, which can be precisely localized by the same methods as are applied to single molecules. The beads are distributed sparsely enough that it is easy to computationally identify which localizations in the two channels correspond to the same bead. By repeating this process over many images (hundreds of images, thousands of individual bead localizations), a map from one channel to the other is established. For the purposes of this thesis, a “local weighted mean” transform is used in a manner similar to that described in (66): for a localization in channel 2, the 12 nearest localizations from the bead sample are found. The corresponding point in channel 1 is determined by a weighted mean of the 12 corresponding bead localizations in channel 1. The weights for the weighted mean are a decreasing function of the distances in channel 2 from the localization of interest. In the context of this thesis, this step is applied between fitting and culling, when applicable.

Despite nearly 20 years of history, SMLM methods continue to be an active research area, and recent developments have been reported throughout the process. These include contributions that improve the image data that the analysis is applied to, for example: labeling schemes that can achieve higher efficiencies and specificities of labeling (67–70); fluorescent probes with superior blinking characteristics (42, 71–73); and microscope design improvements that allow the user to obtain complementary information about the fluorophores such as their axial position (74, 75) or their spectra (76). The analysis algorithms themselves are also an active research area: new analysis algorithms produce more precise localizations by optimizing the statistical and/or computational performance of each analysis step (77–79). The statistical developments allow for improved image quality, so that the experimentalist can more confidently resolve the structure of interest at smaller length scales, and computational developments reduce the barriers to applying sophisticated methods routinely to more and larger datasets. Finally, determination of the quality of an SMLM dataset is an important research area, because it enables unbiased evaluation of the contributions of various new methods to dataset quality (80–82).

1.3 Detecting spatial organization in SMLM datasets

The spatial precision of SMLM enables the experimentalist to study spatial organization of a variety of nanoscale structures. The diversity of structures of interest, and the research questions

one might ask about them, calls for a diversity of approaches as well. In some cases, the labeled molecules are arranged into extended structures with clear geometric features, such as regular spacing of molecular subunits in a lattice-like structure, linear features that branch at specific angles, or tubes with a characteristic diameter. In such cases it may be appropriate to apply dedicated tools to obtain best estimates of the geometric characteristics of individual features. At the opposite extreme, molecules with weak interactions including those mediated by interactions with membrane phase-like domains (23, 83, 84) can give rise to minor but quantifiable enrichment of one type of molecule around another. In this case statistical tools are necessary to distinguish between chance encounters and meaningful colocalization.

Pair correlation functions have proven to be a useful tool for this purpose. Pair correlation functions $g(r)$ can be thought of as a ratio of two densities: the average density of points at a distance r from another point, and the average density of points overall (85–87). Thus for example if proteins are distributed uniformly and independently over a cell's plasma membrane, their correlation function will be $g(r) = 1$, meaning that the density of protein around a given location \mathbf{x} does not depend on whether or not there is a protein at \mathbf{x} . In contrast, a protein that is clustered is denser near other copies of the protein, so that $g(r) > 1$ over a range of r up to about the characteristic size of the clusters. Repulsive interactions would instead imply $g(r) < 1$.

Note that the pair correlation function is non-parametric: computing it does not rely on any assumptions about the mechanism or mechanisms mediating putative clustering or repulsion. Conversely, the presence of non-trivial correlations at a given distance r is not necessarily an indication that the labeled objects interact directly with each other. For example, it could instead indicate that the molecules all interact with some other structure.

When more than one type of point is present in a dataset, for example in a two-color SMLM dataset where the two colors label distinct types of protein, one can compute pair correlation functions as above for each type, often called pair auto-correlation functions for clarity. In addition, one can compute an analogous pair cross-correlation function $c(r)$ between type A and type B, so that $c(r)$ is the ratio of the density of points of type A around points of type B to the overall density of type A. The cross-correlation functions have similar interpretations to the auto-correlation functions: $c(r) = 1$ is evidence that the two types of points are distributed independently of each other, $c(r) > 1$ is evidence that they are co-clustered, and $c(r) < 1$ is evidence that they are segregated from each other.

In the context of SMLM, pair correlation functions are computed from localizations, not protein locations themselves. This fact has three important consequences for correlation analysis. First, localizations are subject to finite spatial errors with respect to the fluorophores that produce them (87). As a result, the pair correlation functions of localizations are effectively blurred relative to the pair correlation functions of the fluorophores, and sharp features in the correlation functions of the proteins give rise to broader and lower-amplitude features in the correlation functions of the localizations.

Second, SMLM datasets typically contain several or many repeat localizations from each detected protein, and it is difficult to attribute particular localizations to particular proteins. This phenomenon is known as overcounting, since counting the localizations overestimates the number of proteins. Therefore, the auto-correlation function is not simply a measurement of the organization of the labeled proteins, but also shows the clustering of the repeat localizations around true fluorophore locations (87). This overcounting feature takes the form of a peak in $g(r)$ as $r \rightarrow 0$, whose shape is related to the distribution of localization errors, and whose amplitude is related to the density of fluorophores and the distribution of numbers of localizations per fluorophore. In a typical SMLM experiment, the overcounting peak is the dominant feature of the autocorrelation function, and can make it difficult to detect clustering of the underlying molecules. However, the cross-correlation function of a two-color SMLM dataset does not have an overcounting peak, because the cross-correlation is derived from pairs of one localization from each color channel, whereas repeat localizations from the same fluorophore can only be found in the same color channel. As a result, a pair cross-correlation function can be more useful for detecting subtle clustering.

Finally, experimental challenges can limit both the efficiency of labeling a target molecule with a fluorophore, and the probability that a given fluorophore gets localized during the imaging process. While high-affinity, high-specificity engineered labeling schemes such as nanobodies, fluorescent fusion proteins and genetically encoded self-labeling tags have done a lot to mitigate labeling efficiency issues in the best case, they tend to require genetic control that is not always feasible. In addition, labeling of plasma membrane proteins in a way that precludes labeling of intracellular populations can present challenges, where cells must be labeled before fixation to prevent labeling of the intracellular pool, live-labeling for a long incubation time leads to excessive internalization of the plasma membrane population, but short incubation times lead instead to

inadequate access of the label to the bottom surface of the cell, so that only a subset of the target molecules are labelled. On the imaging side, an experiment can be more or less effective in detecting the labels that are present. Fluorophores that depend on photoconversion between fluorescent and non-fluorescent states for blinking may produce more or fewer localizations depending sensitively on the illumination intensities, as well as buffer chemical conditions. Separately, poor signal to noise ratios due to background fluorescence or low power excitation will cause some emitting fluorophores to be missed by the segmentation step.

Missed localizations for either kind of reason generally make the correlation functions more noisy, so it is always beneficial to improve the labeling and detection efficiencies when possible (86). If the missed localizations are equally likely to occur at any location, the expected values of the correlation functions are unaffected, so still reflect the correlation functions of the labeled molecules (apart from the contribution of localization precision described above). However, if labeling or detection are spatially non-uniform, the standard correlation functions are modified by a factor that depends on the details of where detection is more and less likely (88).

The purely spatial pair correlation functions described so far are mainly applicable to SMLM datasets from chemically fixed cells, where the dynamics of when localizations are observed are not relevant to the static structures that are being imaged. In cases where dynamics are of interest, such as in live-cell SMLM, the spatiotemporal pair auto- and cross-correlation functions $g(r, \tau)$ and $c(r, \tau)$ can be calculated instead (53, 89, 90). These can also be interpreted as ratios of densities, with the added criteria that the density around a point is considered as a function of the time-delay τ since that point was observed, as well as distance from the point. In a live cell spatio-temporal pair correlation function, the limit $\tau \rightarrow 0$ describes the instantaneous organization of molecules, and so can be directly compared to fixed cell spatial pair correlation functions (23, 83, 91). When the labeled structure is transient, or when it diffuses through the sample, spatiotemporal pair correlation functions generally decay towards 1 with increasing τ : in other words the presence of a molecule at given location and time is less and less informative about the density of other molecules at that location at a later time when the original molecule has diffused away. Live cells also have a kind of overcounting in their autocorrelation $g(r, \tau)$, due to repeat localizations from the same molecule. In this case the overcounting peak of $g(r, \tau)$ characterizes the diffusion of the labeled molecules: the width of the peak is determined by the distribution of displacements of single molecules over a time delay τ . This fact can be used to

estimate diffusion coefficients even when localizations are too dense to precisely track single molecules (23, 53, 91).

1.4 Overview of the dissertation

This dissertation presents several projects that address the subjects discussed above in various aspects. The chapters can be roughly grouped into two parts. The first two chapters relate most directly to membrane phase behavior and physical properties, including new measurements of cholesterol chemical potential in membrane systems near liquid-liquid coexistence. While the availability of cholesterol and the phase behavior of the membrane are both considered to be important for biological function in different contexts, these measurements suggest that they are closely linked, so that treatments that perturb phase behavior could act on biological function through cholesterol availability, or vice versa. The remaining chapters describe methods for improving the processing and interpretation of single molecule localization microscopy data. Taken together, these improvements allow SMLM to produce trustworthy measurements at shorter length-scales, so that smaller structures with less contrast from their surroundings will be distinguishable. In particular, membrane organization by phase-like interactions between membrane components will be observable in systems with smaller clusters and weaker interactions than was previously possible.

Chapter 2 contains a review of what is known about critical phenomena in membranes, including evidence for functional roles of critical phenomena in mammalian plasma membranes. Among other ideas, this review discusses features of membrane phase behavior that are relevant to the rest of the chapters. This includes the essential role of cholesterol in membrane phase behavior, with the suggestion that proximity to the l_o/l_d miscibility critical point could sensitively control the availability of cholesterol as a ligand in biological membranes, which is also the topic of chapter 3. In addition, this review discusses how near-critical membrane phase behavior can act as an organizing force, so that clustering membrane components that prefer one phase (l_o or l_d) will also induce recruitment of other components that prefer that phase, an effect that is amplified near the critical point due to the associated high susceptibility of the membrane. This picture has been tested in the Veatch lab by SMLM experiments that detect co-clustering of membrane components through correlation analysis (83, 84, 91). Chapters 4-7 of the thesis present methods that are motivated by improving these kinds of measurements in various ways.

Chapter 3 describes systematic measurements of the chemical potential of cholesterol in defined mixtures with several phosphatidylcholine lipids at room temperature. The mixtures cover a wide range of cholesterol concentrations, and therefore reveal clear non-idealities in the concentration dependence of the chemical potential. In addition, mixtures in and near the known l_o/l_d liquid-liquid coexistence region were tested. The measurements within the two-phase region display tie line behavior consistent with known phase diagrams. Interestingly, measurements near the two-phase region have similar near-linear contours that fall parallel to tie lines. Finally, the measurements were compared to several regular-solution type thermodynamic models to gain further insight into the major interactions that give rise to the observed chemical potentials. The significance of this work

Chapter 4 presents a new method for estimating microscope stage drift from SMLM data based on a mean shift based algorithm. The method is quite simple to describe and amenable to a computationally efficient implementation. It also proves to be more robust than previously reported methods to the kinds of noise that are present in SMLM data, so that it is capable of estimating drift from fewer localizations. Together, these features make it feasible to estimate drift with high time resolution, so that nonlinear drift can be accurately corrected for. As an added benefit, the low memory requirements of the mean shift algorithm extend to the case of 3D SMLM datasets with drift in all three dimensions, in contrast to previous methods which required much more memory to simultaneously estimate drift in a third dimension.

Chapter 5 introduces the concept of the localization spread function (LSF) of an SMLM experiment as a way of characterizing the quality of an SMLM dataset. By analogy with the point spread function of a microscope, the LSF describes the distribution of localization errors around the true position of a fluorophore in the sample. To estimate the LSF from a static (e.g. fixed cell) SMLM dataset, we take advantage of the time correlated switching of SMLM probes to statistically separate pairs of localizations that come from the same probe, from pairs that come from different probes. This method notably is sensitive to all of the accumulated errors of the SMLM analysis pipeline, for example including imprecise drift correction. Therefore it can be used to compare the effects of different analysis methods, or of different choices of parameters in those methods.

Chapters 6 and 7 both concern correlation analysis of point datasets that are subject to spatially non-uniform labeling or detection probabilities. They describe a novel approach to correcting for these non-uniform sampling effects when estimating pair auto- or cross-correlation

functions, so that the corrected correlation functions on average recover the correlations of the underlying molecules. In particular, chapter 6 gives a rigorous presentation of this correction method in the language of the statistics of point processes. This includes proofs that the method is unbiased in ideal circumstances and a convincing heuristic argument that it typically gives a smaller variance than a previously presented method (88). Furthermore, it gives detailed analysis of bias that can arise in both methods when the spatial profile of sampling probabilities is not known but instead estimated from the data. Finally, a thorough simulation study using standard point process models demonstrates that the bias and variance do hold for real data.

In chapter 7, the same density correction method is applied more specifically to SMLM data. This includes demonstrations of the method on simulated data that is constructed to more closely resemble an SMLM dataset, as well as an application to a real SMLM dataset with noticeable non-uniform sampling. In addition, this chapter presents an extension of the method that improves on the bias characteristics of chapter 6 in a way that gives more accurate results when the density varies on shorter length scales. Finally, the determinants of that bias are described and demonstrated through simulation, with discussion of how best to choose parameters of the method to obtain the most accurate results.

Chapter 2 Critical Phenomena in Plasma Membrane Organization and Function

This review was written with Subhadip Ghosh and Sarah L. Veatch. It presents a historical perspective on our current understanding of the phase behavior of mammalian plasma membranes, with an emphasis on the idea that the plasma membrane resides near the critical point of a liquid-liquid coexistence region. In addition, several mechanisms are presented that could implicate membrane phase behavior in various biological processes, including some that are supported by recent evidence, and others that are under investigation. The writing of this review was collaborative, and my contributions are spread throughout the text and figures. A version of the review is now published as an Annual Review of Physical Chemistry (92).

2.1 Overview

Spatial organization of the plasma membrane on 10-100 nm length scales has been a topic of interest to biology for decades. Heterogeneity has been hypothesized to play important roles in many membrane-associated biological processes, from coordination of signal transduction machinery to endo- and exocytosis to polarization, as iconically described in the lipid raft hypothesis (93). Various biophysical notions have been marshaled to provide explanations of membrane heterogeneity, especially from equilibrium thermodynamics and phase transitions, as coexisting liquid phases are readily observed in both purified membranes and membranes isolated from plasma membranes (94, 95). One important thread of this ongoing conversation explains nanoscale membrane structure as critical phenomena. In contrast to a classical phase separation picture of plasma membrane domains, which implies stable and well-defined discrete regions of defined composition, critical phenomena are subtle, dynamic and malleable, and inhabit the relevant nanoscopic length-scales.

In this review, we conduct a brief historical survey of membrane domains in both model and biological membranes and describe the consensus that has been reached regarding the macroscopic miscibility phase behavior of model membranes as well as remaining controversies regarding the microscopic heterogeneity reported in other regions of phase space. We introduce membrane criticality and the accumulated evidence that eukaryotic plasma membranes are near-

critical. We discuss where the concept of criticality fits into conventional descriptions of raft phenomenology, and describe some unique areas where criticality could play roles in biological function that go beyond simply organizing components.

2.2 An early history of domains in model and cell membranes.

There is a long history of detecting heterogeneity in bilayer membranes containing cholesterol. Some of the earliest studies that interrogated purified membranes in the late 1960s and early 1970s found evidence that liquid membranes containing a single phospholipid species and cholesterol contained structure on 1-100nm length-scales accessible to the spectroscopic methods available at the time (96–101). Over the decades, evidence for these microscopic domains in binary mixtures of phospholipids and cholesterol continued to accumulate, through additional spectroscopic studies, calorimetry, the application of fluorescence techniques such as fluorescence quenching, anisotropy, and Förster resonance energy transfer (FRET), and imaging methodologies such as freeze fracture electron microscopy (12–14, 102–110).

Largely in parallel to this physical chemistry characterization of membranes, cell biologists also postulated that membranes within cells could contain lipid-mediated sub-structures that could be important for cell functions. Early evidence came from the observation that cells could polarize their membrane lipid composition via differential trafficking of lipid species (111). It has also long been appreciated that cholesterol or similar sterols are vital to the proper functioning of most eukaryotic plasma membranes, as manipulation of membrane cholesterol content was shown to interfere with functions including endocytosis, several receptor mediated signaling cascades, and cell cycle control (112–115). In the early 1990s it was discovered that some detergents only partially solubilize cellular membranes, especially the plasma membrane (116, 117), and it was postulated that the insoluble fractions represented distinct membrane domains (93, 109). Researchers found that artificially clustering some membrane components individually could drive macroscopic colocalization, termed co-patching, detectable by conventional fluorescence microscopy (118). Subsequent spectroscopic and fluorescence measurements supported the hypothesis that cell membranes were heterogeneous in their lipid and protein content (119–121), and many studies correlated biological function with presumed membrane heterogeneity probed and perturbed via these and related assays (122–125). Each of these methods had well documented flaws (126), but the remarkable consistency of the conclusions drawn from different

methodologies provided convincing evidence that lipid-driven heterogeneity is relevant to the plasma membrane.

2.3 Liquid-liquid phase separation in model membranes.

2.3.1 Purified membranes:

In 2001, the first observations of macroscopically phase separated fluid domains were reported in bilayer membranes reconstituted from cellular extracts (3, 4). These initial observations emphasized similarities between the composition of the coexisting phases in model membranes and heterogeneity measured in cells by detergent solubilization methods. Soon after, it was appreciated that at least three lipid components were required to form macroscopically phase separated domains in bilayers: a high melting temperature lipid, a low melting temperature lipid, and a sterol such as cholesterol (5). Phase diagrams describing this macroscopic phase transition have now been mapped by numerous groups using a range of methods and are in good qualitative agreement (6–9, 15, 127–130). Since three components are required to observe liquid-immiscibility, phase diagrams have been mapped by varying temperature and the molar fraction of all three components. An introduction to reading and interpreting these phase diagrams is given in Figure 2.1, and several experimental phase diagrams are presented in Figure 2.2.

A range of experimental (10, 131–136) and simulation (17, 18) approaches using different lipid combinations have produced results consistent with phase diagrams topologically similar to those shown in Figure 2.1b, and their detailed characteristics have been described in several comprehensive review articles (95, 137, 138). Liquid immiscibility is most often observed at temperature below the melting temperature (T_m) of the high T_m lipid (6, 9, 127, 128), although there are exceptions (7). In the typical case, reported phase diagrams have a region of liquid-liquid coexistence, a region of three-phase coexistence (2 liquids and a solid), and two regions of liquid-solid coexistence. One of the liquid phases is called the liquid-disordered (Ld) phase, and it resembles the liquid crystalline phase of pure phospholipids (139). The second liquid phase is called liquid-ordered (Lo) (12), which was first characterized in mixtures of saturated phospholipids and cholesterol (13, 100, 107). The solid phase is often called gel, and is most likely $L\beta'$ (2).

The high cholesterol edge of the three-phase triangle is sloped such that the Lo phase contains a higher cholesterol mole fraction than the Ld phase. This edge of this triangle is also the

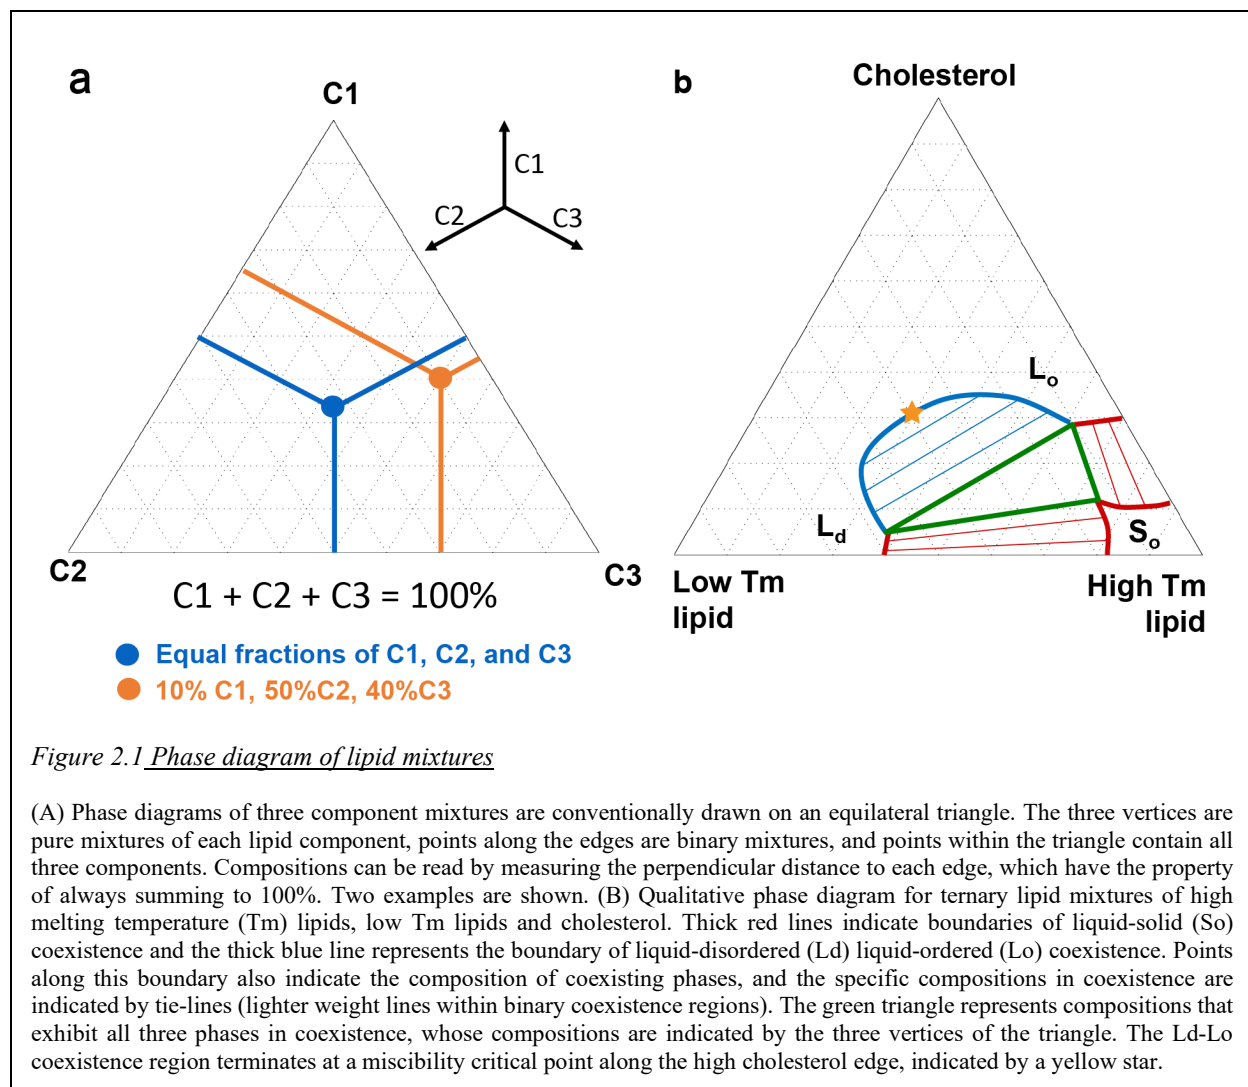


Figure 2.1 *Phase diagram of lipid mixtures*

(A) Phase diagrams of three component mixtures are conventionally drawn on an equilateral triangle. The three vertices are pure mixtures of each lipid component, points along the edges are binary mixtures, and points within the triangle contain all three components. Compositions can be read by measuring the perpendicular distance to each edge, which have the property of always summing to 100%. Two examples are shown. (B) Qualitative phase diagram for ternary lipid mixtures of high melting temperature (T_m) lipids, low T_m lipids and cholesterol. Thick red lines indicate boundaries of liquid-solid (S_o) coexistence and the thick blue line represents the boundary of liquid-disordered (L_d) liquid-ordered (L_o) coexistence. Points along this boundary also indicate the composition of coexisting phases, and the specific compositions in coexistence are indicated by tie-lines (lighter weight lines within binary coexistence regions). The green triangle represents compositions that exhibit all three phases in coexistence, whose compositions are indicated by the three vertices of the triangle. The L_d - L_o coexistence region terminates at a miscibility critical point along the high cholesterol edge, indicated by a yellow star.

first tie-line in the liquid-liquid-coexistence region. As cholesterol is increased further, tie-lines run roughly parallel to one another, meaning that cholesterol concentration increases roughly linearly in both phases. The L_o - L_d coexistence region terminates in a miscibility critical point, where in principle tie-lines merge into a single point. In practice this region of the phase diagram is surprisingly flat, meaning that tie-lines remain long and shorten over a very small range of compositions. As temperature is lowered, the L_o - L_d immiscibility gap extends to higher concentrations of cholesterol and low T_m lipid, as does the concentration of components at the critical point (7, 127, 130). At constant temperature, the miscibility gap expands when the T_m of the high T_m component is increased or when the T_m of the low T_m lipid is decreased (6, 7, 9). No macroscopic miscibility gap is observed for some combinations of low and high T_m lipids (6, 9, 11). A closed loop miscibility gap is found when the extremely low T_m lipid DiPhytanoyl PC is

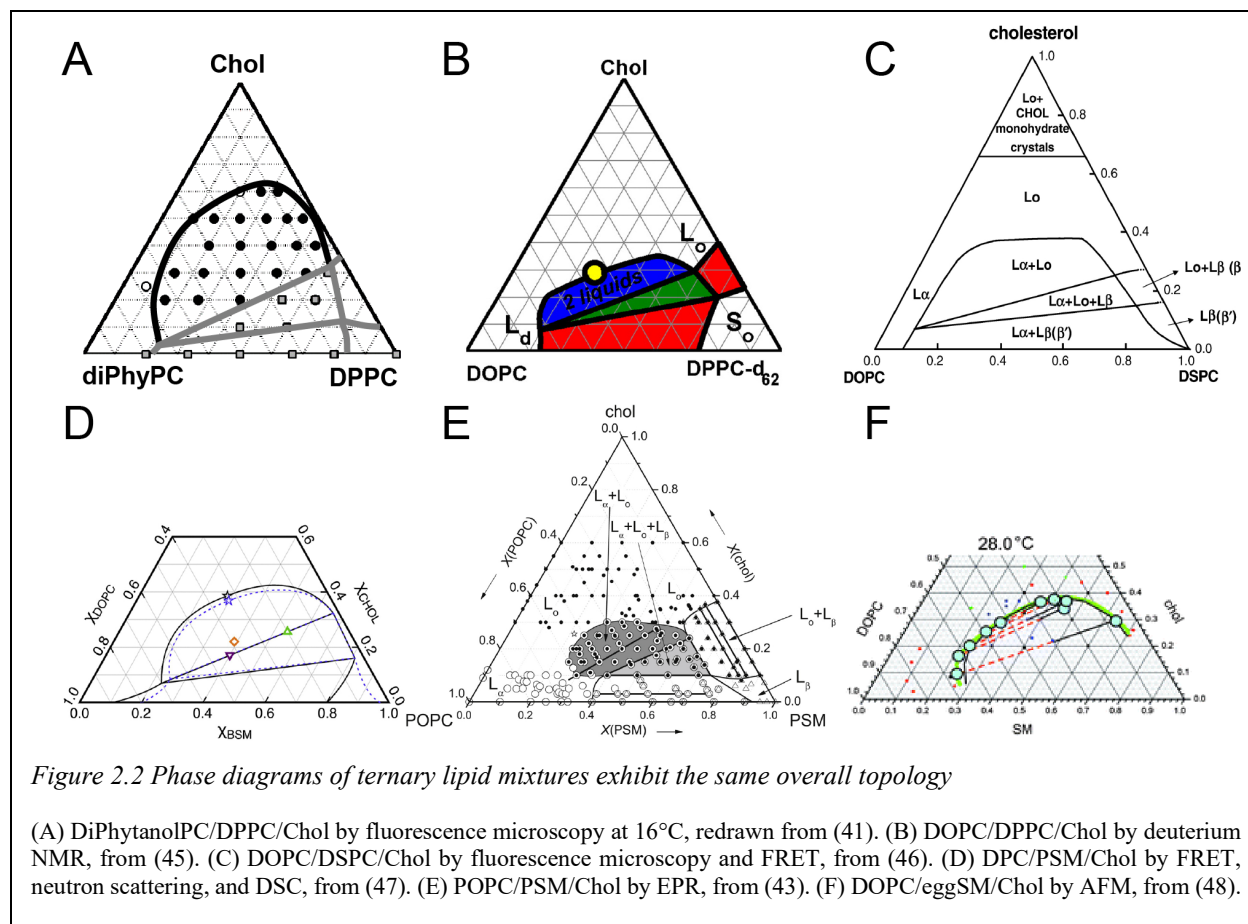
used (7), meaning that Lo-Ld coexistence occurs at temperatures above the T_m of the saturated component and there are two critical points. The phase behavior of the mixed system can depend on more than just the T_m of components. For example, sphingomyelin (SM) lipids are more effective at establishing coexisting phases as compared to glycerol-phospholipid lipids with PC headgroups (9, 18) even when the main chain transition occurs at similar temperatures for SM and PC lipids used.

Although the equilibrium thermodynamics description of the macroscopic miscibility transition is now largely accepted, questions remain regarding the thermodynamic basis of the microscopic heterogeneity also routinely observed in membranes containing cholesterol. These structures are frequently reported at temperatures and compositions where membranes remain uniform on a macroscopic scale using methods sensitive to molecular-scale organization such as FRET, fluorescence quenching, or ESR. Membranes can be tuned from a state with macroscopic phase separation to one with microscopic heterogeneity by raising temperature (140, 141), titrating in an additional component that disrupts the macroscopic phase transition (e.g. (132, 142, 143)), or by probing different ratios of the same lipid species at fixed temperature (141). Submicron structure is also reported in binary mixtures of saturated lipids and cholesterol (99, 144) and in some ternary membranes that do not exhibit macroscopic Lo-Ld immiscibility at any temperature or lipid ratio (11, 110, 129, 145). Recent experimental developments in purified membranes has begun to probe how leaflet asymmetry impacts phase separation and the presence of submicron structure in purified membranes (146, 147).

Experimental observations have motivated numerous theories to explain the presence of microstructure at thermodynamic equilibrium. Critical phenomena provide a possible mechanism to bridge macro- and micro-scales in the form of dynamic fluctuations (127, 148). Other theories enable static, finite-sized domains by including some repulsive mechanism to oppose formation of macroscopic domains (149–151). These theories too predict domains that span macro- to microscales. Recent experimental work has begun to directly test some of these theories explicitly (152).

2.3.2 Isolated membranes:

The first observations of macroscopic liquid immiscibility in vesicles blebbed directly from living cells came in 2007 (94), and these vesicles were named Giant Plasma Membrane Vesicles



(GPMVs) due to their close resemblance to the Giant Unilamellar Vesicles (GUVs) used widely in fluorescence microscopy investigations of purified lipid mixtures. Earlier work using a similar vesicle preparation had characterized their lipid and protein content by mass spectrometry (153), and had observed heterogeneity using ESR (119), a spectroscopic method that can detect heterogeneity on the molecular scale. Baumgart et al (94) detected the sorting of fluorescent lipid analogs and fluorescently tagged proteins with respect to phases in GPMVs at temperatures well below those where cells were grown, leading the authors to conclude that this phase transition was not relevant for cells under normal growth conditions.

Soon after, other methods emerged to isolate plasma membranes from cells and all could yield coexisting liquid phases (154, 155), although the conditions needed to achieve phase separation differed between methods used. A subsequent study found correlations between biochemically defined detergent resistant membranes and the liquid-ordered phase detected in GPMVs (19), and that the surface fraction of ordered phase at low temperature was altered by acute treatments to manipulate cholesterol levels in vesicles (19, 156). In all cases, miscibility

transition temperatures (T_{mix}) remained well below growth temperatures in isolated cells, emphasizing that such macroscopic domains were not likely to form under physiological conditions. A possible explanation came in 2008 when it was shown that freshly isolated GPMVs exhibited hallmarks of criticality, placing them close to a room temperature miscibility critical point (20). Over time, additional studies have documented how GPMV phase behavior is impacted by growth conditions in cells (21, 157, 158) and differentiation into other cell types (159, 160), and lipidomics analysis has begun to characterize the vast compositional complexity of these membranes (21, 158, 161).

While there are many similarities between the phase behavior observed in GPMVs and purified model membranes, there are key differences (16). The coexisting phases detected in GPMVs differ in their physical properties from their purified membrane counterparts. The ‘fluidity’ of phases are more similar in GPMVs compared to GUVs, as measured through diffusion of membrane components or using order sensing fluorophores that report on local hydration within the hydrophobic region of the membrane (162). These different physical properties can be sensed by incorporated proteins. Some transmembrane proteins, particularly those with palmitoylated cysteines, are observed to partition into the Lo phase in GPMVs whereas few are reported to partition into the Lo phase in purified membranes (163).

It should be noted that GPMVs are model membranes that differ from intact plasma membranes in important ways. Plasma membranes exist in close association to the actin cytoskeletal cortex, while GPMVs are missing polymerized cytoskeletal components and tend to be depleted in proteins that associate with the actin (164). Notably, the cell plasma membrane does not macroscopically phase separate even under conditions that cause GPMVs to phase separate, or in any known conditions, even when phase separated GPMVs remain attached to an intact cell membrane (165). GPMVs are depleted of PI(4,5)P₂ (166), which typically makes up several mol% of the inner plasma membrane leaflet in intact cells (167). A recent report demonstrates that isolated GPMVs are frequently permeable to large hydrophilic markers (168), indicating that their membranes contain long-lived defects. While cell membranes are asymmetric in their lipid and protein composition, at least some of this asymmetry is lost in the GPMV generation and isolation process (94, 166). The interpretation of GPMV experiments is also complicated by their sensitivity to methodological choices. For example, the most common method to prepare GPMVs involves incubating cells with a low concentration of formaldehyde and a reducing agent, and the choice of

the reducing agent can greatly impact the transition temperature and physical properties of phases of the resulting GPMVs (155, 169). This is due, at least in part, to the ability of some reducing agents to modify membrane proteins and lipids.

2.4 Phase separated domains are related to but different from ‘raft’ heterogeneity in intact cells.

Phases are robust, macroscopic entities with well-defined compositions, and with the exception of the yeast vacuole (170, 171), structures resembling liquid-liquid phase separation are not observed in intact cells. Instead, the vast majority of membrane domains in cells are microscopic and require significant perturbations in order to be visualized or isolated. In line with this, several recent studies have directly concluded that there is no evidence of a miscibility phase transition in intact cells when cells are examined through the lens of several different experimental observables (172, 173). Nonetheless, a large number of experimental studies provide strong evidence that the heterogeneity reported in intact cells is closely related to the macroscopic phase separation observed in GPMVs. Two excellent recent reviews discuss many of these findings, as well as some exceptions (174, 175). We highlight several lines of evidence below.

1. In many cases, proteins that are associated with live-cell heterogeneity or with detergent-resistant membranes are also found to partition into the Lo-like phase of GPMVs. In particular, single-pass transmembrane and peripheral proteins containing palmitoylations are more likely to be found in detergent resistant membranes and partition into the Lo phase in GPMVs, while proteins containing branched and unsaturated geranylgeranyl or prenyl groups tend to be solubilized by detergents and partition into the Ld phase in GPMVs. Recent work begins to extend to multi-pass proteins which can accommodate more complex protein/lipid interactions (176–178). It is likely that protein partitioning will be dictated by the identity of lipids that solvate transmembrane proteins in complex membranes (179, 180)
2. Cells actively tune their plasma membrane T_{mix} in response to changes in growth conditions and conditions that change T_{mix} lead to different phenotypes. For example, experiments indicate that cells in culture actively tune the miscibility phase transition temperature of their membrane to be a fixed temperature below their growth temperature (21), and that T_{mix} lowers in cells under conditions that inhibit cell growth (157). Other studies have documented that acute treatments with lipophilic small molecules or dietary lipids that alter T_{mix} correlate with changes in signaling outcomes, cellular differentiation, and even the general anesthetic response (159, 181, 182). While it is possible that these correlations with T_{mix} are a result of a mutual correlation with an unrelated

quantity, the accumulated evidence suggest that maintaining T_{mix} plays important roles in cellular processes.

3. In model membranes, T_{mix} predicts structure observed in single phase membranes. While model membranes appear homogeneous at the macroscopic scale above T_{mix} , the value of T_{mix} can predict the presence of structure at smaller scales, or under conditions where membranes are perturbed by clustering one component. A very recent report that documents the temperature dependence of both macroscopic phase behavior and microscopic heterogeneity in GPMVs under a range of perturbation conditions (169). This work found that the microscopic heterogeneity in GPMVs at elevated temperature was highly correlated with their macroscopic transition temperature T_{mix} , suggesting that the same could be true in intact cells. An older study finds that phase marking probes enrich in membrane domains stabilized through adhesion even well above T_{mix} (156). These and other studies have led to the idea that the value of T_{mix} is a measure of ‘raft’ stability under physiological conditions, even though cells do not appear to experience phase separation directly.
4. Theoretical arguments support that coupling to an intact cytoskeletal network will disrupt the macroscopic phase transition. It is well established that ‘quenched disorder’ abolishes first order phase transitions in two dimensional systems like membranes (183). The cortical actin cytoskeleton, which is linked to the plasma membrane through a system of adapter proteins, is likely to play the role of quenched disorder in the intact plasma membrane (165, 184). Experimental studies in model membranes support the main conclusions of this theory, that a broadly distributed cytoskeletal network disrupts macroscopic domains while stabilizing small-scale structure (185). Direct tests in intact cells have yet to be reported, but there is broad evidence for important connections between cortical actin and raft heterogeneity (186–190).
5. There is increasing evidence for phase-driven partitioning within intact cells. Recent work utilizing single molecule fluorescence methods are beginning to draw more direct connections between phases in vesicles and domains in cells (83, 191). These studies monitor the recruitment and exclusion of probes with respect to clustered proteins in intact cells and find that probe concentration in clusters mirrors partitioning with respect to phase separated domains in model membranes, although typically with a smaller magnitude.

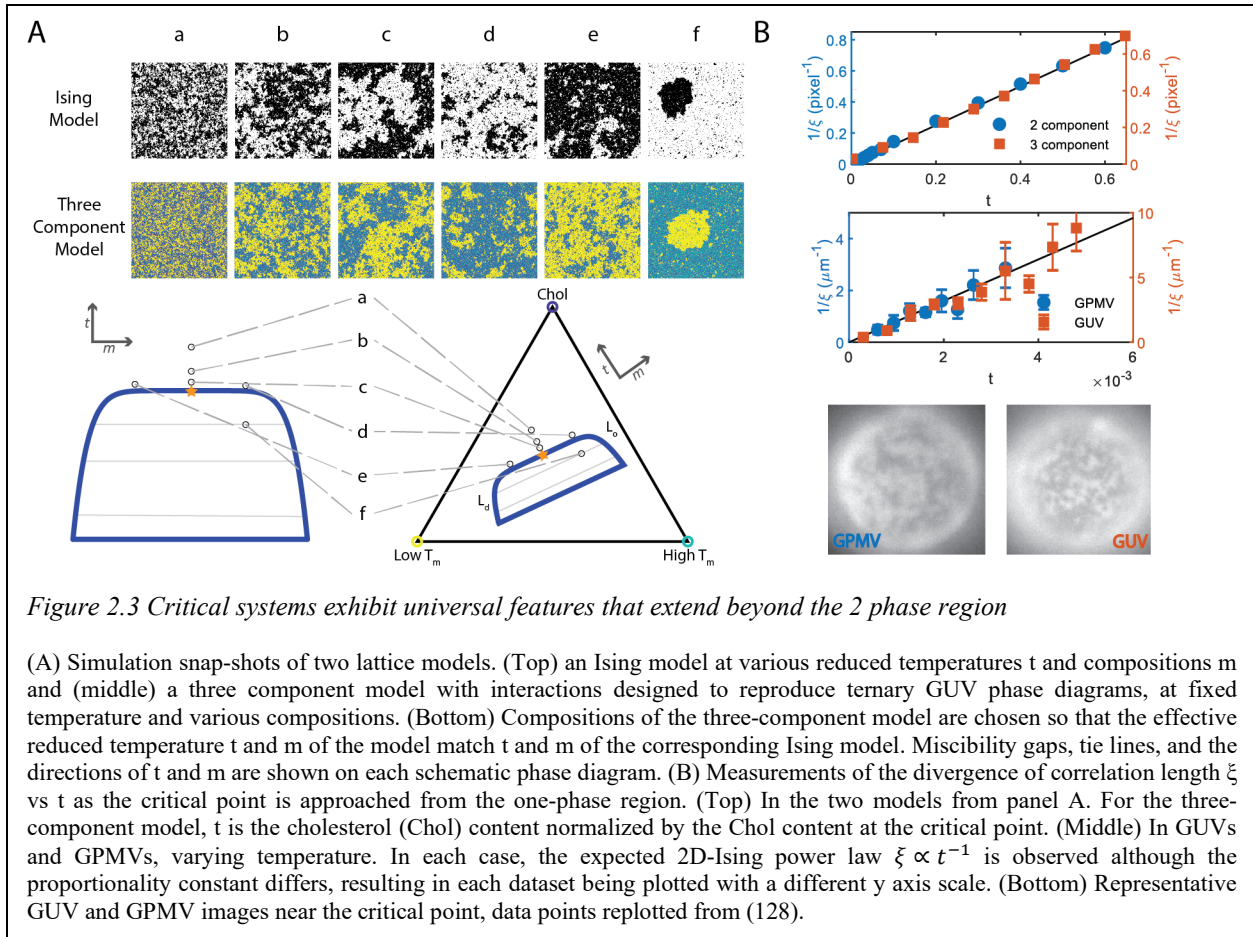
To explain these experimental developments, several theoretical avenues have been explored, with a goal of explaining the lack of a macroscopic phase separation in intact cells. Any complete biophysical model of membrane heterogeneity must first account for this fact, and therefore must be more complex than the simple phase separation picture. These explanations

include coupling to quenched disorder in the form of the cytoskeleton (165, 184), coupling to curvature in a way that produces a microemulsion (149, 151, 192), and nonequilibrium suppression of domain growth, for example by active lipid transport or remodeling of membrane-coupled actin structures (193, 194). Many of these models, including the critical phenomena that are discussed in the next section, are closely related to the canonical phase separation picture (195).

2.5 Criticality and its connection to plasma membrane

Equilibrium critical phenomena are special relationships between thermodynamic properties of a system that emerge due to a diverging correlation length of the system in the vicinity of a critical point. Key physical quantities such as the heat capacity, the susceptibility, and the interfacial energy between domains also exhibit specific behaviors as the critical point is approached. These properties of critical systems are universal, meaning that the dominant behavior is governed by a small number of effective parameters, regardless of the complicated details of microscopic interactions (Figure 2.3). The physics of criticality has been covered in detail in many textbooks (22, 196–198). A useful introduction for biophysicists is given in (199).

The physics of criticality has developed over the last 200 years, beginning with the observation of a critical temperature for several liquids, above which there is no liquid-gas transition (200). The van der Waals equation of state (1873) was the first model that featured a “critical point” with this property (201). Ornstein and Zernicke, in 1918 (202) formalized the study of spatial fluctuations in liquid-gas systems, a key development. In the first half of the 20th century, precise measurements of near-critical liquid-gas systems (203) as well as magnetic systems (204) indicated that classical equations of state were inadequate to explain the near-critical region of phase space, for example that the shape of the phase liquid-gas coexistence curve is qualitatively different than predicted by van der Waals. In 1944, Lars Onsager exactly solved the 2d Ising model, which allowed for the study of critical exponents in that system (205). In 1952, Yang and Lee observed that phase transitions correspond to non-analyticities of the partition function in the thermodynamic limit (206), and Widom and others proposed power-law scaling for various quantities at critical points, and derived relationships between different critical exponents from that hypothesis (207). Real-space (Kadanoff) and momentum-space (Wilson) renormalization group methods explained the emergence of these power laws and provided avenues for computing approximate values for the critical exponents in general systems (208, 209). They also set a



uniform framework for other conceptual advances. A broader class of nonequilibrium critical points can also be defined, generalizing this equilibrium concept. Biological nonequilibrium critical points have received considerable attention in recent years, see, e.g. Mora and Bialek for a review (210).

Coexistence regions generically terminate in a critical point except in special circumstances, and the Lo-Ld miscibility of purified and isolated membranes is no exception. On the phase triangle shown in Figure 2.1B, the critical point occurs on the high-cholesterol edge of the Lo-Ld miscibility gap. Initial evidence of critical behavior in membranes came from NMR studies of multilamellar vesicles, which detected enhanced line-broadening in the vicinity of known critical points (127). This was attributed to the diffusion mediated exchange of lipids between fluctuations with sub-micron dimensions. Later work directly visualized micron-sized critical fluctuations (above the critical temperature; T_C) and fluctuating phases (below T_C) in GUVs of purified lipids, providing evidence that membranes belong to the 2D Ising model universality class, meaning that fluctuations exhibited a temperature dependence that is universal to two

dimensional systems with a one-dimensional order parameter (148). Subsequent measurements confirmed this observation in supported membranes by atomic force microscopy (130), and probed the dynamics of critical membranes (211). At this same time, it was discovered that isolated GPMVs also exhibited critical behaviors in the vicinity of a room temperature critical point (20). Again, fluctuations were consistent with the 2D Ising model universality class. These observations and their implications have been reviewed in greater detail previously (199).

One of the key features of a critical point is that its fingerprints extend well beyond the phase transition itself (Figure 2.3). An important parameter is t , the difference between the temperature of the system and the critical temperature (T_c) normalized by T_c in units of Kelvin. The correlation length ξ , or characteristic size of critical compositions fluctuations, is predicted to vary as $\xi(t) = \xi_0/t$, where ξ_0 is a parameter with dimensions close to the size of molecules in the system and in membranes was measured to be roughly 1 nm. Note that as $T \rightarrow T_c$, $t \rightarrow 0$, so that the correlation length becomes infinite. Extrapolating this relationship using a room temperature critical point ($T_c = 22^\circ\text{C} = 295\text{K}$), then 20 nm sized fluctuations are expected at 37°C , which corresponds to $t = 0.05$. This prediction is in good agreement with recent experimental work probing heterogeneity in GPMVs by FRET, which detects evidence for larger than 10 nm structure in GPMVs over this same temperature range (169). Similar observations have also been made in purified model membranes (212). Note that t need not be a physical temperature. Instead it is any trajectory in the phase diagram that runs perpendicular to tie-lines close to the critical point. Thus, while some Ising model images of Figure 2.3A are obtained by varying temperature in the model, the corresponding three-component lattice model images are obtained by varying composition at fixed temperature, as indicated in the phase diagram.

Another physical property that can extend well beyond the phase transition itself is the susceptibility (χ). The susceptibility measures how large a local composition difference arises from a local force applied to components of one of the phases, e.g. by clustering components that prefer Lo lipids. In other words, in a highly susceptible membrane, a domain of distinct local composition can be stabilized by only clustering a small subset of components, or by weakly biasing the concentration of many components that have the same order preference. In the Ising universality class, $\chi \propto t^{-7/4}$. This too has experimental support in vesicles, where robust domains are stabilized well above T_{mix} by organizing a small subset of components by an actin network or

streptavidin crystal that partially decorates a vesicle surface (187, 213), or by adhesion to a supported membrane (156).

Direct theoretical predictions of critical phenomena such as the scaling of the correlation length and the magnitude of the susceptibility are quite useful for predicting consequences of perturbations to membrane heterogeneity, but the theory quickly becomes intractable when coupled to more complex biological phenomena. Statistical mechanical lattice models based on the Ising model can be useful in this situation. These models typically only contain two components (up and down ‘spins’) positioned on a lattice, where the components at the lattice sites are either allowed to change identity (such that the composition or ‘magnetization’ can vary) or are allowed to exchange with other sites on the lattice (such that the composition remains fixed). Universality guarantees that, so long as the system is close to the critical point, the Ising model captures the relevant mesoscopic heterogeneity of the membrane, for appropriate choices of the Ising reduced temperature t and magnetization m (Figure 2.3). That is, the Ising model accurately recapitulates the thermodynamics of the effective Lo order parameter at length-scales beyond a few lipid diameters, despite the extreme simplicity of the microscopic interaction in the model – a simple nearest-neighbor interaction potential. As a result, when a biological system is coupled to the Lo order parameter, an Ising model modified to include this coupling is expected to reflect the relevant biophysical phenomena. Past work has used this approach to model the coupling of fluctuations to cortical actin (165), to explain changes in phosphorylation steady states upon clustering of a component, for various values of t and m (83, 84), and to predict how proximity to the critical point affects conformational state equilibria of proteins whose boundaries are sensitive to lipid order (214).

2.5.1 Evidence for criticality playing a role in cells:

If intact plasma membranes exhibit similar heterogeneity to that observed in GPMVs, it could easily be relevant to the biological function of membrane proteins. An important line of evidence that criticality plays a role in biological function has come from the tuning of the GPMV critical point. For a system to be near a critical point in the first place, two parameters must be tuned, corresponding to t (temperature) and m (composition) of the (fixed-composition) Ising model. In the extremely large space of lipid mixtures of varying composition, there are many critical points – an $n - 2$ -dimensional manifold in the n -dimensional space. However, there is no

generic reason that tuning the concentration of any given lipid will correspond to tuning just t , or just m , or neither – general perturbations will affect both t and m . Thus it is somewhat surprising that the cell arrives near a critical point if it constructs its membranes without explicitly or implicitly tuning to the critical point, given that lipid composition is modulated by a wide variety of perturbations. In other words, it would be surprising if plasma membrane composition is near-critical simply by coincidence. Furthermore, at least in certain cases, eukaryotic cells adapt to perturbations in ways that preserve the distance to the critical point, and corresponding physical properties. Zebrafish cells cultured at a range of temperatures from 20-32°C produce GPMVs with correspondingly altered T_c (21).

The concept of a high susceptibility near a critical point is useful in interpreting recent single molecule and super-resolution studies documenting the partitioning of phase marking probes to protein clusters in intact cells (83, 84, 191). In these studies, antibodies are used to cross-link a membrane component that prefers either the Lo or Ld phase, then the differential partitioning of probes is monitored with respect to these domains. When proteins are clustered that themselves prefer the Lo phase, then probes that also prefer Lo tend to be recruited and those that prefer Ld tend to be excluded. In contrast, when proteins that prefer the Ld phase are clustered, then probes that prefer Lo are excluded and probes that prefer Ld are recruited. Similar to experiments with vesicles adhered to supported membranes (156), the act of clustering a protein or peptide biases concentration of many components in ways that can be detected when membranes have high susceptibility. In some cases the extent of probe partitioning approaches that observed in phase separated vesicles (84), while in others the sorting of components is much weaker (83). These differences could arise from differences in the coupling of protein clusters to membranes, or differences in the susceptibility of the membrane in different experimental systems.

2.6 Criticality as it relates to biological function

Since the inception of the raft hypothesis, the functional relevance of membrane domains has focused on their ability to compartmentalize protein and lipid components so that they can optimally function within biochemical networks (93). Critical phenomena are in many ways consistent with this framework. A super-critical membrane contains domains resembling ordered and disordered phases, and components that partition with the same phase will colocalize within these domains. The fluctuations are small and dynamic, consistent with evolving descriptions of

rafts over the decades (36, 215–217), but fluctuations alone are not an effective means to strongly colocalize or confine membrane components. This new reality requires us to move beyond the simple mechanisms proposed in the early raft literature to propose and test mechanisms that exploit the unique material properties of critical systems. Several proposals are highlighted in Figure 2.4 and described below.

2.6.1 Interactions between proteins:

Composition fluctuations can mediate forces between membrane proteins, via a process termed ‘critical Casimir forces’ first described between conducting plates in vacuum (218). In essence, proteins will feel an effective attractive potential if they partition into the same phase because their coming into close proximity allows them to share the same local lipids, as shown in Figure 2.4A (219, 220). In contrast, proteins that prefer different environments will feel an effective repulsion, since there is an energetic cost to mixing their local environments. These potentials are weak (on the order of the thermal energy $k_B T$) but have a range given by the correlation length and the size of the protein or protein cluster. This is long-ranged compared to other interaction modes experienced by membrane proteins such as curvature, electrostatics, and van der Waals potentials. It is notable that repulsion of components that prefer different phases has a larger magnitude than attraction between components that prefer the same phase. The Casimir force may contribute to the stability of protein assemblies, including phase separated polymer droplets that assemble on membranes. The Casimir force is also expected to alter biochemistry occurring at the membrane, by increasing or reducing the rates at which the proteins encounter one another. It is tempting to speculate that one functional role of palmitoylation, the post-translational modification that places a saturated acyl chain on proteins, is to tune the magnitude of this Casimir force for specific protein species.

2.6.2 Susceptibility to receptor clustering:

The high susceptibility of a critical membrane provides a means for the cell to sense a redistribution of a subset of membrane constituents by an external force (Figure 2.4B). Clustering a membrane protein that prefers one phase would bias the local lipid composition in proportion to the heightened susceptibility of the system. This effect can impact biochemical reactions that take place within these clusters, drastically altering the chemical steady state of the system. We

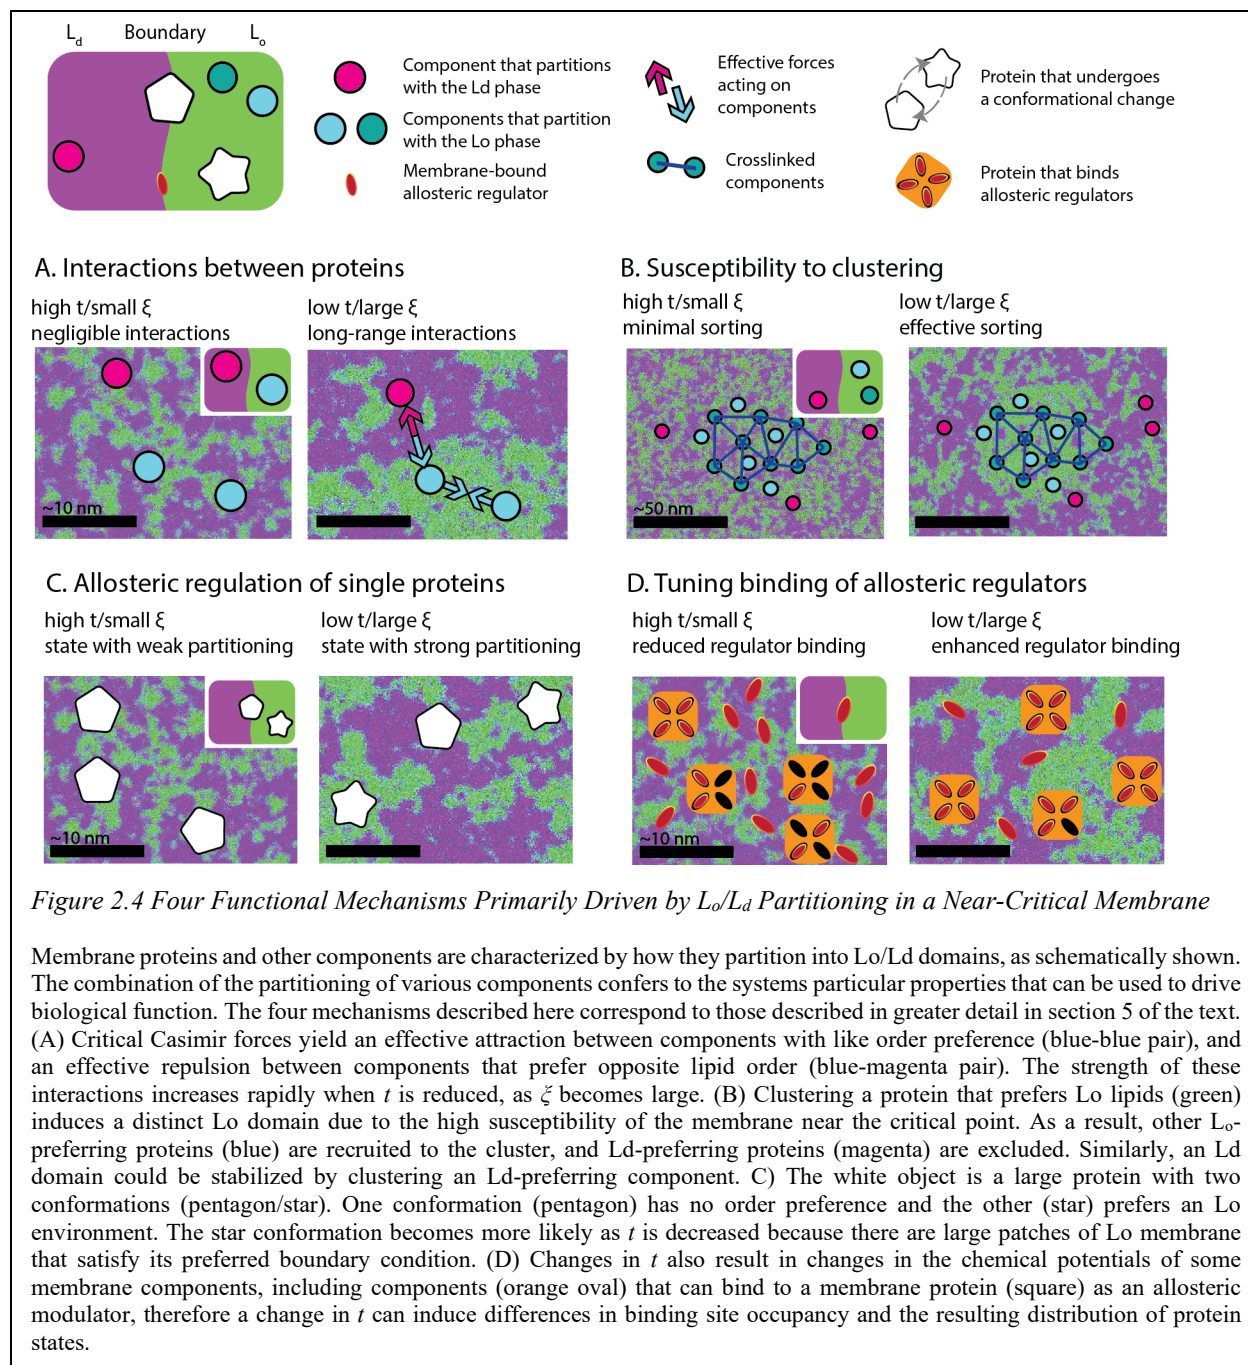


Figure 2.4 Four Functional Mechanisms Primarily Driven by L_o/L_d Partitioning in a Near-Critical Membrane

Membrane proteins and other components are characterized by how they partition into L_o/L_d domains, as schematically shown. The combination of the partitioning of various components confers to the systems particular properties that can be used to drive biological function. The four mechanisms described here correspond to those described in greater detail in section 5 of the text. (A) Critical Casimir forces yield an effective attraction between components with like order preference (blue-blue pair), and an effective repulsion between components that prefer opposite lipid order (blue-magenta pair). The strength of these interactions increases rapidly when t is reduced, as ξ becomes large. (B) Clustering a protein that prefers L_o lipids (green) induces a distinct L_o domain due to the high susceptibility of the membrane near the critical point. As a result, other L_o -preferring proteins (blue) are recruited to the cluster, and L_d -preferring proteins (magenta) are excluded. Similarly, an L_d domain could be stabilized by clustering an L_d -preferring component. (C) The white object is a large protein with two conformations (pentagon/star). One conformation (pentagon) has no order preference and the other (star) prefers an L_o environment. The star conformation becomes more likely as t is decreased because there are large patches of L_o membrane that satisfy its preferred boundary condition. (D) Changes in t also result in changes in the chemical potentials of some membrane components, including components (orange oval) that can bind to a membrane protein (square) as an allosteric modulator, therefore a change in t can induce differences in binding site occupancy and the resulting distribution of protein states.

have studied this effect in the context of B cell receptor (BCR) signaling (83, 84). Here, the act of clustering the BCR or another ordered membrane component by an extracellular ligand stabilizes an ordered domain that contains a higher local concentration of kinase and a lower concentration of phosphatase than the membrane as a whole. This establishes a local environment that favors receptor phosphorylation and activation. In principle, this class of activation mechanism could contribute to a wide range of signaling pathways that are initiated by receptor clustering at the cell

surface. This type of mechanism could also play a role in establishing biochemical environments in membrane regions where components are organized by processes occurring at the inner plasma membrane leaflet, such as at junctions between the ER and plasma membrane (221), or at sites where scaffolding adaptor proteins are anchored to membranes such as in neuronal synapses (222).

2.6.3 Allosteric regulation of single proteins:

Beyond contributing to the organization of proteins, the functioning of single proteins can also be impacted by the size and stability of fluctuations in the membrane. One mechanism that has been proposed requires that two conformational states of the protein in question have different boundary lipid preferences for Lo or Ld lipids (214). If that is the case, then a change in T_C will differentially affect the free energies of the two conformational states. Roughly, a spatially extended lipid preference carries a free energy cost that decreases near T_C , so that a conformational state with strong order preference becomes more probable when fluctuations are large compared to the protein diameter. This model was proposed to explain striking correlations between the T_C -altering effects and anesthetic or anesthetic-reversing potencies of a wide range of treatments, including short and long-chain n-alcohols and hydrostatic pressure (181, 182).

2.6.4 Tuning binding of allosteric regulators:

The chemical potential μ of a component is the thermodynamic parameter that controls the proclivity of that component to enter or exit a system. For example, its availability to bind in a binding pocket. Formally, the chemical potential is the increment of free energy to move one particle into the system from a particle bath. Equivalently, the chemical activity $a \propto e^{\mu/kT}$ can be used. In an ideal gas or ideal dilute solution, the chemical potential has a simple logarithmic relationship to concentration and linear in temperature (so that activity is proportional to concentration), and insensitive to the concentrations of other components(223). However, a near-critical mixture is far from ideal – the critical point is precisely where weak cooperative interactions between the many components lead to strong effects (196). Therefore, we expect strong relationships between the chemical potentials of different components, especially when those components modulate T_C .

Many transmembrane proteins have been shown to be modulated by binding to membrane components, prominently including cholesterol (223, 224) and phosphatidylinositol lipids (225),

and many other signaling lipids (226). If the chemical potential of some of these components is strongly modulated by concentration changes of other components, binding site occupancy will also vary, and we expect to see changes in protein functions that depend on binding of those components. Recent work by Ayuyan and Cohen has developed sensitive methods for measuring and controlling the chemical potential of cholesterol in the plasma membrane, and found that cholesterol chemical potential varies by $\sim 2 k_B T$ in different physiologically relevant cellular conditions (223). That amount is certainly adequate to induce substantial changes in binding site occupancy. Work remains to be done to explore if these differences can be attributed in any way to the critical phase transition, but it is exciting to speculate that perturbations that change membrane criticality may act indirectly by impacting the activity of membrane components.

2.6.5 Criticality coupled to other processes describes a broad array of raft phenomena:

The Ising universal critical phenomena are a good start for understanding membrane heterogeneity, but they are also clearly insufficient to explain all phenomena. As stated above, critical phenomena alone are not expected to give rise to regions of tight clustering or confinement of proteins and lipids, as is sometimes attributed to membrane domains. This said, it is possible that the local membrane environment can impact the conformational states sampled by membrane proteins in ways that facilitate binding through stronger protein binding sites. This type of synergistic effect could underlie a range of cholesterol dependent processes observed at the plasma membrane, including for example the transient pinning observed in studies of membrane protein and lipid dynamics (227).

Another example where there is potential for synergy between criticality and other organizing principles relates to the ‘Active Composite Model’ proposed by Mayor and coworkers (228). This model posits that the plasma membrane interacts with a heterogeneous cortical actin network, composed of both active and passive components. The passive components largely resemble quenched disorder, as described earlier in this review. The active component is composed of motor driven short actin filaments that can actively drive certain membrane proteins and lipids into close proximity, coupling across membrane leaflets. Considering this model in a critical membrane provides a simple means to correlate domain structure across leaflets without requiring strong interactions such as interdigitation, since the cooperativity inherent in a phase transition can amplify weak couplings that may be present. Moreover, the high susceptibility of a critical

membrane allows it to robustly remodel when external forces are applied, including those originating from the actin cortex.

More broadly, we envision that plasma membrane criticality is only one of several organizing principles that contribute to plasma membrane functions. It could be that interactions mediated by curvature and electrostatics superimpose with those mediated via criticality to define the plasma membrane interactome. It is also possible that there is interesting cross-talk between these various interaction modes. For example, studies have shown that the sorting of lipids into curved membranes can be mediated by the binding of curvature-sensing proteins that have preferences for one membrane phase (229). Proteins and peptides can also organize lipids through electrostatics, which in turn can stabilize domains impacted by fluctuations of ordered and disordered phase lipids. The broader implications of this potential cross-talk is largely unexplored, and could give rise to qualitatively new phenomena accessible to cell membranes (230).

2.7 Concluding remarks

While it has long been appreciated that plasma membrane lipids are capable of intriguing, non-ideal behaviors, much of the past literature is clouded by imperfect methods and an incomplete conceptual framework to conceptualize experimental observations. This backdrop led to controversial and often unphysical descriptions of lipid rafts. The past decade or so has brought key advances, including membrane isolations that largely preserve plasma membrane protein and lipid content and super-resolution imaging methods that do not suffer from the same pitfalls that plagued early raft research. Together with this, the membrane community has begun to appreciate the rich phenomena that naturally occur near miscibility critical points, many of which exhibit strong parallels with long-standing observations in both the membrane-biophysics and membrane biology literatures. Moving forward, the challenge will be to isolate these effects to enable a definitive measurement of the role of criticality in cell membranes, and to explore how these immiscibility-mediated interactions work alongside other physical and biochemical organizing principles to contribute to the rich array of biological functions at the cell surface.

Chapter 3 Chemical Potential Measurements Constrain Models of Cholesterol-Phosphatidylcholine Interactions

This chapter presents measurements of the chemical potential of cholesterol in mixtures with one or two phospholipids, with particular attention to membrane compositions in and near the liquid-liquid coexistence region, as well as to the non-ideal cholesterol dependence of the chemical potential curves. The chemical potential of cholesterol is of great interest to membrane biology, where cholesterol is thought to act as an allosteric modulator of many plasma membrane proteins, in addition to its role in determining the physical characteristics of the membrane. My contributions to this work are spread throughout, including helping develop the methods and analysis that are used, interpreting the resulting measurements in terms of principles of physical chemistry, and modeling the measurements with thermodynamic models. Kathleen Wisser and Anna Gaffney performed the experiments, and Taylor Schaffner and Benjamin Machta contributed to the modeling efforts, especially the models that involve stoichiometric complexes. Sarah Veatch devised the central questions of the study and was involved in all aspects of the process. A version of this chapter has been published at *Biophysical Journal* (231).

3.1 Abstract

Bilayer membranes composed of cholesterol and phospholipids exhibit diverse forms of non-ideal mixing. In particular, many previous studies document macroscopic liquid-liquid phase separation as well as nanometer-scale heterogeneity in membranes of phosphatidylcholine (PC) lipids and cholesterol. Here, we present experimental measurements of cholesterol chemical potential (μ_c) in binary membranes containing dioleoyl PC (DOPC), 1-palmitoyl-2-oleoyl PC (POPC), or dipalmitoyl PC (DPPC), and in ternary membranes of DOPC and DPPC, referenced to crystalline cholesterol. μ_c is the thermodynamic quantity that dictates the availability of cholesterol to bind other factors, and notably must be equal between coexisting phases of a phase-separated mixture. It is simply related to concentration under conditions of ideal mixing, but is far from ideal for the majority of lipid mixtures investigated here. Measurements of μ_c can vary with phospholipid composition by $1.5 k_B T$ at constant cholesterol mole-fraction implying a more than

five-fold change in its availability for binding receptors and other reactions. Experimental measurements are fit to thermodynamic models including cholesterol-DPPC complexes or pairwise interactions between lipid species to provide intuition about the magnitude of interactions. These findings reinforce that μ_c depends on membrane composition overall, suggesting avenues for cells to alter the availability of cholesterol without varying cholesterol concentration.

3.2 Introduction

Phospholipid bilayer membranes containing cholesterol are complex fluids that exhibit non-ideal mixing of components that is detectable by a broad range of experimental methods. Non-ideal mixing can take the form of phase separation or nanoscopic domains detected by methods such as Fourier resonance energy transfer, electron spin resonance, nuclear magnetic resonance, neutron scattering, or fluorescence microscopy (e.g. (3, 11, 13, 99, 145, 232–234) reviewed in (92)). This non-ideality is a ubiquitous feature of these membranes, and is therefore expected to contribute to their chemical and material properties.

One fundamental biophysical property of components within membranes is their chemical potential (μ), which describes the thermodynamic availability of these components. The quantity $A = \exp \mu / k_B T$ is called the chemical activity of the component, where $k_B T$ is Boltzmann's constant times temperature, the thermal energy of the system. The chemical activity of a component determines how available it is to bind and influence targets at equilibrium. In dilute solutions, concentration-dependence of the chemical potential μ of a component is determined by its mixing entropy alone because solutes interact nearly exclusively with identical solvent molecules, producing an activity A that is simply proportional to its concentration. Enthalpic deviations from this linear trend can emerge when a component's interactions depend on local composition, with variations comparable to the thermal energy. Nontrivial entropic contributions to the free energy also contribute to nonlinearities when multicomponent systems are heterogeneous. In an extreme example, a condition for multiphase coexistence is that μ for each component is identical across phases, even when phases consist of components at vastly different concentrations and/or physical states.

In this study, we measure the chemical potential of cholesterol (μ_c) within membranes containing different phospholipids or mixtures of phospholipids at room temperature. μ , like any potential, is only thermodynamically defined up to a numerical constant and therefore numerical

values need to be reported with respect to a reference state that defines $\mu = 0$. Here, μ_C is measured against a reference state of cholesterol crystals, which means that $\mu_C = 0$ (and $A_C = 1$) indicates the situation of cholesterol being fully saturated, as it would be for a system equilibrated with cholesterol crystals. Measuring μ_C provides a window into the molecular interactions that underlie heterogeneity and phase separation in these well characterized systems. We measure μ_C in large unilamellar vesicles (LUVs) by suspending them in aqueous solutions of methyl β cyclodextrin (M β CD), a sugar that binds cholesterol making it water soluble, using an equilibration scheme similar to Niu and Litman (34). We extend this method to measure μ_c at substantial cholesterol fractions, and calibrate it against cholesterol crystals following the scheme of Ayuyan and Cohen (223), who carried out similar measurements in cells. These experimental methods make it possible to report μ_c for a wide range of membrane compositions, adding to numerous past measures of cholesterol-lipid interactions in model membranes (34, 33, 235, 35, 236–240). Calibrating experimental findings to a common reference state simplifies a rigorous comparison of trends across measurements and experimental systems. We compare experimental measurements against several models of cholesterol-lipid interactions, with the goal of better understanding how these interactions give rise to phase separation and heterogeneity in multicomponent membranes.

3.3 Materials and Methods

Materials:

Methyl β -cyclodextrin (M β CD) (CAS: 128446-36-6) was purchased from TCI Chemicals (Portland, OR). Cholesterol (Chol), 1,2-dipalmitoyl-sn-glycero-3-phosphocholine (DPPC), 1,2-dioleoyl-sn-glycero-3-phosphocholine (DOPC), and 1-palmitoyl-2-oleoyl-glycero-3-phosphocholine (POPC) were purchased from Avanti Polar Lipids (Birmingham, AL). Phospholipids were ordered as stock solutions in chloroform and cholesterol was a lyophilized powder.

N-(7-Nitrobenz-2-Oxa-1,3-Diazol-4-yl)-1,2-Dihexadecanoyl-sn-Glycero-3-Phosphoethanolamine (NBD-PE) was purchased from ThermoFisher (Waltham, MA). Cholesterol Oxidase from *Streptomyces* was purchased from MP Biomedicals (Santa Ana, CA) and 10-Acetyl-3,7-dihydroxyphenoxazine (Amplex Red) was purchased from Cayman Chemicals (Ann Arbor, MI). All other chemicals and supplies including Horseradish Peroxidase, Raffinose, Millex-VV

Syringe Filter Units, and Amicon Ultra-Centrifugal Filter Unit, 30KDa, were purchased from Sigma-Aldrich (St. Louis, MO) unless otherwise indicated.

3.3.1 Preparing M β CD solutions and saturated solutions of M β CD and M β CD/Chol.

Solutions containing M β CD and M β CD/Chol were prepared as described previously (223) with only minor modifications. Briefly, M β CD (5 mg/ml; 3.8 mM) was dissolved in a buffered saline solution containing 20mM HEPES, 135 mM NaCl, 5 mM KCl, 1 mM MgCl₂, 1.8 mM CaCl₂, 5.6 mM Glucose at pH 7.4. All solutions containing M β CD were degassed for 30 minutes under vacuum and stored under Argon gas to prevent oxidation.

Fully saturated M β CD/Chol solutions were prepared by incubating M β CD solutions with excess cholesterol crystals. These were prepared by first wetting powdered cholesterol (approximately 18 mg for a 20 mL solution) with 200 μ L of methanol and drying under Argon followed by 30 min under vacuum to remove residual solvent. This dried cholesterol aggregate was submerged in the M β CD solution described above, then sonicated two times for 2 min each using a Branson bath Ultrasonifier (model S450A, Process Equipment & Supply Inc, North Olmsted OH) to produce a cloudy suspension of small cholesterol crystals. This suspension was stored at room temperature under Argon with continuous rotation and was typically used at least 12 h after preparation. Immediately prior to an experiment, the equilibrated solution was filtered through stacked 0.2 and 0.1 μ m Millex-VV Syringe Filter Units to remove cholesterol crystals. Within minutes of sonication, M β CD solutions in contact with cholesterol crystals equilibrate, becoming saturated with cholesterol (Supplemental Figure 3.6 – see Section 3.7 for Supplemental Figures and Tables). Aqueous M β CD solutions with a range of cholesterol % saturation levels were prepared by diluting the fully (100%) saturated solution described above with 5 mg/ml M β CD solutions without cholesterol.

The cholesterol standard was prepared by suspending 100 μ g of lyophilized cholesterol in 10 ml of M β CD solution. Because this is such a small quantity, we first suspended dry cholesterol in solvent, measured the appropriate volume corresponding to 100 μ g cholesterol, and then lyophilized away the solvent prior to hydration in M β CD solution. Similar results were obtained when cholesterol in solvent was dried to a thin film under nitrogen then placed under vacuum prior to hydration in M β CD solution.

3.3.2 Vesicle preparation

Lipid mixtures containing 0.1 mol% NBD-PE were assembled in chloroform, dried under Nitrogen while vortexing to form a thin film, then placed under vacuum for 30 minutes to remove residual solvent. The lipid film was hydrated to between 1 and 5 mg/ml in an aqueous buffer containing 300 mM raffinose and 5 mg/ml M β CD, vortexed, then large unilamellar vesicles (LUVs) were formed by extruding 15 times through a 100 nm Polycarbonate Membrane (part 610005) using mini extruder (part 610000) both from Avanti polar lipids (Birmingham, AL). When DPPC was incorporated, lipids were hydrated and extruded at elevated temperature (>60°C) to prevent phase separation. Final molarity of lipid in suspension ranged between 1 and 8 mM, depending on lipid content and stock concentration, but were typically close to 7 mM to give a final lipid:M β CD molar ratio close to 2:1. M β CD is included within the hydration buffer to ensure that M β CD concentration remains constant when vesicles are diluted into calibrated M β CD containing buffers. Including M β CD in the hydration buffer also allows for cholesterol exchange between any multilamellar structures still present after extrusion (241). NBD lipids are not detected in the M β CD containing aqueous phase when vesicle suspensions are filtered to remove vesicles (Supplemental Figure 3.7), in good agreement with past studies demonstrating that much higher M β CD concentrations are required to solubilize lipids (34, 35, 242, 243).

3.3.3 Equilibration and isolation of cholesterol in vesicles and aqueous M β CD solutions.

LUVs were diluted into aqueous M β CD solutions with different % cholesterol saturation levels. For most measurements, 25 μ l of LUVs were added to 475 μ l of M β CD solution with a specified % cholesterol saturation to produce a final lipid concentration of 250 μ g/ml lipid, which is between 300 μ M and 400 μ M lipid for the lipid mixtures investigated. This results in a typical lipid: M β CD ratio of approximately 1:10. Vesicle suspensions were allowed to equilibrate for at least 1 h and up to 24 h under agitation at room temperature to allow for cholesterol to exchange between vesicles and M β CD in the aqueous phase. After equilibration, 100 μ l of the vesicle suspension was extracted and the remaining 400 μ l was pelleted through centrifugation (18,000 \times g for 90 min at 23°C). 350 μ l of the supernatant was extracted representing a suspension depleted in vesicles. The remaining 50 μ l was mixed with the pellet to produce a suspension enriched in vesicles. These three suspensions (Supplemental Figure 3.8A) were then assayed to determine their cholesterol and phospholipid concentrations. The data were plotted as three points on a graph of cholesterol concentration vs. phospholipid concentration (Supplemental Figure 3.8B). The three

points fall on a line because the three samples contain the same aqueous phase with varying LUV concentrations. The % saturation of cholesterol in the aqueous M β CD phase is reported by the y intercept (extrapolating to zero phospholipid) while the slope is the cholesterol to phospholipid ratio within LUVs (Supplemental Figure 3.8C). Measurements conducted on different days and with different initial conditions produce values that collapse onto a single curve for a given phospholipid composition, providing evidence that LUVs and M β CD solutions are equilibrated (Supplemental Figure 3.8D). Past work indicates that equilibration occurs within minutes (34), and this was the case for all but mixtures of DPPC and cholesterol, which did not change in composition after 1h at room temperature but did after 24 h (Supplemental Figure 3.9).

3.3.4 Measurement of cholesterol and phospholipid concentration

The cholesterol concentration of aqueous M β CD solutions and suspensions of vesicles within aqueous M β CD solutions were determined using the Amplex red (AR) cholesterol oxidase assay described previously (223) with minor modifications. Stock solutions of cholesterol oxidase (CO; 200 U/ml in PBS), horseradish peroxidase (HRP; 200 U/ml in Potassium Phosphate buffer, pH 5) were prepared according to manufacturer's recommendations. AR was stored at 5 mg/ml DMSO. The AR reaction buffer was prepared immediately prior to each measurement in the M β CD buffer by adding 1% v/v Triton-X-100, 2 U/mL HRP, and 75 μ g/ml AR either in the presence or absence of 2 U/ml CO. Including a detergent (Triton X-100) in the reaction buffer ensures that cholesterol within vesicles remains accessible to CO.

50 μ l of each vesicle suspension was plated in triplicate in a 96 well plate. Plates included a standard curve prepared from aqueous M β CD solutions with a range of % saturation levels. Plates also included a standard curve of LUVs in M β CD buffer with a range of dilutions. In some cases, samples were diluted with M β CD buffer prior to mixing with reaction buffer to ensure that readings would fall in a sensitive region of the standard curve. The fluorescence intensity of NBD was then measured using an iD3 Microplate Reader (Molecular Devices, San Jose, CA) with 458 nm excitation and emission measured between 510-550 nm. After recording NBD intensities, 50 μ l of AR reaction buffer was added to every well within the plate. In most cases, plates included 2 technical replicates for each sample with reaction mixture containing CO and a single replicate in the reaction mixture lacking CO. Including a replicate lacking the CO enzyme enables background subtraction on a sample-by-sample basis, isolating the effect of the CO enzyme. Plates

were sealed and incubated for 1h at 37°C, followed by at least 30 min at room temperature prior to recording AR fluorescence intensity with 545 nm excitation and emission between 600-650 nm.

To determine the % saturation of cholesterol in samples, fluorescence intensities from wells lacking CO were subtracted from values obtained from CO containing wells (ΔAR). The standard curve is nonlinear for the enzyme and cholesterol levels used, and we fit to the following form with fit parameters A and B:

$$\Delta AR = A(1 - \exp\{-B \times \%sat\})$$

A representative standard curve with fit is shown in Supplemental Figure 3.10A.

The standard curve was converted to units of cholesterol concentration using a cholesterol standard of 10 $\mu\text{g/ml}$ in 5mg/ml M β CD. We find that saturated solutions contain $158 \pm 30 \mu\text{M}$ cholesterol when equilibrated at room temperature, which corresponds to a 25 ± 5 M β CD molecules per cholesterol at saturation (Supplemental Figure 3.10B), in agreement with past measurements (27 ± 3 in 3 mg/ml M β CD at 37°C (223)). We speculate that this dilute incorporation of cholesterol in M β CD in part due to heterogeneity in methylation of the commercial M β CD reagent used, since past studies document the importance of methylation in controlling M β CD/cholesterol binding (244). Past work supports that, when bound, 1 cholesterol interacts with 2 M β CD molecules (35). Supplemental Figure 3.10C shows how the cholesterol concentration of fully saturated solutions varies with temperature.

To determine the phospholipid concentration of samples, cholesterol concentration was first evaluated in wells corresponding to the LUV standard curve. In most cases, the values obtained matched the expected concentrations given the vesicle stock concentration and the mole % of cholesterol in the initial preparation (Supplemental Figure 3.11A), especially considering some loss is expected during extrusion. Occasionally, cholesterol levels were lower than expected, indicating that LUVs were present at lower concentration, possibly due to instability of the vesicle suspension, especially after storage (Supplemental Figure 3.11B). To account for these errors, the phospholipid concentration within standard vesicle wells was inferred from the measured cholesterol concentration using the known cholesterol mole % of LUVs. Plots of NBD fluorescence intensity vs. inferred phospholipid concentration follow a linear trend (Supplemental Figure 3.11C), and a linear fit is used to infer phospholipid concentration from NBD intensity in the remaining wells. We estimate that errors in the cholesterol mole % of prepared vesicles is 2%, and this value is propagated to measurements of phospholipid concentration in all calculations.

In all cases, functions were fit to experimental data points using the fit() function in MATLAB (MathWorks, Natick, MA). Errors arising from uncertainty in the standard curve are applied to predicted values using the predint() function within MATLAB, using the optional input 'functional'. When used this way, predint() calculates confidence intervals on predicted values taking into account errors in parameter values but not scatter of data points. When appropriate, errors are propagated through computations to obtain the reported errors.

3.3.5 Calibrating cholesterol activity in M β CD/Chol mixtures.

The relationship between cholesterol chemical activity and cholesterol concentration using cholesterol crystals as the reference state was measured for aqueous solutions containing 5 mg/ml M β CD following the protocol described previously (223). Briefly, saturated solutions of cholesterol in hydrated hexadecane were prepared by dissolving excess cholesterol in hexadecane and equilibrating overnight. Cholesterol crystals were filtered, then the solution was hydrated in excess water and equilibrated for an additional 24 h, followed by a second filtration to remove additional cholesterol crystals formed due to its lower solubility in hydrated solvent (245). Hydrated hexadecane solutions covering a range of % cholesterol saturations were assembled by mixing saturated and cholesterol free hexadecane solutions at different volume ratios. M β CD solutions over the same range of % cholesterol saturations were assembled by mixing fully saturated and cholesterol free M β CD solutions at the same volume ratios. 500 μ L of the aqueous M β CD solution was then combined with 100 μ L of the organic hexadecane solution at each saturation level, capped under argon, and incubated overnight with shaking to equilibrate. After equilibration, the aqueous phase was separated through centrifugation and retained. Samples were agitated for 30 min along with a single polystyrene bead (1/8 inch Polyballs, Cat. No. 17175; Polysciences, Warrington, PA) to remove residual solvent. Finally, 50 μ L of each solution was then transferred to a multiwell plate in triplicate, along with a standard curve containing dilutions of the saturated cholesterol solution in M β CD buffer. Cholesterol content in each well was measured as described above.

Some hexadecane is soluble in the M β CD aqueous phase and some M β CD is soluble in hexadecane, leading to a slightly lower effective concentration of M β CD in solutions equilibrated against hexadecane compared to freshly prepared M β CD. As a result, the cholesterol concentration measured for the M β CD/Chol sample incubated with the fully saturated hexadecane solution is

slightly lower than in M β CD/Chol solutions incubated with cholesterol crystals alone. To account for this, we normalize to the fully saturated value, as shown in Supplemental Figure 3.12A. This applies the assumption that the cholesterol to M β CD ratio at saturation is independent of M β CD concentration for the narrow range of M β CD concentrations considered. Past work interrogated this issue further, demonstrating that similar calibration curves are obtained with different solvents when this treatment is applied (223). Normalized measured % cholesterol saturation of aqueous M β CD solutions were plotted against the known activity of solutions of cholesterol in hexadecane and fit to a Langmuir isotherm which takes the following functional form (Supplemental Figure 3.12B):

$$\%sat = \frac{FkA}{(1+kA)}.$$

The best fit values for k and F were 0.6 ± 0.1 and 2.6 ± 0.3 respectively. This function and associated errors were used to convert the measured % saturation of cholesterol in M β CD to cholesterol chemical activity.

3.3.6 Evaluating best fit parameters of mean-field models:

The following mean-field Gibbs free energies were used to generate expressions for μ_C which were then fit to experimental observations. The free energies, implied chemical potentials, and other properties of the models are described at length in supplemental notes in Section 3.8. In all cases, μ_C was evaluated as an appropriate derivative to match the definition of the chemical potential: $\mu_C = \frac{\partial G}{\partial N_C}$. More information can be found in Supplemental Note 1. Several of the following belong to the class of models known as regular solution models. General background for regular solution models is given in Supplemental Notes 2 and 3. In the following, the subscript i is either C, S, or U to denote cholesterol, saturated lipid or unsaturated lipid, respectively. N_i and μ_i^0 are the number and chemical potential offset of molecules of each species, and $N = N_C + N_U + N_S$ is the total number of molecules.

Complexes without additional interactions:

$$G = \sum_i N'_i (\mu_i^0 + k_B T \ln N'_i / N') + N_{CS} \ln \frac{N_{CS} / N'}{K_{eq}},$$

where N'_i is the number of molecules of components not found in complexes and the chemical reaction of complex association is: $qC + pS \rightleftharpoons C_qS_p$, with forward equilibrium constant K_{eq} . The N'_i are related to the total numbers of each lipid component by $N'_C = N_C - qN_{CS}$, $N'_S = N_S - pN_{CS}$, $N'_U = N_U$, and $N' = N'_C + N'_S + N'_U + N_{CS}$. N_{CS} is the number of complexes and is determined by minimizing G as a special case of the procedure described in Supplemental Note 4.

Complexes with additional interactions:

$$G = \sum_i N'_i (\mu_i^0 + k_B T \ln N'_i / N') + N_{CS} \ln \frac{N_{CS} / N'}{K_{eq}} + 2k_B N_{CS} (T_{CS-S} N'_S + T_{CS-U} N'_U) / N',$$

where T_{CS-S} and T_{CS-U} are the critical temperatures associated with complex-saturated lipid and the complex-unsaturated lipid binary systems. All other interaction terms are set to 0. See Supplemental Note 4 for a detailed explanation and derivation of this model.

Regular solution without complexes:

$$G = \sum_i N_i (\mu_i^0 + k_B T \ln N_i / N) + \sum_{i < j} z J_{ij} N_i N_j / N,$$

where J_{ij} is the interaction energy between the i^{th} and j^{th} component, and z is the lattice coordination number, which we take as 4 so that we can compare to lattice model simulations on a square lattice. Derivations for this model are laid out in Supplemental Note 3.

Models are fit simultaneously to all observations except those that were deemed to be in liquid-liquid or liquid-solid coexistence regions, by weighted nonlinear least squares with weights given by inverse square experimental errors. All fitting is conducted using MATLAB's fit() function. Weights are normalized to average to 1 overall. Phase separated points are excluded because mean field Gibbs free energies deviate from the true free energy in phase separated regions. All fits to μ_C are accomplished through a 2 step process. In the first fit, weights are determined from experimental errors in μ_C only. In the second fit, total errors are estimated by propagating experimental errors in the cholesterol mole % to μ_C using the initial fit. On plots, error bounds are estimated using predint(), which produces error estimates for predicted values.

3.3.7 Evaluating approximate phase boundaries and tie lines in mean field models:

We determine phase separated regions of mean field models by the approach of (246). Briefly, the phase separated region is the region of composition space for which the surface $G(x_C, x_S)$ is greater than the convex hull of this surface. A discretized convex hull of G is

determined using MATLAB's convexhull function, applied to G evaluated at a fine grid of compositions (x_C, x_S) , with uniform grid spacing of 2^{-10} . The resulting triangulation yields triangles with a long length:width ratio in the phase coexistence region. These long triangles are selected by choosing the triangles where this ratio is more than 20, and the long axis of triangles that meet this criterion are taken as approximate tie lines. The phase boundary is approximated as the (2-dimensional) convex hull of the endpoints of the tie lines. Finally, the chemical potential within the phase separated region is approximated as a linear interpolation across this region of the chemical potential evaluated at the endpoints of the tie lines.

3.3.8 Three-component lattice fluid simulations:

The regular solution model discussed above corresponds to a mean field approximation to a three-component lattice fluid with nearest-neighbor interaction Hamiltonian given by

$$H = \sum_{\langle i,j \rangle} J_{s_i s_j}$$

where $\langle i, j \rangle$ indicates that the sum is to be taken over nearest neighbors, s_i is the component at site i , and J_{kl} is the interaction energy between species k and species l . To compare to the full thermodynamics of this system, it is simulated using home-built C and MATLAB code implementing a fixed-composition Metropolis sampler for this Hamiltonian. The fixed composition criterion is enforced by only proposing updates that swap the components at two sites. The simulations are performed on a periodic $L \times L$ site square lattice with $L=256$, with $1000 \times L \times L$ swaps proposed between samples, and the first 50 samples are discarded to allow for equilibration. For each simulated condition, μ_C is obtained from 7 simulation samples following a procedure derived in detail in Supplemental Note 5. In brief, we define

$$W_{k \rightarrow l} = \sum_i \delta(s_i = k) \exp(-\beta[E_l - E_k]/2),$$

where k, l index chemical species, i indexes sites of the lattice, s_i is the species at site i , and E_k is the value of the Hamiltonian with site i replaced by species k . It may be shown that

$$\frac{W_{j \rightarrow i}}{W_{i \rightarrow j}} = \exp \beta(\mu_j - \mu_i)$$

when the lattice is in equilibrium with a bath of particles at specified chemical potentials μ_i . The calculated chemical potential differences are integrated over composition space to estimate the

overall free energy per molecule $g = G/N$. Finally, μ_C can be obtained from g and the chemical potential differences using the identities derived in Supplemental Note 1.

To determine the phase diagram implied by the simulations, tie lines are estimated from individual samples of each simulated composition. Local composition images are computed by convolving the binary matrix of site occupancies for each component with a disc with a diameter of 25 lattice sites. For a phase-separated composition, 2d histograms of these local compositions with respect to Cholesterol and DPPC mole fractions show two peaks, with compositions equal to those of the ends of the tie-line that the overall simulated composition falls on. Compositions that do not phase separate are characterized by a single peak in their local composition histograms.

Experimental observations of μ_C were fit to this model following a simplified scheme. First, simulations of binary cholesterol-lipid mixtures with a wide range of interaction parameters were carried out. These were fit simultaneously to the binary cholesterol-DOPC and cholesterol-DPPC mixture data points to produce initial estimates of J_{C-S} , J_{C-U} and μ_0 . Then, the remaining parameter J_{S-U} was estimated by conducting a series of simulations using these initial estimates, at fixed cholesterol fraction over a range of DOPC/DPPC ratios, considering a range of J_{S-U} values, and choosing the value that most closely matched the data for those mixtures. Simulations of the full model over the entire composition space were conducted for a 3x3x3 grid of interaction parameter values centered on these initial estimates, with grid spacing of approximately 0.2 $k_B T$. The optimal parameter set was chosen as that with the best simultaneous weighted least-squares fit to all the μ_C data, as described above for the mean field models.

3.4 Results

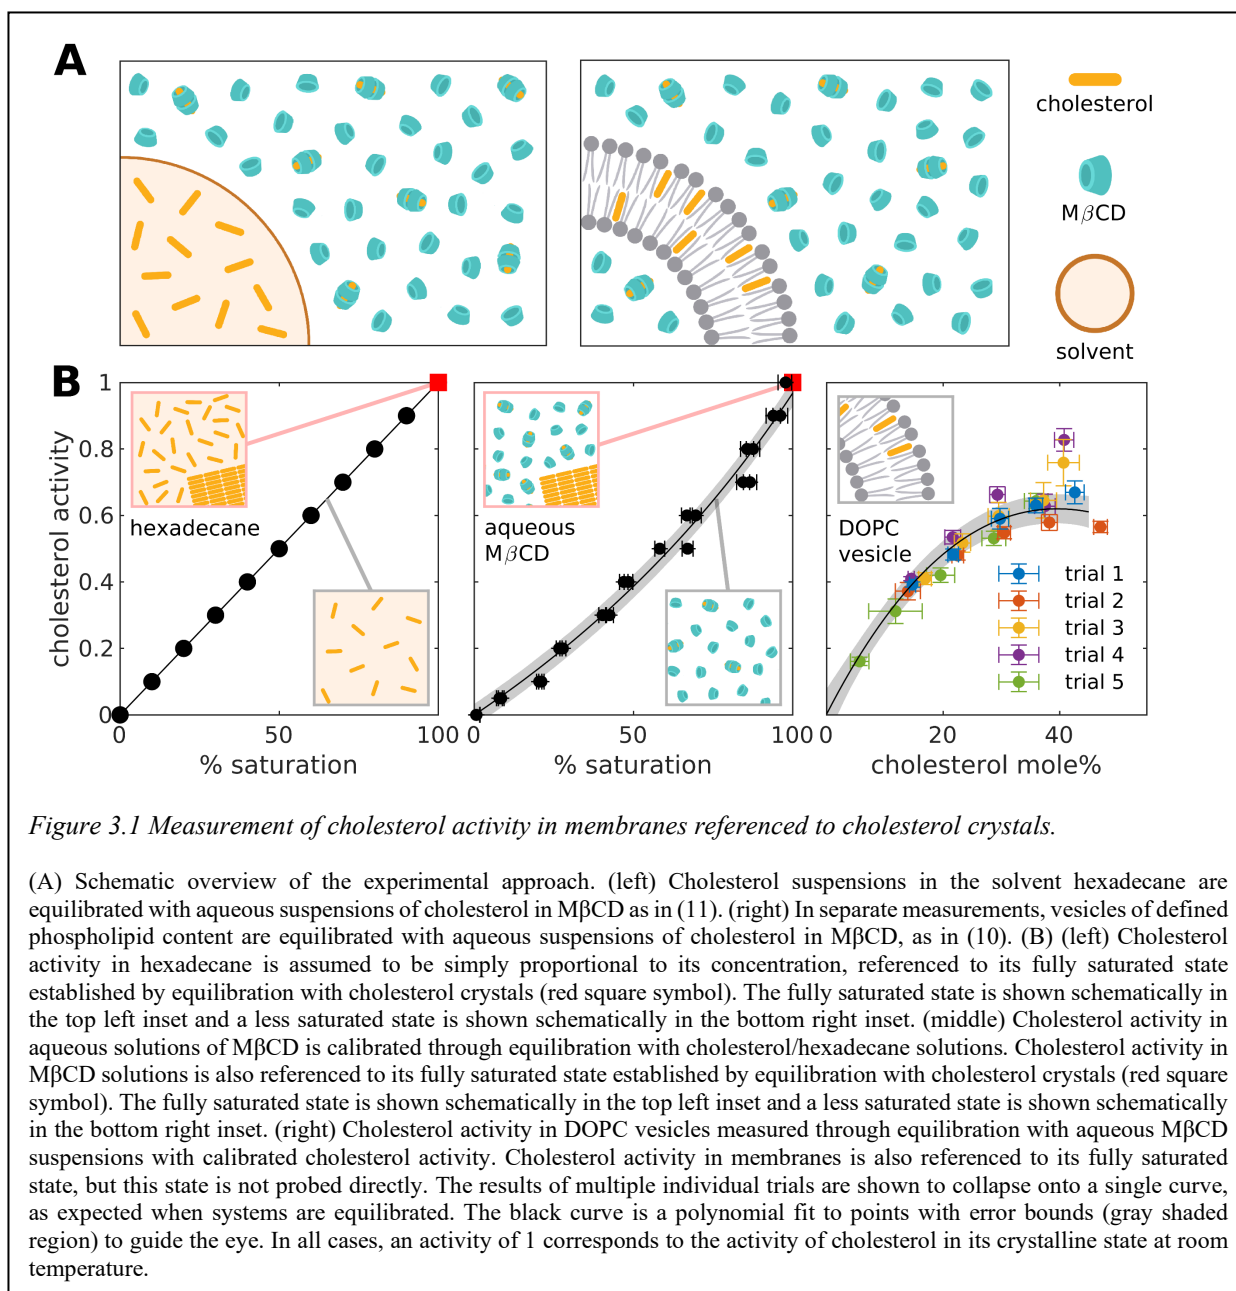
3.4.1 Measurements of μ_C in binary mixtures with DOPC, POPC, and DPPC.

Cholesterol activity in membranes is measured following a procedure that takes advantage of the property that the chemical activity A of a compound is constant across subsystems once those subsystems have come to thermodynamic equilibrium. In this context, if the cholesterol activity can be reliably measured in one state, it can be inferred in the second state when states are equilibrated. Chemical activity is defined relative to a reference state, and by convention we choose this reference state by setting $A = 1$ in cholesterol crystals at room temperature. Fully saturated solution are prepared through equilibration with cholesterol crystals, so that $A = 1$ in saturated solutions as well.

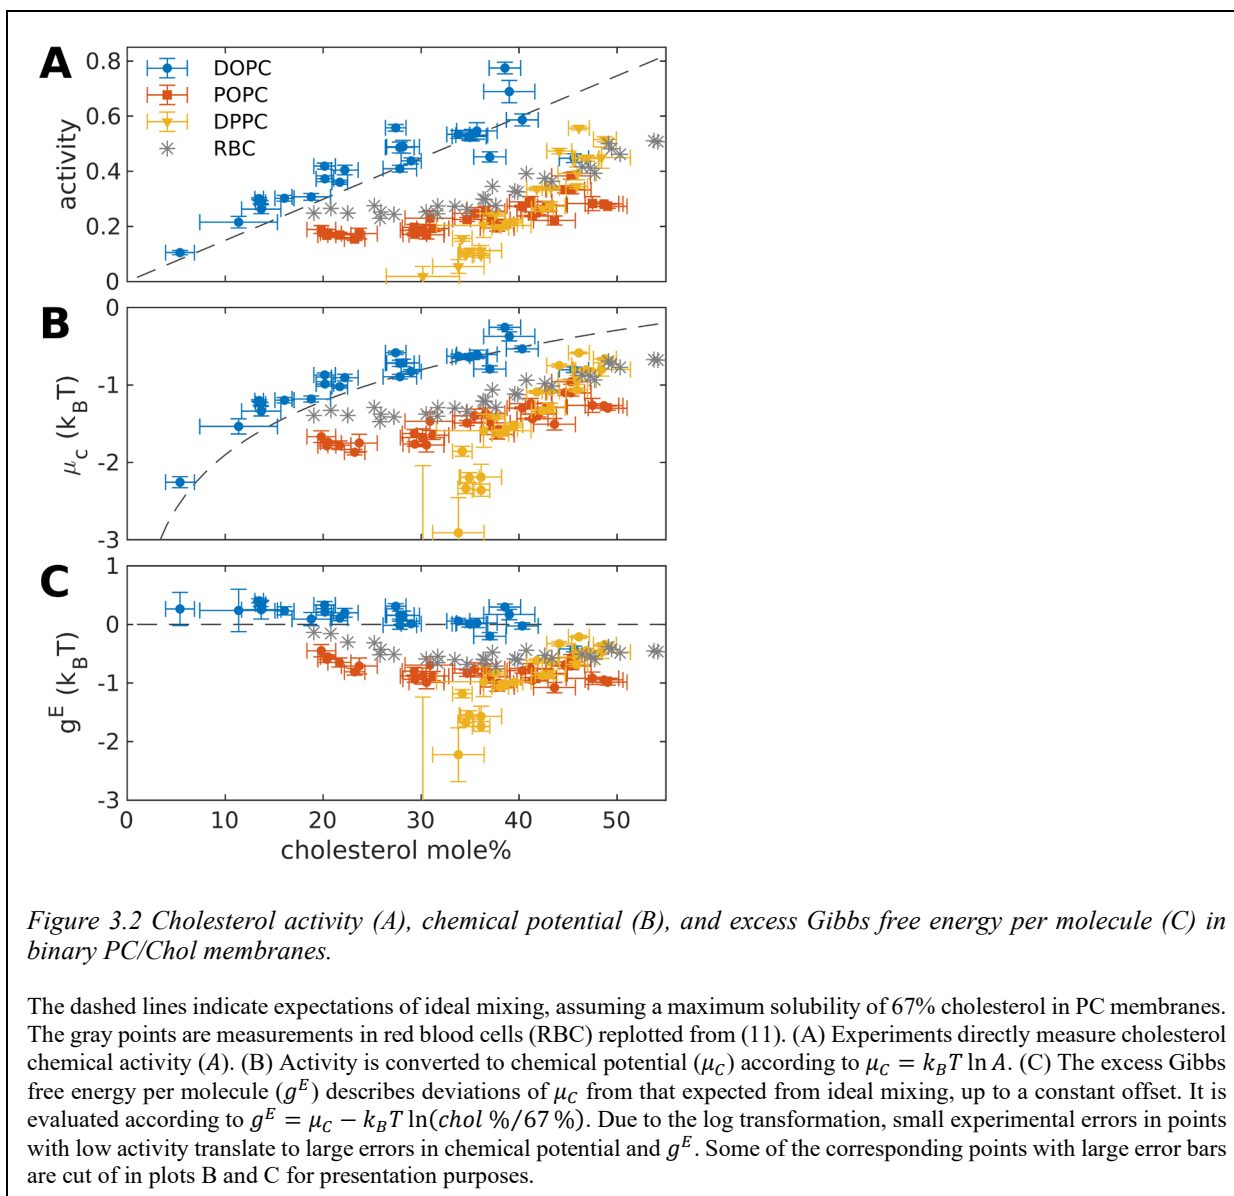
Briefly, we first equilibrated aqueous solutions of cholesterol and 5 mg/ml M β CD with cholesterol dispersed in an organic solvent. In the organic phase, cholesterol activity is proportional to its concentration, therefore measuring cholesterol concentration within the coexisting aqueous M β CD solution calibrates the relationship between cholesterol concentration and activity in this phase. In separate measurements, we equilibrated membranes with aqueous solutions containing M β CD and cholesterol at room temperature, then measured the cholesterol mole % in membranes and the % saturation of cholesterol within the M β CD solutions. Applying the calibration curve for aqueous M β CD solutions, we can then infer the cholesterol activity within the membrane. These steps are presented schematically in Figure 3.1, and detailed experimental procedures are described in Materials and Methods.

Figure 3.2A shows how cholesterol chemical activity varies with cholesterol mole % in vesicles of different phospholipids, highlighting the significant impact of phospholipid chains. For reference, we also plot the expectation from ideal mixing as a dashed line. This dashed line contains a single arbitrary constant, corresponding to the slope of the activity vs concentration or, equivalently, to a constant offset in the chemical potential. We choose this slope to agree with past findings that the maximum solubility of cholesterol in several membrane systems is near 67 mole % (247–250). This solubility limit implies that cholesterol membranes with 67 mole % cholesterol can be equilibrated with cholesterol crystals, and therefore have an activity of $A = 1$. When the chemical activity is converted to chemical potential according to $\mu_C = k_B T \ln A$ (Figure 3.2B), the dashed line corresponding to ideal mixing becomes curved. Figure 3.2C presents the excess Gibbs free energy per molecule, g^E , which is calculated according to $g^E = \mu_C - k_B T \ln(\text{chol \%}/67 \%)$. g^E isolates the non-ideal behavior of μ_C referencing to a maximum solubility of 67% cholesterol in all membranes. Note that while there is a single arbitrary offset to these curves, their variation both with cholesterol fraction, and with the phospholipid forming the rest of the membrane implies non-ideal mixing.

The experimental trends for DOPC shown in Figure 3.2 closely resemble expectations of ideal mixing over a broad range of cholesterol mole %, indicating that the magnitude of interactions between cholesterol molecules is close to the magnitude of interactions between cholesterol and DOPC lipids within membranes. In fact, the trends suggest that DOPC-cholesterol interactions are slightly less favorable than cholesterol-cholesterol interactions since measurements systematically reside above the ideal mixing line in Figure 3.2A,B, and because of



the downward trend in g^E with increasing cholesterol mole % in Figure 3.2C. This means that it becomes relatively easier to add additional cholesterol molecules to DOPC membranes as cholesterol concentration is increased. An alternate possibility is that the maximum solubility of cholesterol in DOPC membranes is somewhat less than 67%, and that the apparent curvature in Figure 3.2A and downward trend in Figure 3.2C arise from subtle systematic errors. A more careful investigation would be needed to distinguish these possibilities.



Experimental trends observed for cholesterol within DPPC membranes are far from expectations of ideal mixing, with values falling systematically below the ideal mixing lines. This indicates that the magnitude of attractive interactions between cholesterol and DPPC far exceed those between cholesterol molecules. The steep positive slope of g^E vs cholesterol mole% in DPPC membranes indicates that the favorable cholesterol-membrane interactions are substantially attenuated at higher cholesterol mole %. This effect could be due to unfavorable cholesterol—cholesterol interactions screening the more favorable DPPC—cholesterol interactions that dominate at lower chol mole %, to cholesterol molecules saturating a limited pool of DPPC molecules, to changes in membrane structure associated with higher cholesterol content, or

combinations of these factors. DPPC/Cholesterol mixtures phase separate into liquid-ordered and solid phases below 30% cholesterol at room temperature (13), and our measurements in Figure 3.2 are consistent with these past observations. Cholesterol activity remains low even at elevated cholesterol concentrations, consistent with cholesterol's exchange between two equilibrated states, one of which is a gel phase with very low cholesterol concentration and therefore very low cholesterol activity.

Figure 3.2 also shows measurements for POPC membranes at room temperature and past measurements for red blood cell (RBC) membranes conducted at 37°C, replotted from (223). These both exhibit intermediate and more complex behaviors than observed for cholesterol within DOPC or DPPC membranes. Both cases are not ideally mixed, exhibiting more attractive interactions between cholesterol and the average membrane component than between cholesterol molecules themselves. For POPC, cholesterol activity is only weakly dependent on cholesterol mole % in vesicles, meaning that the difference in interactions nearly cancels the difference in mixing entropy over a range of cholesterol concentrations. When these contributions exactly balance, phase separation occurs, suggesting that POPC/Cholesterol membranes exhibit properties close to that of a phase separated system, in good agreement with past studies (145, 251, 252). Cholesterol activity is independent of cholesterol concentration for RBC membranes with 15–30 mole % cholesterol, consistent with phase coexistence in that region. We note that measurements in RBCs were conducted at 37°C, with a reference state of cholesterol crystals at this same temperature. Because of this, it is appropriate to compare trends present in RBC measurements to those conducted in purified vesicles at room temperature but not absolute values.

Various past studies have measured cholesterol partitioning between membranes of different compositions or between membranes and M β CD, and their results imply chemical potentials with respect to various reference states (33–35, 235, 238–240). Much of this past work is not directly comparable, because it assumes simple Nernst partitioning with constant partition coefficient, equivalent to assuming g^E is independent of cholesterol concentration. In addition, most of the existing literature reports on different lipids or lipid mixtures than those presented here. However, cholesterol-dependent measurements equivalent to chemical potentials have been determined in POPC-cholesterol mixtures at 25°C using calorimetric methods, which simultaneously measure partitioning and the enthalpy of transfer between M β CD and LUVs (35).

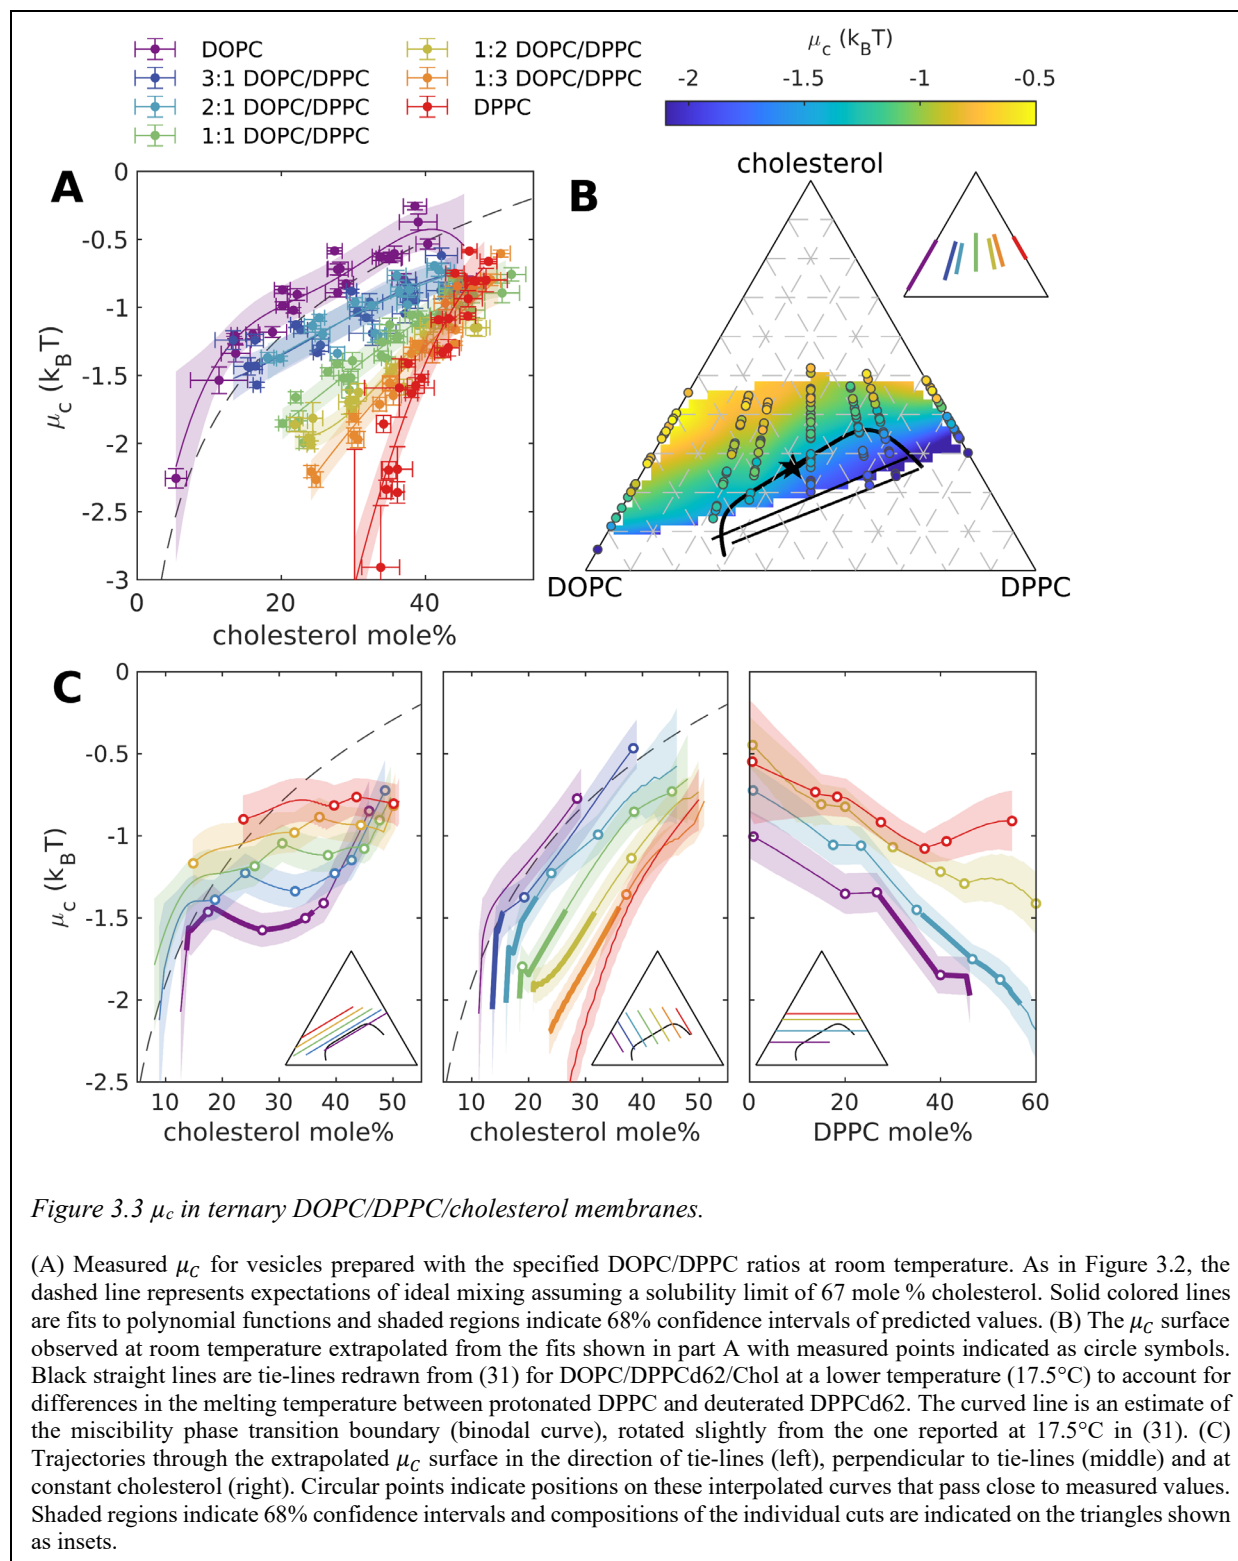
We find good quantitative agreement with these past studies, as demonstrated in Supplemental Figure 3.13. Supplemental figures and tables are included in Section 3.7.

3.4.2 Measurements of μ_C in ternary mixtures of DOPC, DPPC, and cholesterol.

We have also measured μ_C in vesicles containing both DOPC and DPPC, and results are summarized in Figure 3.3. LUVs were prepared with phospholipids at the specified molar ratio and incubation in aqueous M β CD/cholesterol mixtures produced vesicles that retained the same molar ratio of DOPC to DPPC but with varying cholesterol content. To smooth noise, results for specific DOPC/DPPC ratios were fit to polynomials to capture both the trend and the confidence interval of the measurement (solid lines and shaded regions in Figure 3.3A respectively). These fits are used to interpolate the chemical potential surface shown in Figure 3.3B.

Mixtures of DOPC, DPPC, and cholesterol undergo liquid-liquid and solid-liquid phase separation, and Figure 3.3B also includes tie-lines measured for a closely related lipid mixture by deuterium NMR (127) and an estimated phase boundary. The measured μ_C surface includes regions both inside and outside of the estimated miscibility gap but does not extend to compositions where gel or solid phases are reported at lower cholesterol and higher DPPC concentrations.

Figure 3.3C shows linear trajectories through the measured μ_C surface. These include trajectories that run in the direction of estimated tie-lines within the miscibility gap. We note that the estimated phase boundary and therefore tie-line direction drawn here is rotated slightly from those determined in past work using ^2H NMR. This was done to enforce that μ_C remains constant along the tie-line within the coexistence region, a requirement for chemical equilibrium in phase-separated systems (253). Minor differences in the phase boundary and tie-lines are expected here both because a protonated DPPC lipid is used and because a small mole % of NBD-PE is included (0.1%) (254). Trajectories that run parallel to tie-lines but pass outside of the miscibility gap also retain a shallow slope over the range of compositions interrogated, similar to those observed for cholesterol within POPC membranes in Figure 3.2. Trajectories through the μ_C surface that run perpendicular to tie-lines increase with increasing cholesterol concentration. Trajectories at constant cholesterol concentration vary by several $k_B T$ as phospholipid ratios are changed, highlighting that differential interactions between cholesterol and different phospholipids play important thermodynamic roles in this system. In addition, some constant cholesterol



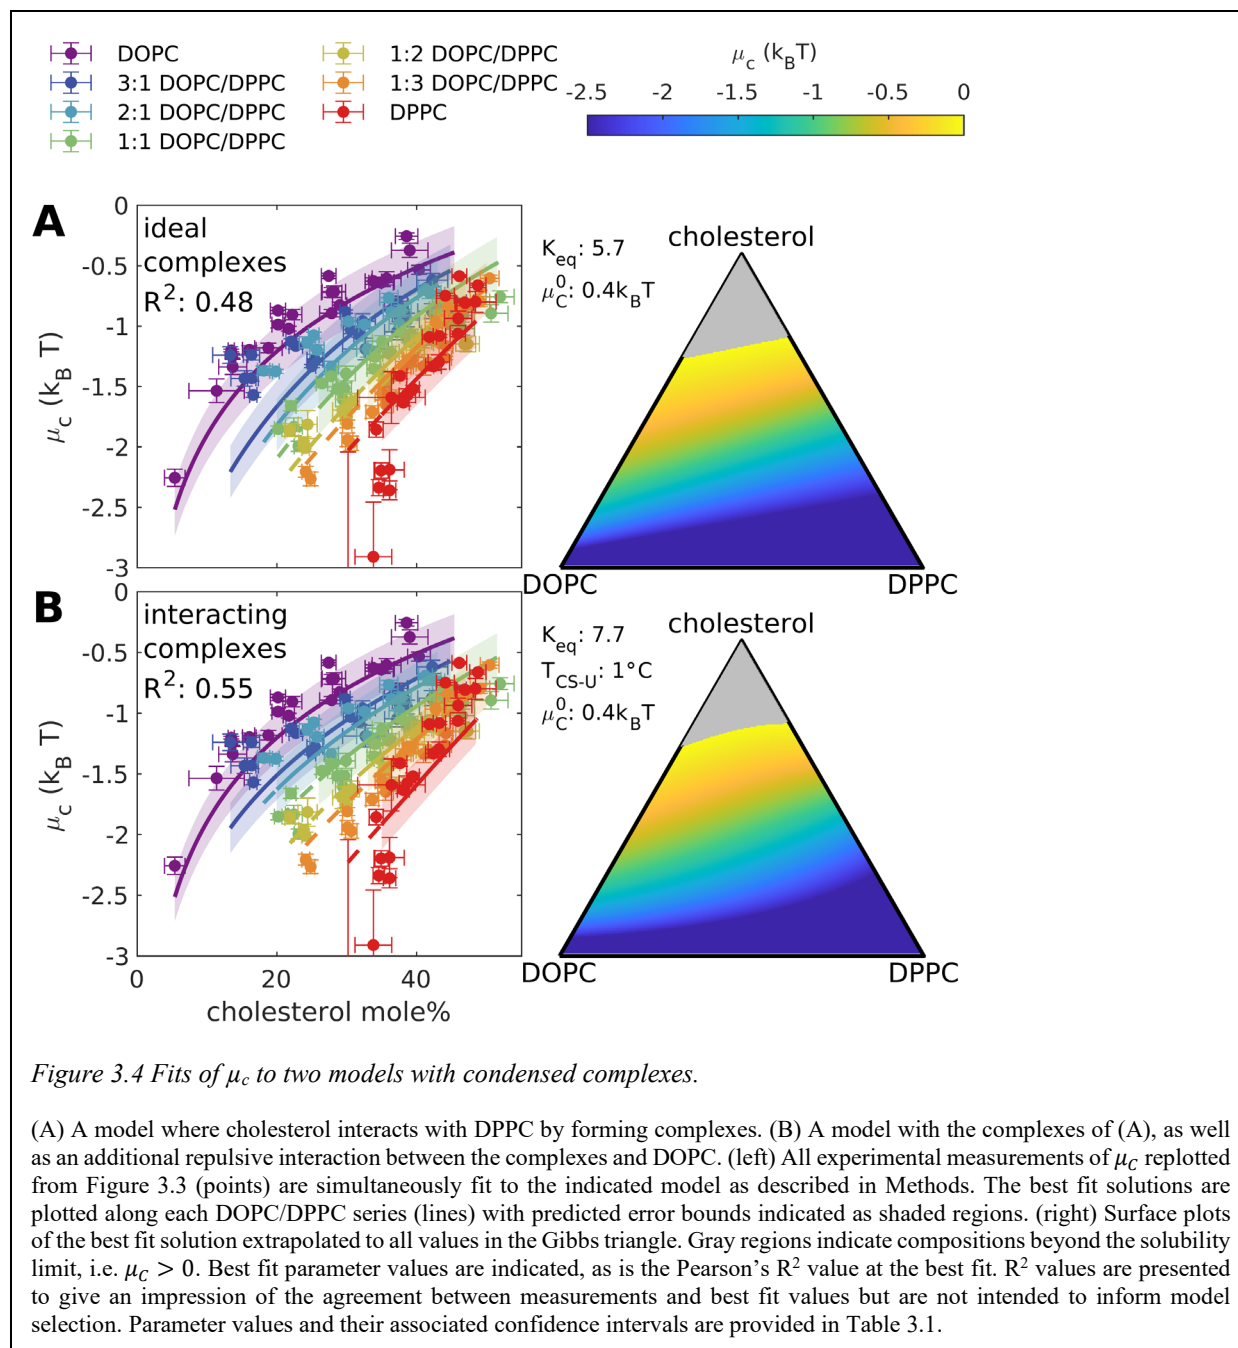
trajectories appear nonlinear, potentially indicating nontrivial DOPC—DPPC interactions. In the absence of these interactions, a linear trend would be expected as the prevalence of cholesterol—

DOPC interactions is linearly exchanged for cholesterol—DPPC interactions. Plots resembling Figure 3.3 but reporting results as activity or the excess Gibbs free energy per molecule are provided in Supplemental Figure 3.14.

μ_C is the derivative of the Gibbs free energy with respect to the number of cholesterol molecules and is therefore specified by thermodynamic models. The figures that follow use measurements of μ_C shown in Figure 3.3 to constrain thermodynamic models of DOPC, DPPC, cholesterol membranes. Our goal is to better understand the magnitude and type of interactions that give rise to the phase behavior of these systems, as well as their implications for other physical properties that are not measured here, rather than to select the most appropriate model. Models are fit to all measurements simultaneously, excluding points corresponding to compositions within the phase separation region, as described in Materials and Methods.

3.4.3 Measurements of μ_C limit the binding affinity of cholesterol-phospholipid complexes in a condensed complex model.

One way to explain non-ideal mixing is in a model where other lipids compete for cholesterol by forming stoichiometric complexes, thereby lowering the concentration of unbound cholesterol. Previous work proposed a condensed complex model to describe aspects of phospholipid-cholesterol phase diagrams in monolayer and bilayer membranes (255–258). In this model, cholesterol and saturated phospholipids interact by assembling into complexes of fixed stoichiometry with an affinity characterized by an equilibrium constant K_{eq} , as described in Materials and Methods and derived in Supplemental Note 4. Figure 3.4A presents a model that allows for the formation of a condensed complex made up of cholesterol and DPPC with a fixed stoichiometry of 1:1, but does not contain additional interactions between components or between components and complexes, as done in past work (257, 258). The model used in Figure 3.4A contains 2 parameters: a constant offset μ_C^0 which defines the solubility limit (mole % cholesterol at 100% saturation), and the equilibrium constant of complex formation (K_{eq}). This model captures many aspects of experimental μ_C measurements. For example, the observed shifts towards larger cholesterol mole % for smaller DOPC/DPPC ratios. Also, the best fit μ_C^0 corresponds to a solubility limit of 69% chol in DOPC and 75% in DPPC, in reasonable agreement with past reports of 67% in both systems (247). The best fit value for K_{eq} is 5.7, corresponding to a binding affinity of 1.7



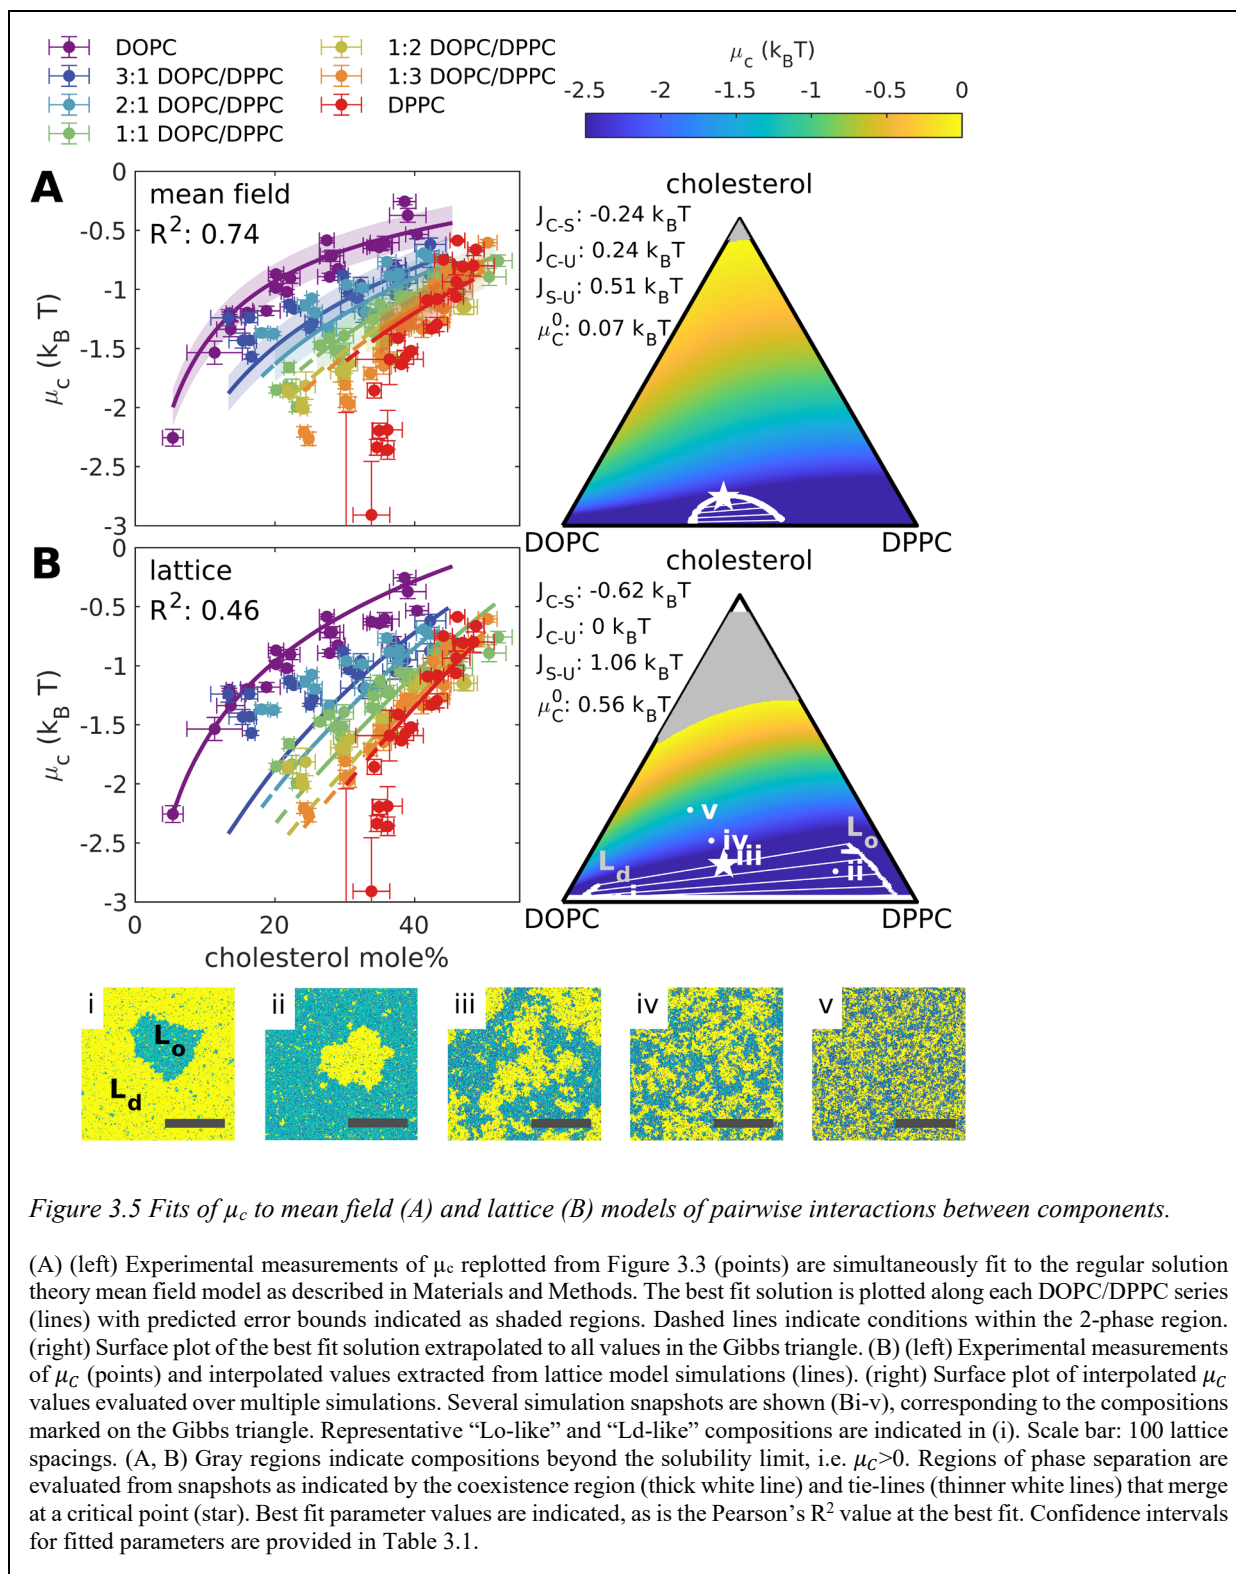
$k_B T$, which is much lower than assumed in some previous studies (255–258). This weak association energy means that the fraction of cholesterol not bound in complexes remains substantial across all lipid compositions, as shown in Supplemental Figure 3.15A. The model of Figure 3.4A does not support phase separation by construction, since the Gibbs free energy does not include interactions between components beyond those required for complex formation.

Figure 3.4B fits the measured μ_C to a model that includes both complex formation between DPPC and cholesterol and repulsive interactions between complexes and DOPC, and can in principle produce phase separation. The model used in Figure 3.4B contains 3 parameters, μ_C^0 and K_{eq} as in Figure 3.4A, and an interaction energy represented as a critical temperature for phase separation between complexes and unsaturated lipid (T_{CS-U}). As expected, including the additional fit parameter allows for a somewhat better fit to experimental measurements, as indicated by the slightly higher R^2 . More importantly, the fitted equilibrium constant K_{eq} remains small ($K_{eq} = 7.7$), again indicating weak binding of this complex and a meaningful fraction of uncomplexed cholesterol across all compositions (Supplemental Figure 3.15B). The best fit parameters for this model are not consistent with a phase separating mixture at room temperature. Instead, the best fit value for T_{CS-U} indicates that temperature would need to be lowered below 1°C for phase separation to occur anywhere in this mixture.

We also considered other elaborations of this class of model, including one where the interaction energy between the complex and DPPC was allowed to vary ($T_{CS-S} \neq 0$), and a second with a complex stoichiometry of 1 cholesterol to 2 DPPC lipids, as shown in Supplemental Figure 3.16. In both cases results were either not improved compared to the models of Figure 3.4 or were considered unphysical.

3.4.4 Phase separation and measurements of μ_c are described by models capturing pairwise interactions between components.

We next investigated whether experimental observations could be described by models that include pairwise interactions between all three components without the explicit formation of complexes (259). Two approaches are shown in Figure 3.5 and described in Materials and Methods, with detailed explanations in Supplemental Notes 2, 3 and 5. In the first, we employ a mean field regular solution theory model of the ternary system and the best fit solution is shown in Figure 3.5A. There are 4 parameters in this model: μ_C^0 and interaction energies between all three pairs of components, between saturated and unsaturated (J_{S-U}), saturated and cholesterol (J_{C-S}) and unsaturated and cholesterol (J_{C-U}). This model accurately captures trends in measured μ_C , with a higher R^2 value than the model of Figure 3.4B, although part of this improvement may be due to the presence of an additional fit parameter. We find weak repulsive J_{C-U} and weak attractive J_{C-S} ($J_{C-U} = 0.24 k_B T$, $J_{C-S} = -0.24 k_B T$), reflecting the preferred partitioning of cholesterol



with DPPC over DOPC. Notably, the best fit J_{S-U} is substantially repulsive ($J_{S-U} = 0.51 k_B T$), consistent with interactions that favor phase separation of coexisting high-DPPC and low-DPPC

phases. Indeed, the best fit regular solution theory model predicts phase separation, albeit only at low cholesterol mole % and with a narrower two phase coexistence region. The phase separated region of the model and several tie lines are shown. The tie-lines are roughly parallel to those observed experimentally. This model does not incorporate interactions that would give rise to a solid phase, therefore the phase diagram does not capture regions of solid-liquid coexistence or the 3 phase region detected experimentally (127). In supplemental note 6, we compare to a past report that has fit an equivalent model to POPC/sphingomyelin/cholesterol interactions at 37°C using a complementary method (237).

The analytical models described in Figure 3.4 and Figure 3.5A are mean-field models that consider average effective concentrations. Membranes are expected to exhibit deviations from this mean-field approach, since the mean field approach ignores fluctuations, which are particularly important in two-dimensional (2D) systems, especially near critical points (196). To address this, we additionally conducted simulations of pairwise interactions between components on a 2D square lattice and extracted μ_C as described in Materials and Methods and Supplemental Note 5. Because these μ_C estimates are computationally expensive, we carried out an approximate fit as detailed in Materials and Methods and the resulting μ_C estimates are shown in Figure 3.5B. Although these parameters are likely not the global best fit parameters, they give a decent fit to the data. The R^2 value is somewhat lower than the mean field model of Figure 3.5A, but in the same range as those of the complex models of Figure 3.4. Several representative simulation snapshots are also shown.

The optimal interaction parameters were qualitatively similar to those observed in the mean field model, and J_{S-U} was again the dominant repulsive interaction. We found that the repulsive DOPC—cholesterol interaction of the mean field model is absent here, which leads to a slightly improved fit at the DOPC—cholesterol edge (purple points), where mixing is nearly ideal as seen in Figure 3.2. In addition, the other two interactions J_{C-S} and J_{S-U} are more than twice as strong as they are in the mean field model, which is required in this model to achieve phase separation at similar temperatures (196). Similar to the corresponding mean field model from Figure 3.5A, this model supports a two-phase coexistence region at low cholesterol mole % that exhibits tie-lines roughly parallel to those observed experimentally. However, in this case, the coexistence region is much broader. This is expected because the shape of the phase boundary is flatter near the critical point in 2D compared to mean-field predictions (196).

3.5 Discussion

Cholesterol is vital for a broad range of cellular functions, and extensive past work demonstrates its multifaceted roles in tuning the physical properties of membranes. In this study, we have measured the chemical potential of cholesterol μ_C in purified binary and ternary membranes by equilibrating them with an aqueous M β CD phase. These measurements of μ_C provide a window into the thermodynamics of these systems, as interactions with the other lipids in the bilayer are reflected in μ_C .

Several influential past studies have studied the partitioning of cholesterol between membranes of different compositions or between membranes and cyclodextrin (33–35, 235, 240, 260). These studies have observed important trends, one of which is that cholesterol partitions more strongly into membranes with greater lipid saturation, indicating stronger interactions between cholesterol and saturated compared to unsaturated acyl chains. With a few exceptions (35, 236, 237, 251), these past studies focused on the dilute limit where cholesterol partitioning can be modeled using a constant partition coefficient. The present study extends this type of measurement to new lipid mixtures, focusing on ternary mixtures of DOPC, DPPC, and cholesterol, a system with well characterized and complex phase behavior and where the partition coefficient varies substantially with membrane composition. A second distinguishing feature of the current work is that measurements are reported in terms of a chemical potential referenced to the well-defined and reproducible standard state of crystalline cholesterol. Reporting results in this way simplifies the comparison of measurements made in different experimental contexts.

Chemical potential differences report on the direction that cholesterol currents flow between membranes when they are allowed to come to thermodynamic equilibrium. Chemical potential differences also directly inform the relative availability of cholesterol to bind to proteins embedded within different membranes. For example, a $1.6 k_B T$ increase in μ_C corresponds to cholesterol being 5 times more available to bind to a cholesterol binding site. In our measurements, even different membranes with the same cholesterol mole % can have μ_C shifts of this magnitude.

In ideal mixtures, the chemical activity of a component is simply proportional to its concentration, or equivalently the excess Gibbs free energy g^E of a component is constant. The measurements presented in Figure 3.2 and Figure 3.3 provide many instances that violate this simple picture. Non-ideality is apparent when taking constant cholesterol slices through the DOPC/DPPC/cholesterol μ_C surface, where it is observed that μ_C can vary by several $k_B T$, and in

some regions depends nonlinearly on the DOPC/DPPC ratio. This is also apparent when examining that μ_c remains constant over a range of cholesterol concentrations when traversing a tie-line in the miscibility gap, as required for two-phase coexistence. These nearly constant contours persists well beyond the limits of the phase separated region, emphasizing that the same interactions that give rise to phase separation also impact the physical properties of membranes outside of phase separation.

3.5.1 Experimental trends in μ_c can constrain simple models

To guide intuition about the important interactions in DOPC/DPPC/Chol membranes, we applied simple phenomenological models. We note that none of these models are constructed to replicate interactions that give rise to regions of solid-liquid coexistence, as is experimentally observed in this system at high DPPC and low cholesterol compositions, therefore we anticipate disagreement between models and experiment in this regime. However, we find that fitting the data to various models provides useful constraints on model parameters.

We first modified a model introduced by Radhakrishnan and McConnell (255–258), in which interactions between cholesterol and saturated lipids are described by a reversible but energetically favored binding reaction to form a stoichiometric condensed cholesterol-phospholipid complex. Measurements of μ_c are reasonably well explained by the model when the association constant for complexes is weak compared to past estimates (257, 258), implying a substantial fraction of the cholesterol remains unbound in this model (Supplemental Figure 3.15). In addition, we find that the fitted interaction strengths are not consistent with room-temperature phase separation of these mixtures as is observed experimentally (6, 127).

We have also presented both mean field and lattice models where deviations from ideal mixing come from local interactions between lipid molecules and their nearest neighbors. Here, the model parameters that best explain the μ_c measurements reproduce the favorable interactions between cholesterol and DPPC that are expected, and also predict substantial repulsive interactions between DPPC and DOPC. In addition, the model reproduces the qualitative shape of the experimental phase coexistence region for this lipid mixture, although some features of the coexistence region are poorly matched to experiment, such as the width of the mean-field coexistence region and the cholesterol content at the critical point in both models. While these models lack an explicit complexation reaction, fitted parameters imply that cholesterol and DPPC

will often be found adjacent to each other. In this sense, while complex formation brings a different microscopic picture to mind, it leads to similar thermodynamic predictions in the limit where complexes are relatively weak.

3.5.2 Comparison to past results in model membranes and cells

A comparison of past results in red blood cell membranes (223) and current results in model membranes is shown in Figure 3.2. The nearly common reference state between this study and ours enables direct comparisons across measurements. Comparing results, it is apparent that RBC membranes more closely resemble POPC and DPPC membranes than they do DOPC membranes. The non-ideality of RBC membranes is notable, with nearly constant cholesterol activity over a wide range of cholesterol concentration (15 – 30 mole %). Past work also measures relatively low μ_C in two mammalian cell culture lines to be between -1.4 and $-2.2 k_B T$ (223). These values correspond to <15 mole % cholesterol in DOPC membranes or less than 40 mole % cholesterol in DPPC membranes. Cells in culture typically contain 30-40 mole % cholesterol and a large fraction of unsaturated and polyunsaturated lipids (161), therefore it is striking that the availability of cholesterol in these membranes is much lower than is measured in a model membrane with similar cholesterol concentration. This comparison stresses the importance of matching chemical potential rather than concentration in reconstituted systems. One structural difference between the model membranes studied here and cell membranes measured previously is the presence of phospholipid and cholesterol asymmetry across leaflets. Recent reports suggest that cholesterol plays a role in supporting this asymmetry (261, 262), and it is possible that the surprisingly low μ_C detected for cells is related to this phenomenon.

There is a body of past work that reports cholesterol chemical activity inferred in model membranes and cells from kinetic measurements of cholesterol extraction or oxidation with M β CD or cholesterol oxidase respectively (239, 248, 263–266). These past studies argue that initial rates of extraction or oxidation are proportional to the chemical activity in the equilibrium state. While this measurement almost certainly correlates with the activity of cholesterol, it could also be impacted by kinetic properties leading it to differ from the true equilibrium activity, especially where membrane structure is changing (266). For example, tighter head-group packing, thicker hydrophobic regions, or a reduction in the number or lifetime of defects could impede access of M β CD or enzymes to cholesterol and slow entry of cholesterol into binding sites. Another, related

body of work uses the binding of soluble molecules to membranes to infer cholesterol activity (238, 263, 267–270). Here too we agree that the localization of these peptides or proteins to membranes via cholesterol binding will depend strongly on the activity of cholesterol, but it will also depend on any other interactions that proteins or peptides make in the membrane bound state, and on the energetics of any membrane deformations that are induced by binding. Again, it is reasonable to expect that these other interactions might depend on cholesterol composition in complicated ways. These caveats do not diminish the value of past conclusions but emphasize that these studies do not directly report on chemical activity in the thermodynamic sense even though this same language is used.

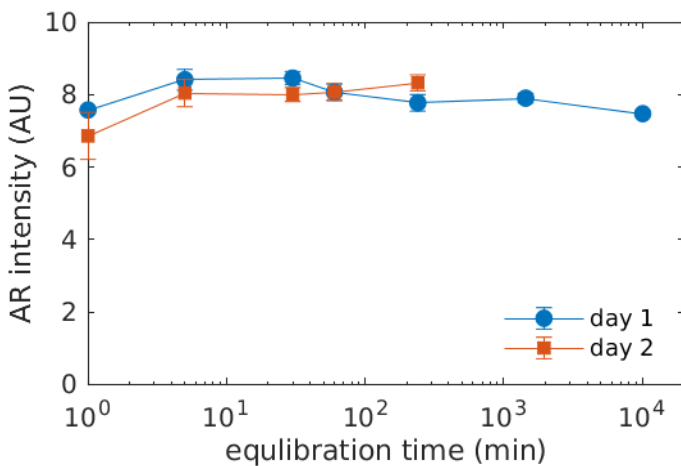
3.6 Conclusions

Here we demonstrate that μ_C within mixtures of cholesterol with one or two phospholipids is far from ideal, meaning that it depends on the full composition of the system rather than just the concentration of cholesterol. The composition dependence of μ_C is directly connected to the physical properties of cholesterol's interactions in each membrane. Since our measurement of μ_C is calibrated, we are able to compare our findings to past studies in model and cell membranes and explore whether our results are consistent with complementary models of interactions between cholesterol and saturated lipids.

While biological membranes are considerably more complex than these model systems, these measurements motivate new questions relevant to the functional roles of cholesterol in cells. μ_C is a direct determinant of cholesterol-protein binding interactions and the equilibrium partitioning of cholesterol between different cellular membranes. Several groups have begun to explore the functional consequences of cholesterol availability in cells (223, 263, 270). Our current work emphasizes that μ_C is strongly composition dependent and closely related to phase behavior, suggesting that even minor components could have an outsized impact on μ_C especially if they induce large changes in the phase behavior. Isolated cell plasma membranes phase separate at biologically tuned transition temperatures that are sensitive to acute treatments with small molecules (21, 158, 182). Connecting this to the present work suggests that these treatments might functionally act through their impact on μ_C . This is especially exciting in the context of recent reports of cholesterol binding sites within membrane protein structures. We are intrigued by the

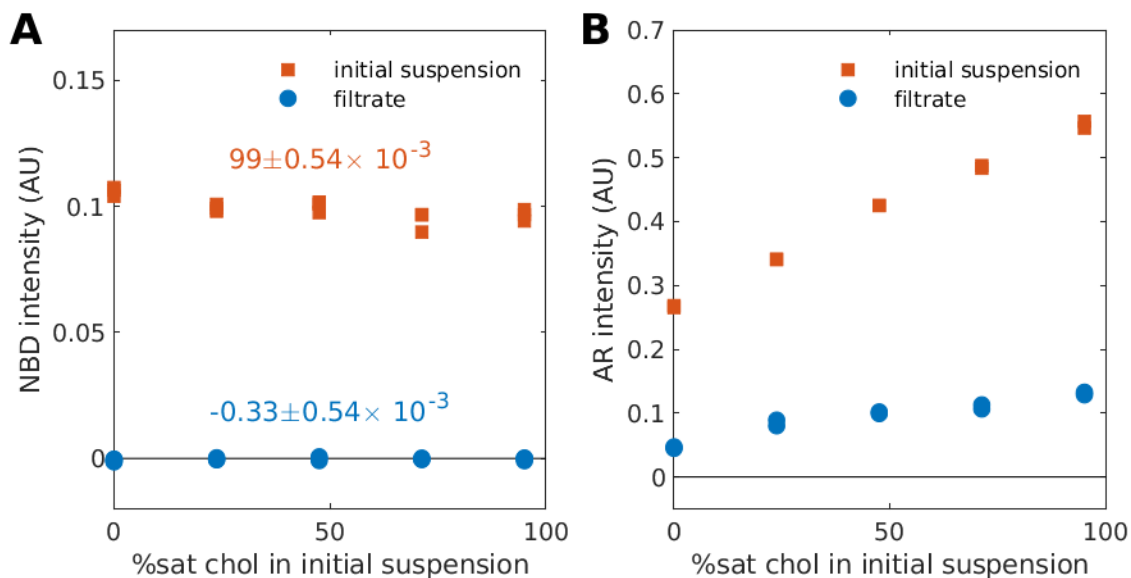
possibility that future measurement of μ_c could connect phase behavior to functional roles of cholesterol in cells.

3.7 Supplemental Figures and Tables



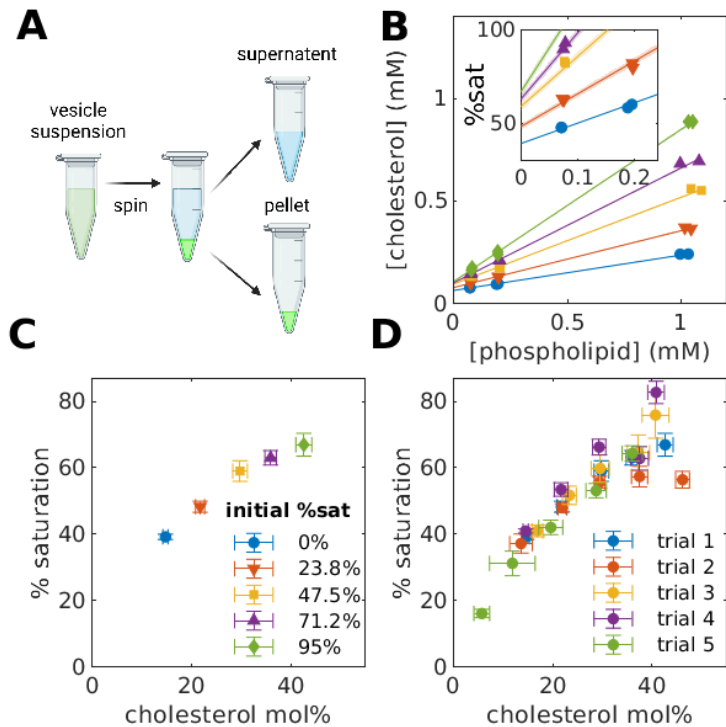
Supplemental Figure 3.6 Saturated solutions of cholesterol in 5 mg/ml M β CD equilibrate within 5 min and remain stable over extended times.

Cholesterol crystals were prepared as described in the main text and filtered at the specified time-points after sonication. Raw background corrected Amplex red (AR) signals were for filtered samples following the same protocol on different days. Immediately after sonication (1 min), AR intensities are systemically lower, but reach their maximal value by 5 min incubation. On day 1, samples reported for 24 h and 1 week were prepared using different M β CD solutions.



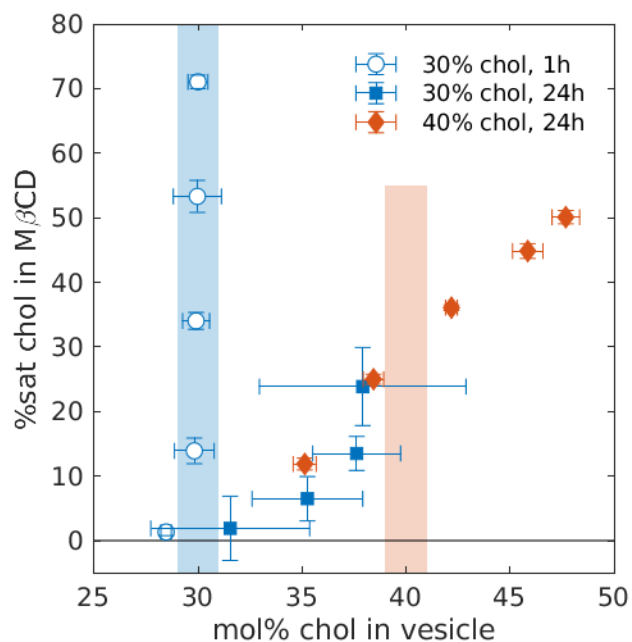
Supplemental Figure 3.7 NBD-PE is not detected in the M β CD containing aqueous phase.

(A) Background subtracted NBD intensities detected in initial vesicle suspensions and in filtrates passed through an Amicon Ultra-Centrifugal Filter Unit with a 30 kDa cutoff. Vesicles were diluted in aqueous M β CD solutions with the specified % saturation of cholesterol prior to filtration. The average and SE value for filtrate over all conditions is shown as blue text and is zero within error. (B) Background subtracted AR intensities for the same samples as in A. Cholesterol from vesicles partition into the M β CD containing aqueous phase, as indicated by the finite AR intensity found even when vesicles are diluted in M β CD without added cholesterol (0% sat). This example used vesicles that initially contained 1:2 DOPC/DPPC + 30% cholesterol.



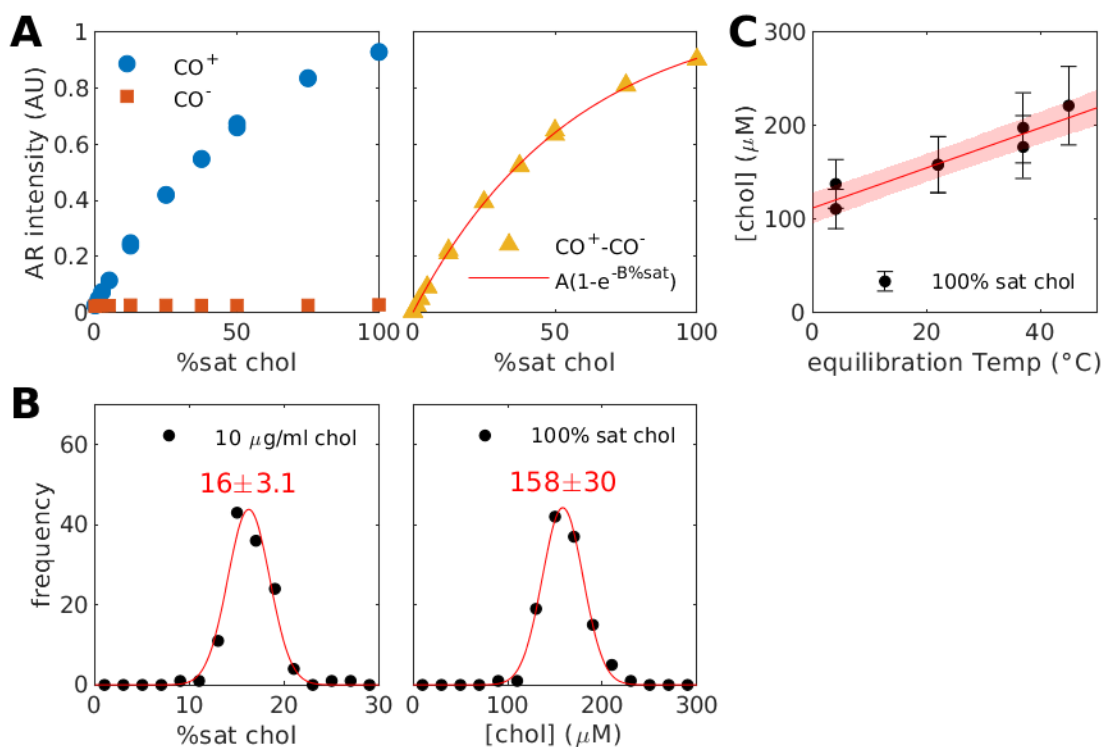
Supplemental Figure 3.8 Experimental determination cholesterol content of vesicles equilibrated with M β CD solutions.

A) Schematic representation of the measurement scheme. Vesicles were suspended in aqueous M β CD solutions with varying % saturation of cholesterol. After an equilibration period at room temperature, suspensions were spun to produce samples enriched or depleted in vesicles as described in Methods. B) Cholesterol and phospholipid concentrations were measured for supernatant, pellet, and initial suspensions as described in Methods. This example shows values for vesicles that initially contained 70% DOPC and 30% cholesterol suspended in M β CD solutions with the % saturation levels shown in the legend of part C. Points are fit to lines to determine the y intercept and the slope. The y intercept reports on the [chol] or % sat of the M β CD solution without lipid. The slope reports on the cholesterol to phospholipid ratio in vesicles. C) Slopes and intercepts from B plotted as % saturation of cholesterol in M β CD vs. cholesterol mole % in vesicles. D) The results of multiple trials with different DOPC/Chol vesicle preparations collapse into a single curve, as expected when samples are fully equilibrated. Measurements from parts B and C are redrawn in the series labeled trial 1.



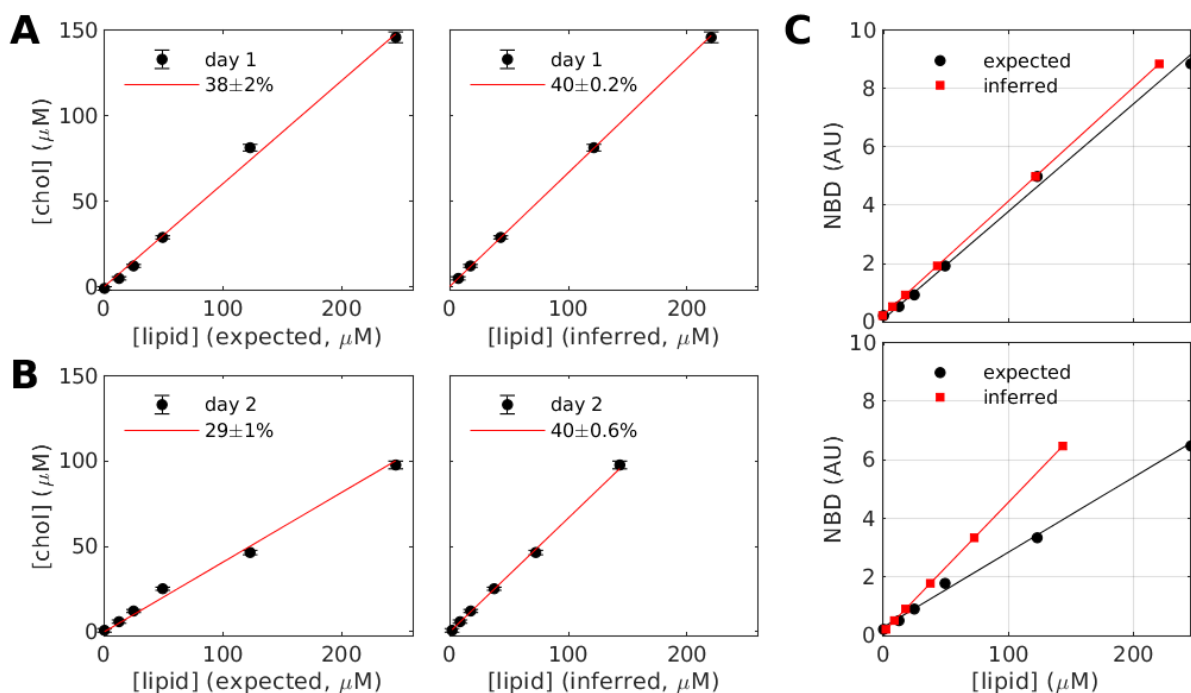
Supplemental Figure 3.9 Equilibration of systems of DPPC and cholesterol.

LUVs initially prepared with 70% DPPC and 30% cholesterol (blue shaded area) did not change composition after incubation with MβCD solutions for 1h (open blue circles). The same vesicle preparation did vary in composition after incubation with MβCD solutions for 24h (filled blue squares). Large error bars in this specific example arise from errors in the reading of NBD levels and not related the sample or preparation. LUVs initially prepared with 60% DPPC and 40% cholesterol changed in composition after 24h incubation with MβCD solutions. Data points from vesicles with different starting cholesterol compositions appear to follow the same master curve, providing evidence that samples have equilibrated.



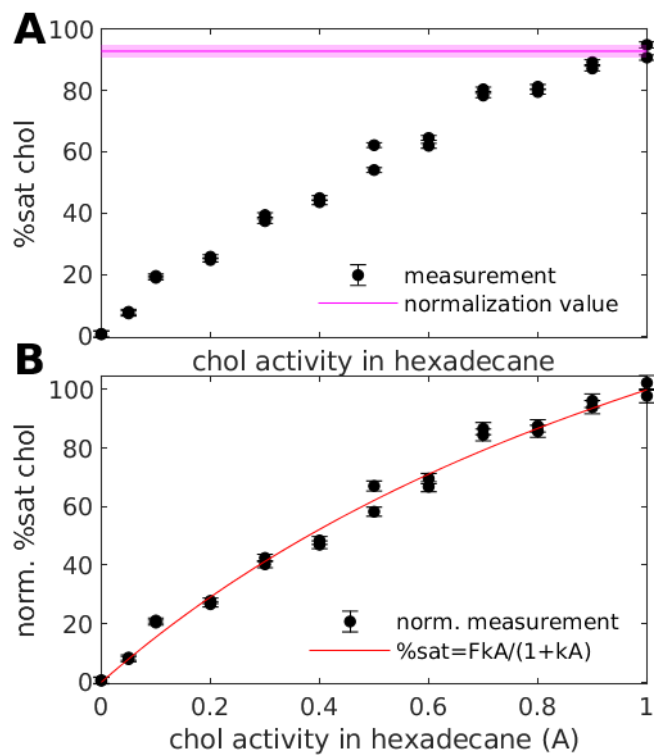
Supplemental Figure 3.10 Characterization of saturated cholesterol solutions in 5 mg/ml MβCD.

(A) Representative standard curve relating Amplex Red (AR) intensity to % saturation of cholesterol in 5 mg/ml MβCD solution at room temperature. (left) Samples are prepared by diluting fully (100%) saturated cholesterol with MβCD solutions prepared without cholesterol and incubated with a reaction buffer that either contains cholesterol oxidase (CO⁺) or not (CO⁻) then fluorescence intensity of the AR is measured. (right) The curves are subtracted and fit to the nonlinear form $AR = A(1 - \exp -B \cdot \%sat)$. In this example, the best fit parameter values are $A = 1.09 \pm 0.01$ and $B = 0.0176 \pm 0.0003$. (B, left) Histogram indicating measurements of the saturation of a 10 μg/ml cholesterol standard solution over 122 individual measurements at room temperature. (right) Histogram showing the same measurements presented as cholesterol concentration of the 100% saturated solution. Both histograms are fit to a Gaussian distribution to extract the mean and standard deviation shown. (C) The cholesterol concentration of saturated solutions in 5mg/ml MβCD varies with temperature. In these measurements, saturated solutions were prepared by equilibrating with cholesterol crystals overnight, then crystals were removed through filtration accomplished at the specified temperature.



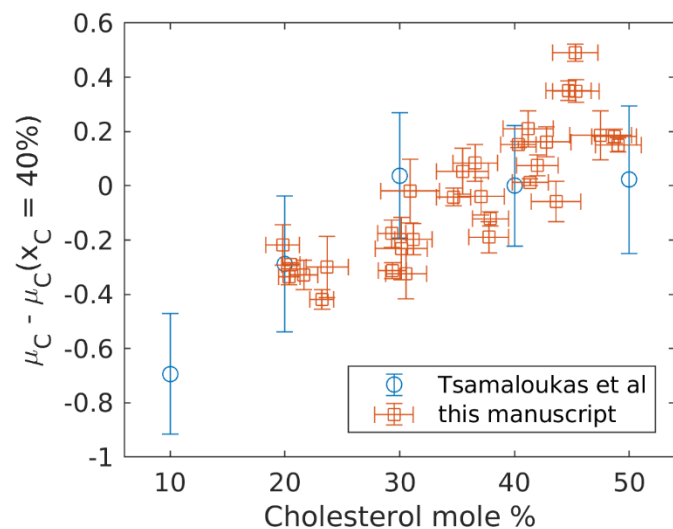
Supplemental Figure 3.11 Calibration of phospholipid concentration using NBD-PE.

Vesicles from the same stock solution initially prepared with 1:1 DOPC/DPPC + 40% cholesterol and interrogated over 2 days. On each day, the vesicle stock solution was diluted in cholesterol free M β CD buffer to yield the expected phospholipid concentration, and cholesterol concentration was measured using the methods described in the main text. (A) (left) Fitting the standard curve using expected phospholipid concentrations yields vesicles containing 38 \pm 2 mole % cholesterol, in good agreement with expectations of 40 mole %. (right) The lipid concentration can also be inferred directly from the measured cholesterol concentration, enforcing that the vesicle sample contains 40% cholesterol. (B) The same analysis from A but with the same vesicle stock interrogated on a different day. On this day, the cholesterol mole % in vesicles appeared to be 29 \pm 1% when the phospholipid concentration was calculated from the stock solution, suggesting the vesicle stock solution had a lower concentration than expected. (C) (top) on the first day, the standard curve connecting NBD intensity to phospholipid concentration is largely unchanged when using expected or inferred phospholipid concentration. (bottom) On the second day, the expected and inferred phospholipid concentrations yield different standard curves. Across days, NBD signals were read with the same parameters. Lower absolute values in day 2 vs. day 1 supports the conclusion that the stock concentration was lower than expected. In all cases, the inferred curve is used to convert NBD intensity to phospholipid concentration for subsequent analysis.



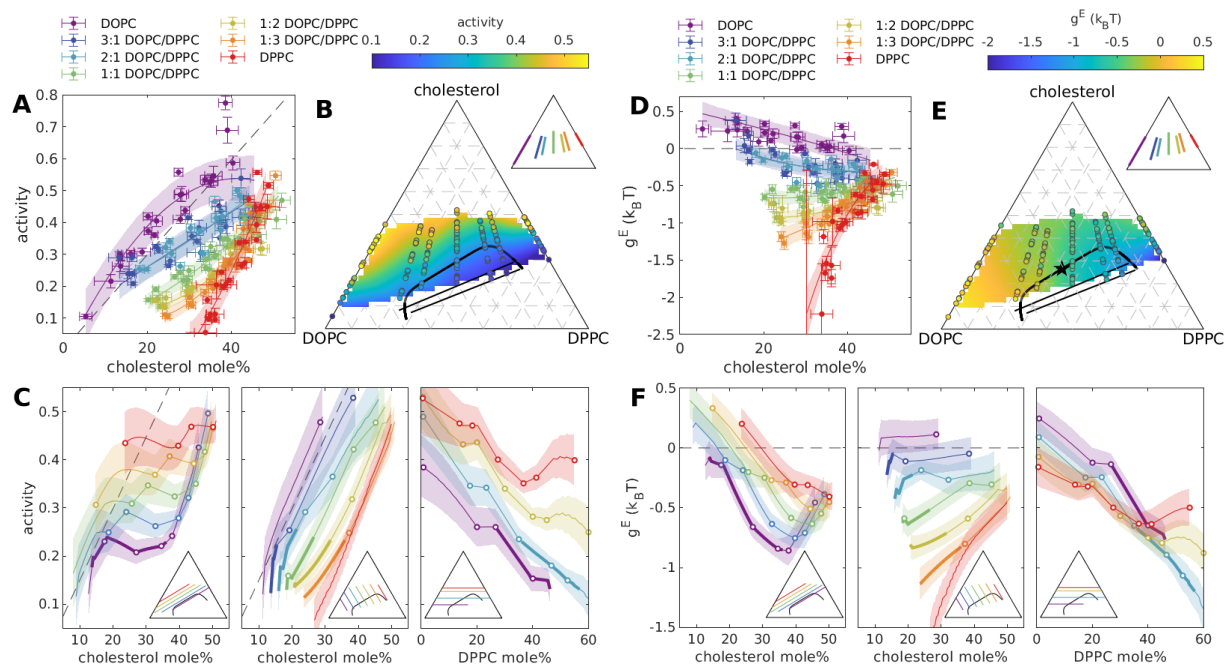
Supplemental Figure 3.12 Calibration relating % saturation of cholesterol in 5 mg/ml M β CD solutions to cholesterol activity.

(A) Measured % sat cholesterol values from aqueous solutions equilibrated with cholesterol in hexadecane as described in the main text. These values are obtained using a standard curve prepared from a 100% saturated solution equilibrated with cholesterol crystals in the absence of hexadecane. Values come close to but do not reach 100% saturation, likely because some M β CD can dissolve in hexadecane and some hexadecane could bind to M β CD in aqueous solution. To account for this, values are normalized to the levels measured for M β CD solutions equilibrated against fully saturated hexadecane solutions (magenta line at 93 ± 2 % sat chol). (B) Normalized values from A represent the chemical activity of cholesterol in 5 mg/ml M β CD solutions. Points are fit to a Langmuir isotherm of functional form $\%sat = FkA/(1+kA)$ to obtain the best fit values of $F = 2.6 \pm 0.3$ and $k = 0.6 \pm 0.1$.




Supplemental Figure 3.13 Comparison of μ_C measurements from the present manuscript to past work from Tsamaloukas et al (35).

μ_C measurements as a function of cholesterol mole % in POPC/cholesterol mixtures, referenced to μ_C at 40 mole % cholesterol. Measurements from (35) were obtained by transforming plotted partition coefficients K_x , with $\mu_C = k_B T (\ln x_C - \ln K_x)$, and error bars were obtained by propagating plotted error bars for K_x through that transformation.



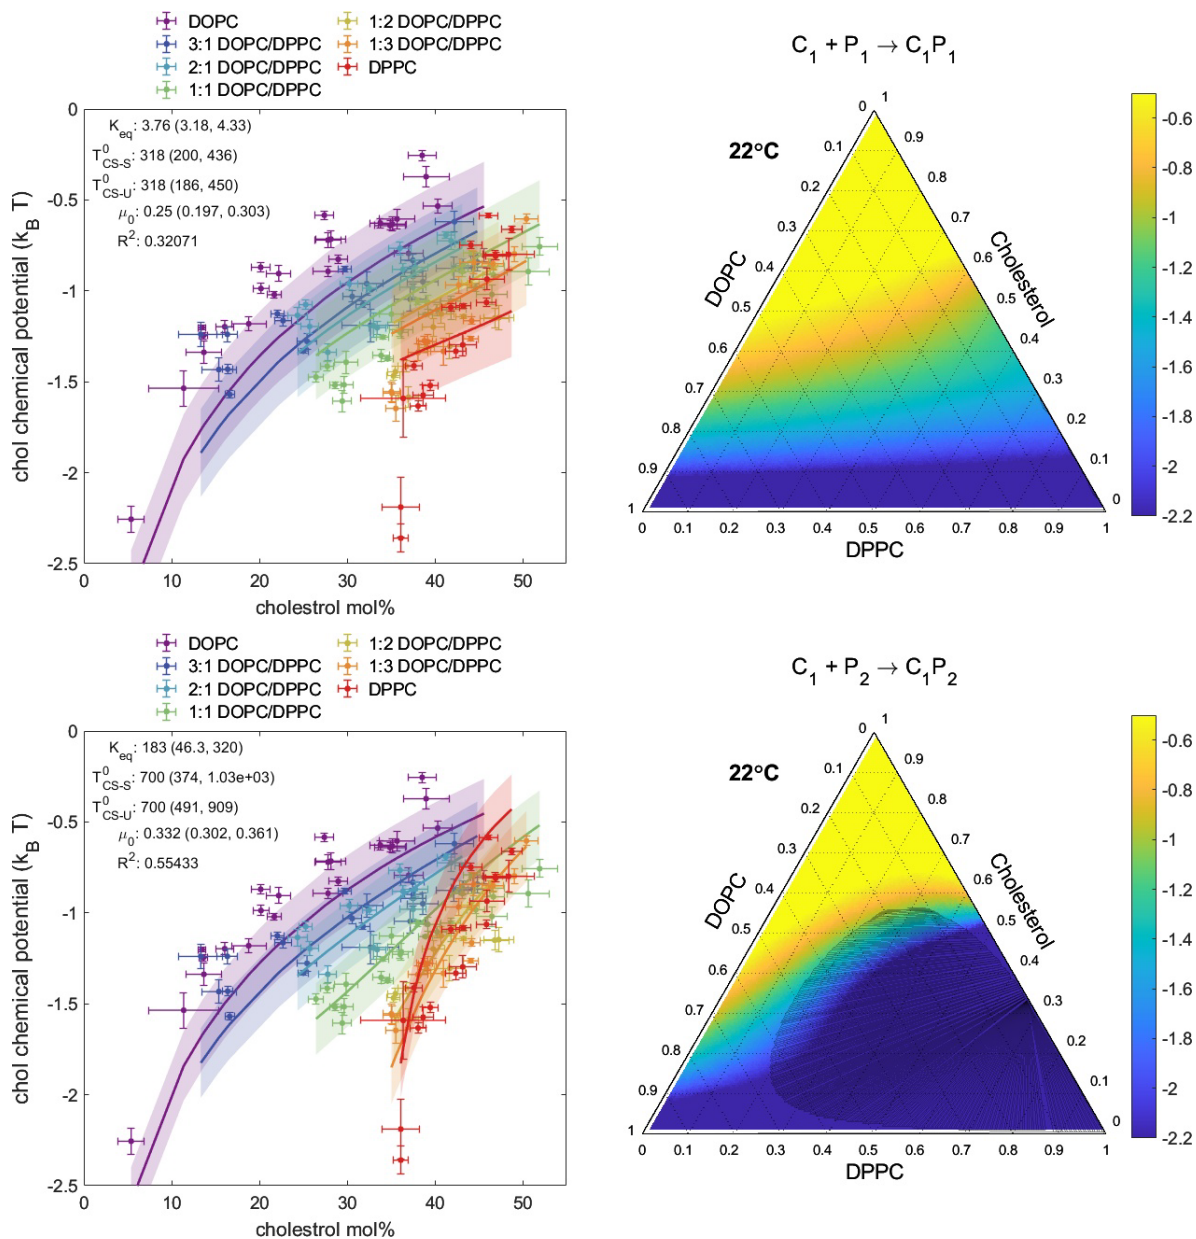
Supplemental Figure 3.14 Measurements of cholesterol chemical potential from Figure 3.3 replotted as chemical activity (A-C) and excess Gibbs free energy per molecule (D-F).

Measurements are processed exactly as in Figure 3.3 of the main text.



Supplemental Figure 3.15 Fraction of cholesterol that remains free.

Free cholesterol fraction for the models of Figure 3.4, where cholesterol reversibly forms a 1:1 complex with a DPPC molecule. Free (uncomplexed) cholesterol fraction is calculated directly from best-fit model parameters. (A) corresponds to the model of Figure 3.4A, where the only modelled interaction is the complex formation. (B) corresponds to the model of Figure 3.4B, where an additional repulsive interaction between complexes and DOPC is included. In the limit of low cholesterol and high DPPC, the model of (A) implies 16% of cholesterol remains free, and the model of (B) implies 12% of cholesterol remains free.



Supplemental Figure 3.16 Fitting μ_C to two regular solutions models with complexes and two repulsive interactions.

In addition to the repulsive interaction between complexes and unsaturated lipids that is considered in Figure 3.4B, these models include a repulsive interaction between complexes and saturated lipids as well. The strength of this interaction is given in terms of the critical temperature of a binary mixture of these components: T_{CS-S}^0 . The cholesterol:saturated lipid stoichiometries are 1:1 (top) or 1:2 (bottom). (left) Experimental data, fitted curves, and shaded confidence intervals as determined by MATLAB's predint. Fitted model parameters, confidence intervals, and Pearson's R^2 for this fit are also indicated. (right) Predicted μ_C surface for all ternary mixtures of cholesterol, DOPC and DPPC, using the fitted model. We rejected the bottom fit with 1:2 stoichiometry because a large part of the data is in a phase-separated region of the model, so that the mean-field chemical potentials calculated for those compositions are inaccurate – the true chemical potentials within the phase separated region would be those of the endpoints of the tie lines.

Table 3.1 Fitted parameter values and confidence intervals.

Parameter values and confidence intervals for the four fitted models from Figure 3.4 and Figure 3.5. Confidence intervals on parameters are for nonlinear least squares fits to the models as described in Methods. MF indicates a mean-field approximation. The interaction parameters for the lattice simulation were obtained by semi-manual optimization, so we did not obtain confidence intervals for those parameters. In addition, each model implies a region for which $\mu_C > 0$, implying spontaneous formation of cholesterol crystals at that such a condition. We consider this to imply a solubility limit for cholesterol, also shown in the table for each model. Separate values are shown for the Cholesterol-DOPC edge and the Cholesterol-DPPC edge. Finally, we indicate whether the fitted model has a phase-separated region.

Model:	Shown in Fig:	Fitted parameters (95% confidence interval)	Pearson's R ²	Implied solubility limits of cholesterol (mole %)	Supports room temperature phase separation?
Ideal solution with complexes	3.4A	$K_{eq}: 5.7 (4.3, 7.1)$ $\mu_C^0: 0.40 (0.34, 0.46) k_B T$	0.48	DOPC edge: 67% DPPC edge: 74%	N/A (model does not support phase separation)
MF Interacting solution with complexes	3.4B	$K_{eq}: 7.8 (5.8, 9.5)$ $\mu_C^0: 0.41 (0.35, 0.47) k_B T$ $T_{CS-U}: 274 (115, 433) K$	0.55	DOPC edge: 67% DPPC edge: 74%	No
MF Interacting solution without complexes	3.5A	$J_{CS}: -0.24 (-0.36, -0.13) k_B T$ $J_{CU}: 0.24 (0.18, 0.30) k_B T$ $J_{SU}: 0.51 (0.34, 0.68) k_B T$ $\mu_C^0: 0.07 (-0.05, 0.19) k_B T$	0.78	DOPC edge: 93% DPPC edge: 94%	Yes
Lattice simulation of Interacting solution without complexes	3.5B	$J_{CS}: -0.62 k_B T$ $J_{CU}: 0 k_B T$ $J_{SU}: 1.06 k_B T$ $\mu_C^0: 0.54 k_B T$	0.50	DOPC edge: 53% DPPC edge: 66%	Yes

3.8 Supplemental Notes

3.8.1 Supplemental Note 1: computation of chemical potentials in a mole fraction basis

While chemical potentials are typically defined in terms of partial derivatives with respect to number of each component, it is convenient here to compute derivatives with respect to mole fractions instead. This supplemental note derives expressions for the chemical potentials in terms of those derivatives.

For a system containing r molecular species, standard definitions of Gibbs free energy G and the chemical potentials give (assuming const. pressure and temperature for simplicity):

$$dG(N_1, N_2, \dots, N_r) = \sum_{i=1}^r \mu_i dN_i.$$

It is convenient to rewrite this in terms of $N = \sum_i N_i$, and $x_i = N_i/N$, for $i = 2, \dots, r$. For consistency, $x_1 = 1 - \sum_{i>1} x_i$ is taken to be dependent on the other x_i , and is no longer a formal variable of G . Note that the choice of the first component as the dependent variable is arbitrary, and the following forms apply whichever component is chosen. However, it is important to keep track of which one is dependent, because the partial derivatives obtained below are dependent on this choice. Then we have

$$\begin{aligned} dN_i &= N dx_i + x_i dN, i > 1 \\ dN_1 &= dN - \sum_{i>1} dN_i \\ &= dN \left(1 - \sum_{i>1} x_i \right) - N \left(\sum_{i>1} dx_i \right). \end{aligned}$$

Rewriting the exact differential in terms of the new variables, we obtain:

$$dG(N, x_2, \dots, x_r) = \sum_i \mu_i x_i dN + N \sum_{i>1} (\mu_i - \mu_1) dx_i$$

We can then read off the partial derivatives with respect to the new variables. The derivatives with respect to $x_i, i > 1$ now give differences of chemical potentials:

$$\frac{\partial(G/N)}{\partial x_i} = \mu_i - \mu_1, \quad i > 1$$

The final chemical potential can be obtained from the fact that G is an extensive function, so is proportional to N (at fixed mole fractions of the components):

$$\frac{\partial G}{\partial N} = \sum_i \mu_i x_i = \frac{G}{N}.$$

Putting this together with the above and defining $g = G/N$, we may conclude:

$$g - \sum_{i>1} x_i \frac{\partial g}{\partial x_i} = \mu_U \left(1 - \sum_{i>1} x_i \right) + \mu_U \sum_{i>1} x_i = \mu_1.$$

This expression for μ_1 can then be substituted into the above difference to obtain expressions for the $\mu_i, i > 1$:

$$\mu_i(x_2, \dots, x_r) = g(x_2, \dots, x_r) + \frac{\partial g}{\partial x_i} - \sum_{i>1} x_i \frac{\partial g}{\partial x_i}, i > 1.$$

These expressions will be used in the following to derive the function forms of the relevant chemical potentials of specific models.

3.8.2 Supplemental Note 2: Regular solution models

The theoretical models discussed in the main text take the form of regular solution models. A regular solution model for a mixture considers nearest-neighbor interactions and internal degrees of freedom of each of a set of components, while neglecting oriented interactions. Define the Hamiltonian for such a model by

$$\mathcal{H} = \sum_{i=1}^r \mu_i^0 + \sum_{i<j} N_{ij} J_{ij},$$

where N_i indicates the number of species i , μ_i^0 is the free energy per molecule of the internal degrees of freedom for species i , N_{ij} is the number of $i - j$ nearest-neighbor pairs, and J_{ij} specifies the interaction energy of such a pair. Note that we need not include $i - i$ interactions explicitly, because by an appropriate transformation of the J_{ij} and μ_i^0 , the J_{ii} can be set to 0 without altering the Hamiltonian overall(259).

In a mean field, sometimes known as zeroth order, approximation, we take the entropy of mixing to be that of an ideal mixture, and the average energy to be the energy of a well-mixed configuration, so that the free energy per molecule becomes

$$g_N = \frac{G}{N} = \sum_i \{k_B T x_i \log x_i + x_i \mu_i^0\} + z \sum_{i<j} x_i x_j J_{ij},$$

where, as above, $x_i = N_i/N$ is the mole fraction of each component, and z is the coordination number of the lattice in question, which we take to be 4 in all of our models. In some cases, such as the models that include cholesterol-phospholipid complexes that are described in Supplemental Note 4, where all non-zero interactions are repulsive, we find it convenient to rewrite the interaction parameters as (mean-field) critical temperatures for phase separation, by the transformation $zJ_{ij} = 2k_B T_{i-j}$, where T_{i-j} is then the critical temperature for phase separation of binary $i - j$ mixtures.


3.8.3 Supplemental Note 3: Mean field regular solution without complexes

The model that is fitted in Figure 3.5A corresponds to a regular solution, defined as above, with three components only, C, S, and U, corresponding to cholesterol, DPPC (saturated phospholipid), and DOPC (unsaturated phospholipid). All of the interaction parameters are allowed to vary freely: J_{CS} , J_{CU} and J_{SU} . We consider the regular solution free energy as a function of x_S and x_U and use the chemical potential identities derived above to obtain:

$$\begin{aligned} \frac{\mu_C}{k_B T} &= \frac{1}{k_B T} \left(g - x_S \frac{\partial g}{\partial x_S} - x_U \frac{\partial g}{\partial x_U} \right) \\ &= \log x_C + \frac{\mu_C^0}{k_B T} + z \frac{J_{CS}}{k_B T} (x_S^2 + x_U x_S) + z \frac{J_{CU}}{k_B T} (x_U^2 + x_U x_S) - z \frac{J_{SU}}{k_B T} (x_S x_U) \end{aligned}$$

3.8.4 Supplemental Note 4: mean field regular solution models with complexes

Following McConnell and Radhakrishnan(255, 257, 258), we model a lipid bilayer composed of cholesterol, a saturated lipid (DPPC), and an unsaturated lipid (DOPC). As above, C, S, and U are used as shorthand for cholesterol, DPPC (saturated lipid) and DOPC (unsaturated lipid), respectively. A condensed complex written CS is allowed to form between cholesterol and DPPC in a reversible reaction assumed to have $q:p$ stoichiometry with equilibrium constant K_{eq} .



For later convenience we distinguish N_i , the total number of lipid species $i \in \{C, S, U\}$ in the system before the reaction, from N_i' , the number of species i after the reaction comes to equilibrium. Furthermore, $N = N_C + N_S + N_U$, and $x_i = N_i/N$, and $N' = N'_C + N'_S + N'_U + N_{CS}$,

and $x'_i = N'_i/N'$. As above, we begin with the regular solution free energy per molecule (N') of an equilibrium mixture of C, S, U, and CS:

$$g_{N'}^{RS} = \frac{G}{N'} = \sum_i x'_i (k_B T \ln x'_i) + 2k_B \sum_{i<j} x'_i x'_j T_{i-j}^0,$$

where μ_i^0 is standard chemical potential of pure component i , x'_i is equilibrium mole fraction of i , and T_{ij}^0 is the critical temperature of the $i - j$ binary pair.

To account for the complexation reaction we add to this a term:

$$G_{complex} = N_{CS} \Delta G_{complex} = -k_B T N_{CS} \ln K_{eq},$$

and finally allow for a non-zero μ_C^0 in order to compare to refer the chemical potential to the energy scale set by taking the cholesterol crystal as standard state:

$$G_{standard} = N'_C \mu_C^0$$

Finally, following (258), we assume all critical temperatures are zero except T_{CS-S}^0 (immiscibility between CS and S) and T_{CS-U}^0 (immiscibility between CS and U).

Following these claims, the free energy per molecule (of which there are N' in number) denoted as g' takes the form

$$\frac{g'}{k_B T} = \left(x'_C (\mu_C^0/k_B T + \ln x'_C) + x'_S \ln x'_S + x'_U \ln x'_U + x_{CS} \ln \frac{x_{CS}}{K_{eq}} \right) + \frac{2x_{CS}}{T} (T_{CS-S} x'_S + T_{CS-U} x'_U) \quad (1)$$

It is useful to be able to write the free energy in terms of the mole fractions that are inserted at the beginning of the experiment x_i , rather than the ones after the reaction reaches equilibrium x'_i . Begin by defining the system in terms of a reaction progress parameter γ which is minimized in all calculations of the total free energy G^*

$$\gamma = \frac{N_{CS}}{N'_C + N'_S + N'_U + (q + p)N_{CS}}$$

We can write the transform in terms of this parameter and a parameter λ as follows

$$x_{CS} = \gamma \lambda$$

$$x_{C'} = (x_C - q\gamma)\lambda$$

$$x_{S'} = (x_S - p\gamma)\lambda$$

$$x_{U'} = x_U \lambda.$$

λ can then be determined by noting our constraints: $x_C + x_S + x_U = x_{C'} + x_{S'} + x_{U'} + x_{CS} = 1$.

Solving with these constraints yields.

$$\lambda = [1 + (1 - q - p)\gamma]^{-1}$$

Using the same constraints, we can solve for gamma in terms of x_{cs}

$$\gamma = \frac{x_{cs}}{1 + (p + q - 1)x_{cs}}$$

and then use these two relations to eliminate γ and λ from our transformations entirely

$$\begin{aligned}x_{c'} &= x_c[1 + (p + q - 1)x_{cs}] - qx_{cs} \\x_{s'} &= x_s[1 + (p + q - 1)x_{cs}] - px_{cs} \\x_{u'} &= (1 - x_c - x_s)[1 + (p + q - 1)x_{cs}].\end{aligned}$$

Because the number of particles in the system is not conserved through the reaction it is useful to know how the initial number of molecules inserted into the system N relates to the final number N' . We can use our relation for $x_{u'}$ combined with the fact that $N'_u = N_u$ to find this relation.

$$\begin{aligned}x'_{u'} &= x_u \lambda \\x'_{U'} &= x_U \lambda \\ \lambda &= \frac{N}{N'}.\end{aligned}$$

It is worth underlining that with this relation, $g = G/N$ may be obtained from $g' = G/N'$ by

$$g = g' \lambda^{-1}$$

Letting the reaction equilibrate:

To extract information from g or g' we need to minimize G with respect to the fraction of complex x_{cs} . This is equivalent to minimizing $g = G/N$, because N does not vary as the reaction progresses. So we define $g^* = g(x_c, x_s, x_{cs}^*)$ such that

$$\left. \frac{\partial g}{\partial x_{cs}} \right|_{x_{cs}=x_{cs}^*} = \lambda^{-1} \left. \frac{\partial g'}{\partial x_{cs}} \right|_{x_{cs}=x_{cs}^*} = 0$$

Unfortunately computing such an x_{cs}^* leads to a transcendental equation which does not have an analytical solution. Thus, we compute x_{cs}^* numerically using MATLAB's numerical minimization routine `fminbnd`.

Calculating μ'_c :

With numerical x_{cs}^* , we may return to general stoichiometry $q:p$ and the primed basis for simplicity to calculate the chemical potential of free cholesterol (now ignoring the denominator $k_B T$) using the relations of Supplemental Note 1:

$$\mu_{c'} - \mu_{u'} = \left(\frac{\partial g^*}{\partial x'_c} \right)_{T,P,x'_{i \neq c}}$$

$$\mu_{u'} = g^* - \sum_{i \neq u} x'_i \frac{\partial g^*}{\partial x'_i}$$

We can rewrite our expression for $\mu_{c'}$ as

$$\begin{aligned} \mu_{c'} &= \left(\frac{\partial g^*}{\partial x'_c} \right)_{T,P,x'_{i \neq c}} + \mu_{u'} \\ \mu_{c'} &= \left(\frac{\partial g^*}{\partial x'_c} \right)_{T,P,x'_{i \neq c}} + g^* - x'_s \frac{\partial g^*}{\partial x'_s} - x'_c \frac{\partial g^*}{\partial x'_c} - x_{cs}^* \frac{\partial g^*}{\partial x_{cs}^*} \\ \mu_{c'} &= g^* + (1 - x'_c) \frac{\partial g^*}{\partial x'_c} - x'_s \frac{\partial g^*}{\partial x'_s} - x_{cs}^* \frac{\partial g^*}{\partial x_{cs}^*} \end{aligned}$$

where $g^*(x'_c, x'_s, x_{cs}^*)$ in units of $k_B T$ is

$$\begin{aligned} g^* &= \left(x'_c (\mu_c^0 + \ln x'_c) + x'_s \ln x'_s + (1 - x'_c - x'_s - x_{cs}^*) \ln(1 - x'_c - x'_s - x_{cs}^*) + x_{cs}^* \ln \frac{x_{cs}^*}{K_{eq}} \right) \\ &\quad + \frac{2x_{cs}^*}{T} (T_{cs-s} x'_s + T_{cs-u} (1 - x'_c - x'_s - x_{cs}^*)). \end{aligned}$$

Proceeding by first looking at each partial derivative in $\mu_{c'}$

$$\begin{aligned} \frac{\partial g^*}{\partial x'_c} &= \mu_c^0 + \log \left(\frac{x'_c}{1 - x'_c - x'_s - x_{cs}^*} \right) - \frac{2x_{cs}^*}{T} T_{cs-u} \\ \frac{\partial g^*}{\partial x'_s} &= \log \left(\frac{x'_s}{1 - x'_c - x'_s - x_{cs}^*} \right) + \frac{2x_{cs}^*}{T} (T_{cs-s} - T_{cs-u}) \\ \frac{\partial g^*}{\partial x_{cs}^*} &= \log \left(\frac{x_{cs}^*}{K_{eq} (1 - x'_c - x'_s - x_{cs}^*)} \right) + \frac{2x'_s}{T} T_{cs-s} - \frac{2(x'_c + x'_s + 2x_{cs}^* - 1)}{T} T_{cs-u}. \end{aligned}$$

Then combining these results, we arrive at

$$\mu_{c'} = \mu_c^0 + \log(x'_c) - \frac{2x_{cs}^*}{T} [(1 - x'_c - x'_s - x_{cs}^*) T_{cs-u} + x'_s T_{cs-s}]$$

Note that for boundary data (where x_s or $x_u = 0$) we need to define separate free energies g_{DPPC}^* and g_{DOPC}^* in the limit as the relevant mole fraction goes to 0, yielding ultimately:

$$\begin{aligned} \mu_{c'}^{DPPC} &= \mu_c^0 + \log(x'_c) - \frac{2x_{cs}^* x'_s}{T} T_{cs-s} \\ \mu_{c'}^{DOPC} &= \mu_c^0 + \log(x'_c). \end{aligned}$$

In equilibrium $\mu_{c'}$ will be equivalent to μ_c giving us a final piece-wise expression for μ_c as a function of the x_s, x_c, x_{cs} basis (but written here in the x'_i basis for brevity):

$$\mu_c = \begin{cases} \mu_c^0 + \log(x'_c) & x_s = 0 \\ \mu_c^0 + \log(x'_c) - \frac{2x_{cs}^*}{T} [(1 - x'_c - x'_s - x_{cs}^*)T_{cs-u} + x'_s T_{cs-s}] & x_i \neq 0 \\ \mu_c^0 + \log(x'_c) - \frac{2x_{cs}^* x'_s}{T} T_{cs-s} & x_u = 0 \end{cases}$$

3.8.5 Supplemental Note 5: Estimating chemical potentials from lattice model simulations of the regular solution model without complexes

The model that is shown in Figure 3.5B corresponds to the same model as Figure 3.5A, but without making the mean field approximation. Therefore, we cannot write an explicit expression for the free energy or its derivatives, and have to estimate these from the simulations themselves. Simulations are carried out as described in Methods, and then analyzed as follows.

Estimation of derivatives $\partial g / \partial x_i$ in a lattice model using a Monte Carlo scheme:

Consider a Monte Carlo update scheme that allows components from the lattice simulation to be exchanged with a bath of particles. It must satisfy detailed balance in order for it to equilibrate. This means

$$P_{A \rightarrow B} P_A = P_{B \rightarrow A} P_B,$$

where A, B indicate states of the system, P_A is the probability of state A in equilibrium, and $P_{A \rightarrow B}$ is the probability under the Monte Carlo scheme of a state transition from A to B . The equilibrium distribution is given by the grand canonical ensemble

$$P_A \propto \exp\left(-\beta \left[E_A - \sum_k \mu_k N_k^{(A)} \right]\right),$$

where E_A indicates the (total) energy of state A , k indexes the chemical species of the mixture, μ_k is the chemical potential of species k , and $N_k^{(A)}$ is the number of molecules of species k that are present in state A . The proportionality constant is the grand partition function, which need not be estimated here because we will only need ratios of probabilities.

We will only consider schemes with moves that swap the species that is present at a random single site i . Write $(s_i, k \rightarrow l)$ to denote such a swap. For such a swap, the ratio P_B/P_A only depends on the interaction energies between site i and its four neighbors. In particular,

$$\frac{P_{s_i,l \rightarrow k}}{P_{s_i,k \rightarrow l}} = \frac{P_{s_i=k}}{P_{s_i=l}} = \exp(-\beta[(E_k - E_l) - (\mu_k - \mu_l)]),$$

where the first equality comes from rearranging the detailed balance condition, and the second comes from the grand canonical ensemble. E_k now indicates the sum of the interactions of site i with its neighbors if the species at that site is k . The energy of the rest of the lattice need not be considered, because it does not depend on the species at site i , and thus cancels in the ratio.

One proposal for transition probabilities (conditioned on site i being chosen as the random site to swap) that satisfies the above condition is

$$P_{s_i,k \rightarrow l} = C \cdot \exp(-\beta[(E_l - E_k) - (\mu_l - \mu_k)]/2)$$

with the constant C chosen so that no transition probability exceeds 1. Then

$$\frac{P_{s_i,l \rightarrow k}}{P_{s_i,k \rightarrow l}} = \frac{\exp(-\beta[(E_k - E_l) - (\mu_k - \mu_l)]/2)}{\exp(-\beta[(E_l - E_k) - (\mu_l - \mu_k)]/2)} = \exp(-\beta[(E_k - E_l) - (\mu_k - \mu_l)])$$

as required. Consider the simulations as a sample from the equilibrium distribution, i.e. under the assumption that the μ_k have been chosen such that the average number of molecules of species k , $\langle N_k \rangle$ under the grand canonical ensemble is equal to the (fixed) number that are actually present in the given simulation. Consider the sum

$$\begin{aligned} W_{k \rightarrow l} &= \sum_i \delta(s_i = k) \exp(-\beta[E_l - E_k]/2) \\ &= \exp(-\beta(\mu_l - \mu_k)/2) \sum_i \delta(s_i = k) \exp[-\beta[(E_l - E_k) - (\mu_l - \mu_k)]/2] \\ &\propto \exp(-\beta(\mu_l - \mu_k) / 2) P_{s_i,k \rightarrow l} P_{s_i=k}. \end{aligned}$$

Therefore, by taking advantage of detailed balance, we can conclude that

$$\frac{W_{k \rightarrow l}}{W_{l \rightarrow k}} = \exp(\beta(\mu_k - \mu_l)) \frac{P_{s_i,k \rightarrow l} P_{s_i=k}}{P_{s_i,l \rightarrow k} P_{s_i=l}} = \exp(\beta(\mu_k - \mu_l)).$$

so that

$$\mu_k - \mu_l = k_B T \log \left(\frac{W_{k \rightarrow l}}{W_{l \rightarrow k}} \right).$$

To obtain the chemical potentials themselves from these differences, Supplemental Note 1 gives us a way to write these differences as partial derivatives of the free energy per molecule g . For the three species C,S and U, taking as dependent mole fraction $x_U = 1 - x_C - x_S$, we have

$$\frac{\partial g(x_C, x_S)}{\partial x_C} = \mu_C - \mu_S.$$

Then we can estimate the overall free energy by numerically integrating these partial derivatives:

$$g(x_C, x_S) = \mu_0 + \int_{x_C^*}^{x_C} \frac{\partial g}{\partial x_C}(x, x_S^*) dx + \int_{x_S^*}^{x_S} \frac{\partial g}{\partial x_S}(x_C, x) dx.$$

where x_C^* and x_S^* are chosen as the origin of the integration and μ_0 is an integration constant. We evaluate the derivatives using simulations that are closely spaced in composition space, with spacing of at most .05 in both x_C and x_S . At the edges of composition space where one component vanishes, finer spacing is chosen. These samples are then interpolated using cubic spline interpolation to obtain the derivatives at arbitrary interior points of composition space. Integration is carried out as a Riemann sum with spacing of at most .002 in x_C and x_S .

Finally, with g and the various $\mu_i - \mu_j$ in hand, we can compute for any composition:

$$\mu_C = g - x_U(\mu_U - \mu_C) - x_S(\mu_S - \mu_C)$$

as derived in Supplemental Note 1.

3.8.6 Supplemental Note 6: Comparison of fitted mean field regular solution model interaction energies to past estimates for a similar mixture.

Tsamaloukas et al have reported regular solution interaction parameters for mixtures of POPC/sphingomyelin/cholesterol at 37C, as inferred from differential partitioning of the detergent Triton X-100 into these membranes (237). Their model is equivalent to the mean field regular solution model we have fit, although the definitions of their interaction parameters differ from ours by a factor of 4, the lattice coordination number of our model. Converting their fitted interaction parameters to our definitions thus yields $J_{SM-Chol} = -1.5 \pm .75$ kBT, $J_{POPC-Chol} = 0.5 \pm .25$ kBT and $J_{SM-POPC} = 0 \pm .25$ kBT. While that system has important differences to our room temperature DOPC/DPPC/cholesterol vesicles, we observed cholesterol-lipid interaction energies J_{C-S} and J_{C-U} similar to $J_{SM-Chol}$ and $J_{POPC-Chol}$, respectively. A notable difference is in the third parameter J_{S-U} , which we found to be substantially repulsive. These discrepancies may be related to the choice of lipids or to the higher temperatures at which those experiments were conducted.

Chapter 4 A Mean Shift Algorithm for Drift Correction in Localization Microscopy

This work proposes, validates, and applies a method for correcting sample drift in SMLM data. It is now published in *Biophysical Reports* (271). Drift correction is an important step in the analysis of SMLM data because data collection is slow, and the microscope stage can drift tens to hundreds of nm over relevant acquisition times – much more than the localization precision of the technique. Our method has several advantages over the existing alternatives, including lower computational resource requirements and improved robustness allowing for improved time resolution. These advantages mainly derive from the simplicity of the method. This chapter was written collaboratively with Frank Fazekas, Sumin Kim, Ryan Bogucki, and Sarah Veatch. I provided the idea for the method and much of the theoretical foundation. Frank Fazekas wrote much of the code and performed many of the analyses, with supervision from Sarah Veatch and myself. The three of us then wrote the manuscript together. The datasets of figures 3 and 4 were contributed by Sumin Kim and Ryan Bogucki, respectively.

4.1 Abstract

Single molecule localization microscopy (SMLM) techniques transcend the diffraction limit of visible light by localizing isolated emitters sampled stochastically. This time-lapse imaging necessitates long acquisition times, over which sample drift can become large relative to the localization precision. Here we present a novel, efficient, and robust method for estimating drift using a simple peak-finding algorithm based on mean shifts that is effective for SMLM in 2 or 3 dimensions.

4.2 Introduction

Stochastic super-resolution microscopy techniques such as STORM (38, 40) and PALM (37, 39) exploit photoswitching of fluorescent probes to enable imaging of densely labeled samples with resolutions an order of magnitude smaller than the diffraction limit of visible light. Sparsely distributed point spread functions (PSFs) of single emitters are identified in individual image frames, and their centroids are determined according to an appropriate fitting algorithm. The axial

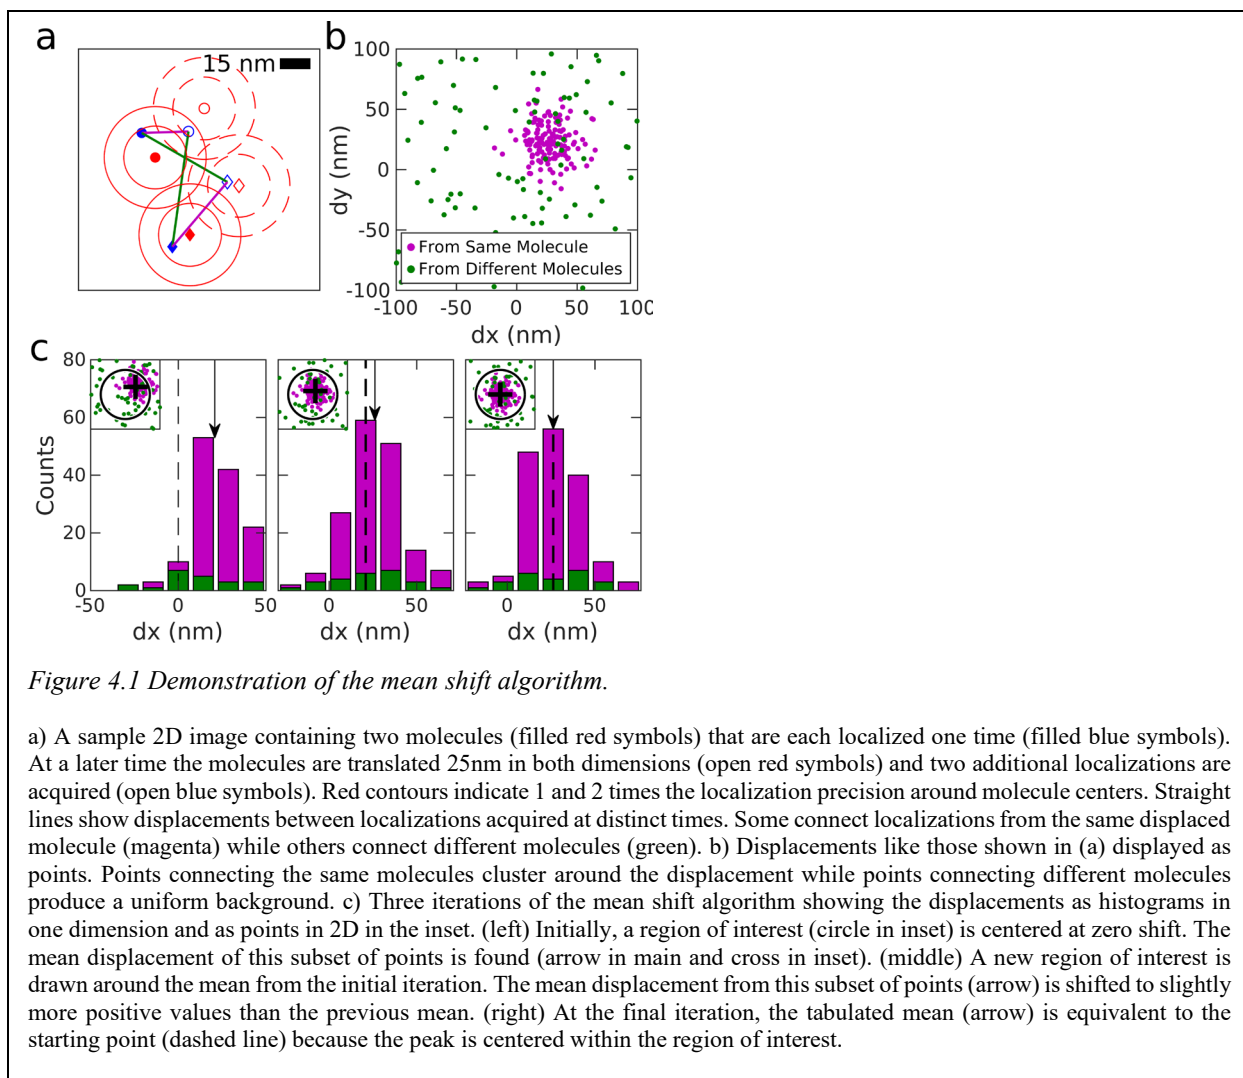
position of molecules can be encoded in their PSFs through engineering measures utilizing astigmatism (272, 273), multifocal plane imaging (274), or a double helix PSF (74). The final reconstruction is typically a 2D or 3D histogram of these single-molecule positions.

Drift due to thermal expansion or mechanical instabilities can degrade image quality over the course of image acquisition, which typically occurs on the timescale of minutes. Drift compensation requires either active stabilization of the microscope (275–279) or *a posteriori* computation of the drift curves either using fiducial markers (62, 63, 280–282) or the acquired single molecule localizations (59–61, 64, 65, 283–285). In this report, we present a mathematically simple approach to drift correction using a mean shift (MS) algorithm (286–288) for static SMLM datasets without fiducial markers, with some advantages over past approaches that use nonlinear least squares (NLLS) fitting of image-based cross-correlations (59–61).

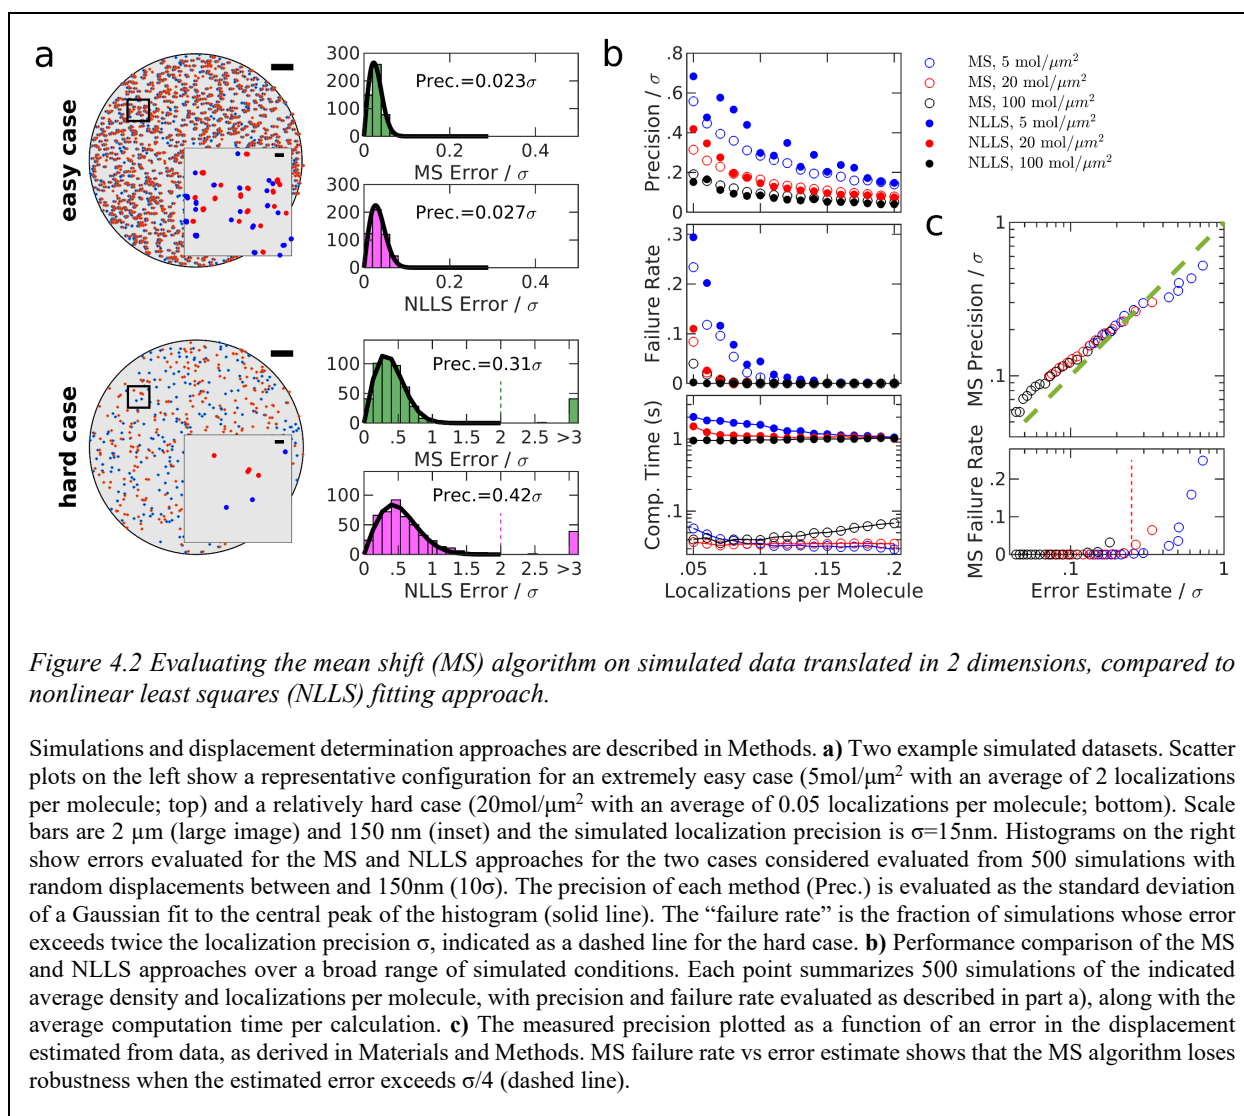
4.3 Results

A graphical illustration of the mean shift (MS) algorithm as applied to sample 2D localizations is presented in Figure 4.1. The localizations all lie in one of two datasets which sample the same uniformly distributed emitters, but with a constant relative shift $\mathbf{r}_{\text{shift}}$ in space. The first step of the algorithm is to extract pairwise displacements between all localizations across the two datasets. When individual displacements are plotted as points (Figure 4.1b), displacements arising from the same labeled objects (magenta points) cluster around $\mathbf{r}_{\text{shift}}$, while displacements arising from different objects (green points) distribute randomly over space. The mean shift algorithm determines the center of the peak of the distribution through iteration (286–288). At each iteration, all pairs within the radius of consideration are extracted, and the updated shift estimate is the centroid of these pairs. The uniformly distributed background will tend to bias the centroid towards the center of the observation window, while the peak moves the mean toward $\mathbf{r}_{\text{shift}}$. The observation window is then redrawn around the new mean and the process is repeated until the peak is centered in the observation window. Three iterations of the algorithm are visualized in Figure 4.1c.

While the emitters of Figure 4.1 are distributed uniformly in space, leading to the uniform distribution of the pairs from different emitters, the MS method does not depend on this assumption. In samples where emitters are organized into structures or randomly clustered, the pairs arising from different emitters are also more likely to be at shorter distances, so that the



distribution of green points in Figure 4.1b will also be peaked at $\mathbf{r}_{\text{shift}}$. However, in our experience, pairs of localizations from the same emitter are more important for the MS and other drift estimates. We also note that our analysis assumes that emitters that are localized in one dataset remain within the field of view in the second dataset, and vice versa. This may not always be the case, and could in principle lead to bias in shift estimates, but in practice this is typically a negligible effect. Roughly, the contribution of fluorophores near the edge of the field of view may be biased by up to about the localization precision, and the fraction of fluorophores that are affected is restricted to those that lie within about a localization precision from the edge of the field of view, in the direction of the drift. So, for example, in a 100 μm field of view with localization precision of 15 nm, we would expect this bias to be on the order of a picometer.



In order to benchmark this mean shift (MS) approach, we evaluated the ability of the algorithm to detect known shifts of simulated datasets of a circular test cell, as summarized in Figure 4.2. Shifts were estimated by both the MS algorithm and by nonlinear least squares (NLLS) fitting of a Gaussian to the spatial cross-correlation function of the two datasets, as implemented in the supplementary software provided with (61). The performance of each algorithm was similar for easy cases that produce a well-defined peak at $\mathbf{r}_{\text{shift}}$. An extremely easy case is depicted at the top of Figure 4.2a, where single molecules are well spaced (surface density = 5/ μm^2) and where their positions are well sampled in both frames (twice per molecule on average). In this case, the shift can be clearly identified by eye, and both algorithms reliably and accurately identify the displacement between frames. The simulation depicted at the bottom of Figure 4.2a represents a

much harder case, where molecules are present at higher surface density ($20/\mu\text{m}^2$) and only 1 in 20 molecules are imaged on average in a given dataset. In this case, MS modestly outperforms NLLS fitting, both by locating the peak with improved precision and by more reliably finding the peak overall. These trends hold over simulations conducted over a broad range of molecular densities and localizations per molecule (Figure 4.2b). We also estimated shifts from the overall center of mass of each dataset, which yielded precisions more than an order of magnitude worse than both the MS and NLLS methods. Moreover, MS is more computationally efficient than NLLS, largely because Fast Fourier transforms (FFTs) are not computed in the mean shift approach. This improvement in speed is enabled through the use of a particularly efficient algorithm from the R package spatstat (289) to extract pairwise displacements between nearby points (see Methods).

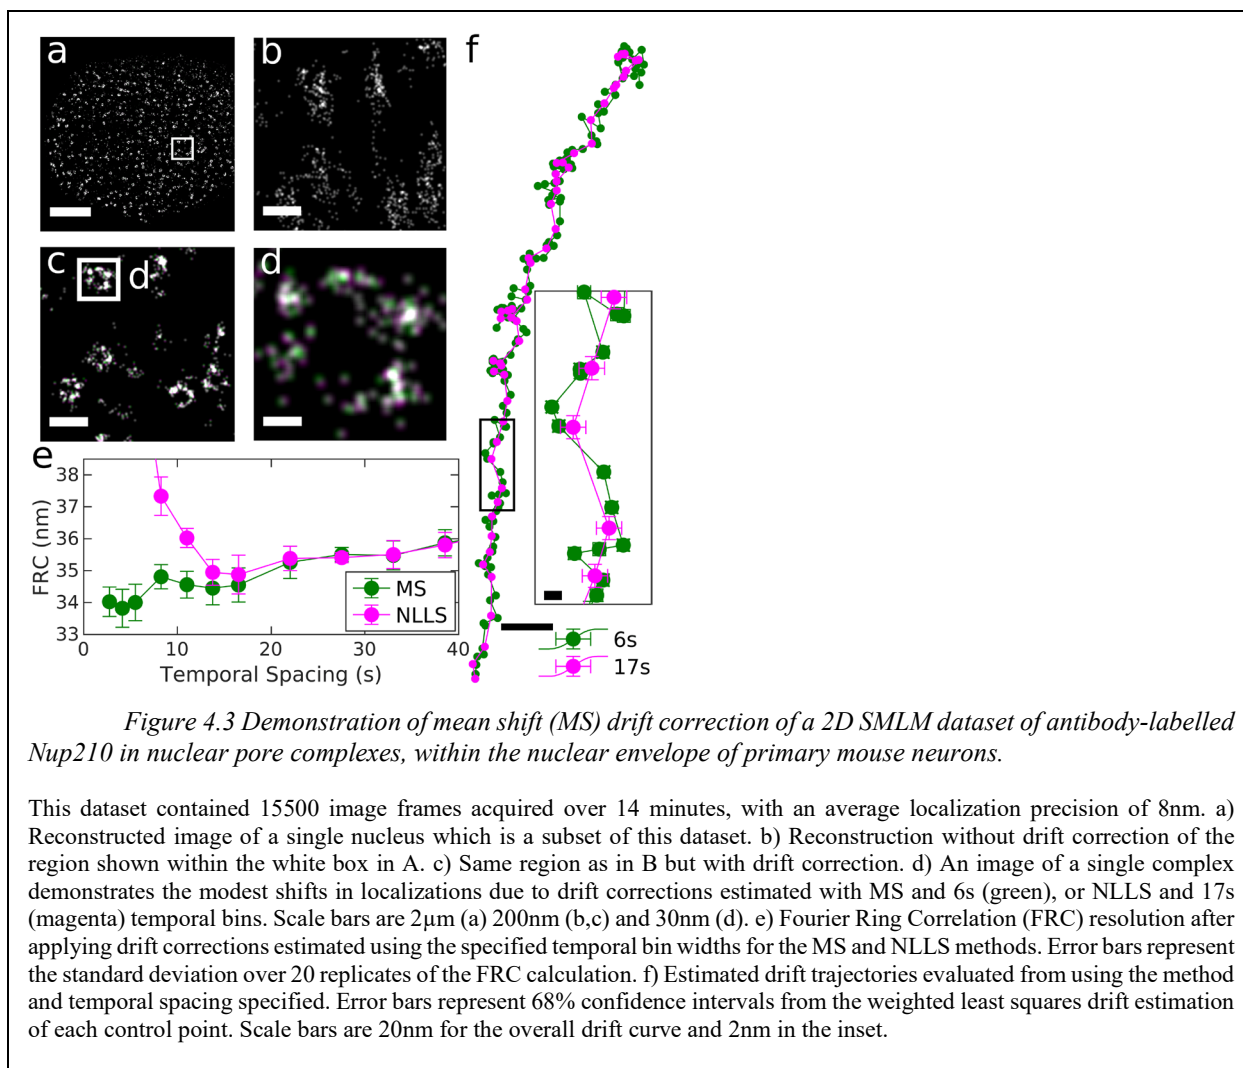
For large displacements, both the MS and NLLS algorithms applied in Figure 4.2 require an initial step to identify an approximate starting point for the higher accuracy calculation. Supplemental Figure 4.5 (see Section 4.5 for supplemental figure) shows the failure rate of each algorithm as a function of the distance of the start point from the true shift. MS robustly identifies the main peak over a broad range of simulation conditions as long as it resides within the initial observation window, so large shifts can be identified simply using a large window in the first iteration. This window is typically 100nm for experimental localizations and 150nm for the simulations of Figure 4.2. NLLS robustly identifies the main peak when the starting point for the computation falls within the localization precision of the peak of the cross-correlation function. In many practical cases the peak is much farther from the origin than the localization precision, so a separate method is needed to identify a suitable starting point. This is accomplished here using a particularly effective algorithm that identifies the global maximum in a smoothed cross-correlation function, as described in the supplemental material of (61). The robustness of the NLLS fitting approach is dependent on the ability of this algorithm to identify a suitable starting point over a broad range of simulation conditions. Note that datasets where the emitter distribution is highly structured or clustered typically lead to improved performance of the start point identification routine, by introducing a broad peak in the cross-correlation function in addition to the sharp peak that represents repeat localizations of the same fluorophore.

This mean shift approach is applied to SMLM localizations that experience continuous drift by distributing localizations into non-overlapping temporal bins with equal numbers of

frames, and displacement estimates are tabulated between all possible pairs of bins. The number of frames in each temporal bin is an important parameter: short temporal bins have few localizations per molecule, so individual displacements may be estimated imprecisely. Long temporal bins have more localizations per molecule and more precise drift estimates, but reduce the time resolution of the drift estimate. A linear least squares fitting algorithm is then used to generate a trajectory that passes through control points positioned at times centered on each temporal bin, as described previously (61), taking advantage of the high redundancy to improve precision of the control points. We have slightly modified this past approach by including weights in the linear least squares fitting, where weights are determined directly from data using a relation that approximates error in the mean-displacement (described in Methods) as demonstrated in simulated datasets (Figure 4.2c). Briefly, errors are reduced when there are more pairs originating from the same molecules (magenta points in Figure 4.1) and errors increase when more pairs originate from different molecules within the observation window (green points in Figure 4.1). Estimated errors can also act as a proxy for overall reliability of the algorithm. Figure 4.2c also shows that MS reliably finds the desired peak when the estimated error remains smaller than one quarter of the localization precision. This observation can act as a guide when selecting the number of frames included in temporal bins.

It is tempting to distribute frames into overlapping temporal bins, which in principle could improve time resolution while retaining a sufficient number of localizations to accurately determine displacements. However, we find that drift estimates from overlapping time bins are subject to substantial bias, underestimating the actual displacements accrued over time (Supplemental Figure 4.6a). This occurs because the same localizations are present in adjacent bins, biasing the result towards $\mathbf{r}_{\text{shift}} = 0$. Similar bias can arise even in the absence of overlapping time bins because SMLM data frequently contains time-correlated localizations arising from the finite off rates of fluorescent blinking (PALM/dSTORM) or binding (PAINT). These factors mean that pairs of localizations from the same fluorophore are mostly from time separations that are shorter than the time difference between the bin centers, and therefore underestimate the average drift between the bins.

The MS approach is applied to a 2D experimental dataset of nuclear pore complexes (NPCs) in Figure 4.3. NPC assemblies are labeled with primary and secondary antibodies against Nup210 within the nuclear envelope of intact primary mouse neurons and Figure 4.3a-d show



reconstructed images at various magnifications. Figure 4.3b is a reconstruction produced without drift correction in which localizations from single NPCs are smeared over a large area, highlighting the importance of drift correction.

The performance of the mean shift algorithm was tested on this dataset by generating multiple drift trajectories through binning with different temporal resolutions. These trajectories were each applied to the full SMLM dataset, and Fourier Ring Correlation (FRC) (80, 81) was used to quantify image resolution (Figure 4.3e). For comparison, we conducted drift corrections using the redundant cross-correlation (RCC) NLLS approach as described previously (61). In this case, MS modestly outperforms NLLS fitting, allowing for accurate drift correction with smaller temporal bins and modestly improving the resolution of the reconstructed image. We used this dataset to explore possible bias introduced due to temporal correlations of single fluorophore blinking by

running the linear least squares algorithm including or excluding adjacent pairs of bins on the data of Figure 4.2. We found no significant difference between the two cases (Supplemental Figure 4.6b), indicating that the impact of this bias is negligible within experimental errors. Additional diagnostics for the MS and NLLS approaches are shown in Supplemental Figure 4.7.

Drift trajectories are shown in Figure 4.3f for temporal bin-widths that produce accurate FRC metrics for the MS and NLLS approaches. For MS, a temporal bin slightly larger than the minimum from the FRC curve is used, since this produces smaller errors on individual control points. As expected, the drift trajectories follow the same general shape but the trajectory generated from MS has improved time resolution. In parts of the trajectory, the errors of the control points are smaller than the distance between the trajectories. In these regions, higher time-resolution yields improved spatial resolution in the final reconstructed image. While the differences in the trajectories are significant, their impact is not apparent when viewing reconstructed images of entire nuclei or collections of NPCs, as in Figure 4.3a,c. Differences become more apparent in images of individual pores, where displacements of several nm shift the relative positions of labeled subunits (Figure 4.3d).

The MS algorithm is easily extended to localizations acquired in 3D, where performance improvements are more evident compared to the established NLLS approach. Since the MS algorithm uses points instead of reconstructed images and FFTs, it can be extended into 3D without needing expanded memory resources, which limits the practical application of NLLS in 3D. Instead, the 3D application of NLLS drift correction is typically accomplished by generating 2D projections that contain less information than the 3D localizations they are produced from (61). To see why, consider a pair of emitters that are close together in x-y but far apart in z. Pairs of localizations from this pair of fluorophores will be included in 2D MS drift estimation when using data projected into the x-y plane, but excluded from the full 3D drift estimation method, by virtue of their large separation in z. We compare the precision and robustness of 3D MS and NLLS on simulated localizations spread over a cylindrical volume in Supplemental Figure 4.8, where the NLLS correction is performed on projections into the x-y, x-z, and y-z planes as described in (61). We also directly compare the x-y performance of the full 3D MS method to the 2D MS method performed on data projected into the x-y plane (Supplemental Figure 4.9).

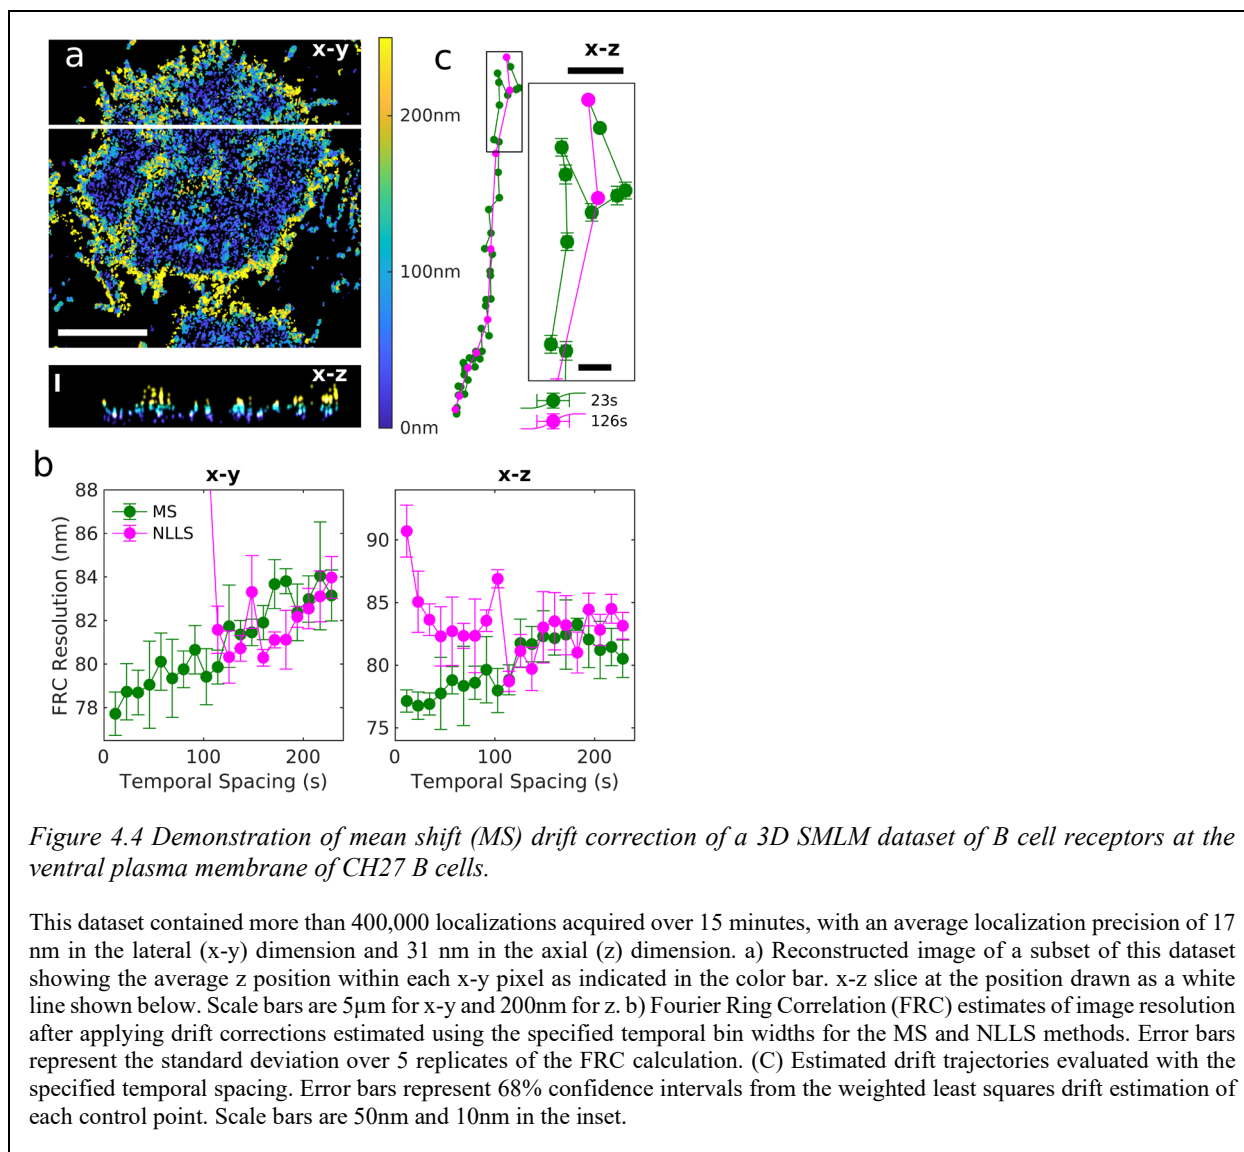


Figure 4.4 applies the MS approach to an experimental SMLM dataset of labeled B cell receptors on the ventral membrane of B cells imaged using a phase mask in the emission path to localize fluorophores in 3D (74). As was the case for simulated datasets, the differences in the performance of the mean shift and NLLS fitting methods are more pronounced than in the 2D dataset of Figure 4.3. Additional diagnostics for the 3D case are shown in Supplemental Figure 4.10.

In summary, a mathematically simple mean shift algorithm modestly outperforms cross-correlation based estimates of drift correction in 2D and more significantly improves the time-resolution of drift-corrections in 3D. The approach is computationally efficient, is robust without sophisticated methods to estimate start-points, and does not require image reconstruction with

memory and pixelation limitations. The metric provided to estimate error and predict robustness directly from data provides users with a means to evaluate the quality of a drift correction within an SMLM analysis pipeline. For the example datasets explored, modest improvements in resolution lead to adjustments of localized molecule positions relevant for evaluating the structure of protein complexes in cells.

4.4 Methods and Materials

4.4.1 Simulated Datasets

An idealized 2D SMLM dataset was simulated as a spatially random set of fluorophores on a 20 μm diameter circular cell, with each fluorophore giving rise to a Poisson-distributed number of localizations with isotropic Gaussian localization error $\sigma_x = \sigma_y = 15/\sqrt{2}$ nm. We define $\sigma = \sqrt{\sigma_x^2 + \sigma_y^2} = 15$ to denote the total root-mean-square localization error. A second dataset was generated from the same fluorophore locations, localization precision, and average number of localizations per fluorophore, and shifted between 0 and 150 nm in a random direction. Simulated datasets were generated over a range of densities (5 to 100 per μm^2) and a range of localizations per molecule (.05 to .2).

An idealized 3D SMLM dataset was simulated in a similar fashion. Fluorophores were distributed uniformly on a cylinder 20 μm in diameter and 2 μm deep. Each fluorophore produces a Poisson-distributed number of localizations with $\sigma_x = \sigma_y = 15/\sqrt{2}$ nm as before, and with $\sigma_z = 30/\sqrt{2}$ nm. One dataset is translated by a random distance between 0 and 150 nm in a random direction in x, y, and z.

4.4.2 Extracting Close Pairs of Coordinates Between Datasets

Consider two point sets $\mathbf{u}_i = (u_{ix}, u_{iy})$ and $\mathbf{v}_j = (v_{jx}, v_{jy})$, for $i = 1, \dots, n_1$ and $j = 1, \dots, n_2$. We wish to quickly determine which pairs (i, j) are closer than some maximum distance r_{max} ; i.e. which pairs satisfy $\|\mathbf{u}_i - \mathbf{v}_j\| < r_{\text{max}}$. The algorithm is adapted from the code for the `closepairs()` and `crosspairs()` functions of the R package `spatstat` (289), and implemented in C with a MATLAB

interface. We first sort each dataset with respect to its x -coordinate, so that $u_{kx} \leq u_{lx}$ whenever $k \leq l$. Then the algorithm proceeds as follows:

1. Let $i = 1$ and $j_{\text{left}} = 1$.
2. Let $x_{\text{left}} = u_{ix} - r_{\text{max}}$. All close pairs of \mathbf{u}_i must satisfy $v_{jx} > x_{\text{left}}$.
3. Increment j_{left} until $v_{j_{\text{left}}x} \geq x_{\text{left}}$.
4. For each $j = j_{\text{left}}, \dots, n_2$, if $v_{jx} - u_{ix} > r_{\text{max}}$, increment i and return to step 2. Otherwise, compute $r_{ij}^2 = (u_{ix} - v_{ix})^2 + (u_{iy} - v_{iy})^2$. If $r_{ij}^2 \leq r_{\text{max}}^2$, add (i, j) to the list of results.

This algorithm avoids computing pairwise distances between most pairs in the dataset, and so is much faster and more memory efficient than a brute force approach. It can be readily adapted to higher dimensions by applying the appropriate n -dimensional distance metric in step 4. For convenience, our implementation returns the displacements $\Delta \mathbf{r}_{ij} = \mathbf{v}_j - \mathbf{u}_i$, and total distance $r_{ij} = \|\Delta \mathbf{r}_{ij}\|$ for each pair (i, j) , instead of the indices themselves.

4.4.3 Determining Shifts between Translated Datasets Using a Mean Shift Algorithm

Given the set of displacements $\Delta \mathbf{r}_{ij} = \mathbf{u}_i - \mathbf{v}_j$ between two point sets \mathbf{u}_i and \mathbf{v}_j , a mean shift clustering algorithm (286–288) can be applied to search for the peak of the displacement density function. Briefly, let $\mathbf{r}_{\text{shift},0}$ be an initial guess to initialize the shift estimate, and δ a radius of consideration to use in the optimization procedure. The algorithm proceeds by iteration, by setting

$$\mathbf{r}_{\text{shift},t+1} = \left\langle \mathbf{r}_{ij} \right\rangle_{\|\mathbf{r}_{ij} - \mathbf{r}_{\text{shift},t}\| \leq \delta},$$

where the average is restricted to the subset of displacements \mathbf{r}_{ij} that satisfy the subscript, i.e. that are within a radius δ from the previous shift estimate $\mathbf{r}_{\text{shift},t}$. The algorithm terminates when the distance $\|\mathbf{r}_{\text{shift},t+1} - \mathbf{r}_{\text{shift},t}\|$ between subsequent shift estimates becomes smaller than machine precision, or when the number of iterations exceeds a user-defined maximum number. δ must be sufficiently large so that the true shift resides within the explored area when centered at the starting-point. In practice, we apply the algorithm twice: first with a large δ to determine the rough shift, and then with a smaller δ , using the first estimate as a starting point, to refine the estimate.

While the above can be applied directly to 3-dimensional data by taking the average over a 3-dimensional ball of radius δ instead of the 2-dimensional disc, we find it is advisable to consider an ellipsoid that is stretched in the z-direction, to account for the larger axial localization errors present in our 3-dimensional simulated and experimental datasets. In the present work, we let the semimajor axis of the ellipsoid be $\sqrt{2}\delta$, in the z direction, and the semiminor axes both δ , so that x-y cross-sections of the regions of consideration are discs.

4.4.4 Estimates of mean shift error

We model the distribution of pairs around the true shift as a Gaussian-distributed peak with standard deviation ζ , centered on the true shift $\mathbf{r}_{\text{shift}}$ on a uniformly distributed background. Assuming $\mathbf{r}_{\text{shift},t}$ is sufficiently close to $\mathbf{r}_{\text{shift}}$ that most of the Gaussian peak falls within the region of consideration, the variance ξ^2 of the two components of $\mathbf{r}_{\text{shift},t}$ is given by

$$\xi^2 = \text{Var}[r_{\text{shift},t,x}] = \text{Var}[r_{\text{shift},t,y}] = \frac{n_{\text{true}}\zeta^2 + n_{\text{false}}\delta^2/4}{(n_{\text{true}} + n_{\text{false}})^2},$$

where δ is the radius of consideration for the MS algorithm, and n_{true} and n_{false} are respectively the number of “true pairs” that are drawn from the Gaussian part of the distribution (displacements between different localizations of the same molecules) and the number of “false pairs” that are drawn from the uniform part (displacements between different molecules), that fall within the region of consideration. Furthermore, the expected value after one more step can be derived:

$$\begin{aligned} \mathbb{E}[\mathbf{r}_{\text{shift},t+1} - \mathbf{r}_{\text{shift},t}] &= \frac{n_{\text{false}}\mathbf{r}_{\text{shift},t} + n_{\text{true}}\mathbf{r}_{\text{shift}}}{n_{\text{true}} + n_{\text{false}}} - \mathbf{r}_{\text{shift},t} \\ &= \frac{n_{\text{true}}}{n_{\text{true}} + n_{\text{false}}}(\mathbf{r}_{\text{shift}} - \mathbf{r}_{\text{shift},t}). \end{aligned}$$

Suppose t is the final step of the algorithm, i.e. $\mathbf{r}_{\text{shift},t+1} - \mathbf{r}_{\text{shift},t} = 0$. Then by hypothesis, $\mathbf{r}_{\text{shift},t}$ deviates from its expected value by

$$\frac{n_{\text{true}}}{n_{\text{true}} + n_{\text{false}}} \|\mathbf{r}_{\text{shift},t} - \mathbf{r}_{\text{shift}}\|.$$

This deviation will typically take on values comparable to the standard deviation ξ shown above. Thus, we estimate the error of the MS algorithm by:

$$\text{Predicted Error} = \frac{\sqrt{n_{\text{true}}\zeta^2 + n_{\text{false}}\delta^2 / 4}}{n_{\text{true}}}$$

This predicted error is to be interpreted as an estimate of the standard deviation of the shift estimate in each direction.

In practice, the parameters n_{true} , n_{false} , and ζ are not known, so we estimate them from data. Specifically, we construct the isotropic cross-correlation function $c(r)$ from the pair separations \mathbf{r}_{ij} , determine the baseline of $c(r)$ from its long-range median value, and use the baseline to infer n_{true} and n_{false} . Finally, we fit $c(r)$ to a Gaussian plus a constant to estimate ζ . This error estimate is derived from a heuristic argument and is not exact. However, its performance is adequate in practice. See Figure 4.2c for a comparison to observed standard deviations of MS shift estimates.

For 3D data, we compute lateral and axial predicted errors separately, by projecting the data from the ellipsoidal region of consideration into the x-y plane or onto the z axis, respectively. n_{true} , n_{false} , and ζ are estimated separately for the lateral and axial directions from the two projections.

4.4.5 Evaluating displacements using nonlinear least squares (NLLS) fitting

Displacements $\mathbf{r}_{\text{shift}}$ between pairs of localization datasets were also estimated by NLLS fitting of a Gaussian to the spatial cross-correlation function of the two datasets. NLLS fitting was accomplished using software published as Supplemental material of (61). Images were first reconstructed from simulated localizations with a pixel size of 15nm for simulated localizations, or from acquired data with a pixel size of 8nm for Nup210 or 15nm for B cell receptor experimental localizations. Cross-correlations are tabulated using 2D Fast Fourier Transforms (FFTs) and then fit a 2D Gaussian function to a subset of the cross-correlation centered at the start-point of the NLLS algorithm. The software from (61) finds the start-point using an elegant smoothing step to reduce noise then uses the largest local maximum of the smoothed cross-correlation as the start-point for fitting.

For localizations acquired in 3D, multiple 2D projections were constructed from 3D localizations, then the procedures described for 2D images were applied to determine displacements. First, images projecting on the lateral dimension (x-y plane) were generated and the lateral displacement was determined. To compute the z displacement, both the xz and yz projections were used, and the final z displacement was the average determined from the two projections.

4.4.6 Correcting continuous drift

Continuous drift was corrected by temporally dividing the data into N bins, each having the same number of frames. For each of the $N(N-1)/2$ pairs (m, n) of temporal bins, the mean shift or NLLS algorithm is applied to estimate the shift $\mathbf{r}_{\text{shift}, m \rightarrow n}$ from temporal bin m to n , corresponding to the drift between the bins. Drift at each of the N time points is calculated from the $N(N-1)/2$ pairwise shifts using a least-squares minimization algorithm (61); this takes advantage of the overdetermined nature of the drift calculation to improve the precision of the measurement. Outlier shifts, whose residual with respect to the least-squares estimate exceeds a user-defined threshold, can also be discarded as described in (61). These shifts typically correspond to “failures” of the shift estimation method. The final drift curve at each frame is determined by linear interpolation and extrapolation from the N basis points.

4.4.7 Evaluating performance of displacement algorithms

For simulated localizations, errors away from known displacements were tabulated for each simulated configuration. The 2D precision of each method is defined as the standard deviation of a centered, isotropic 2D Gaussian fit to the central peak of the histogram of these values, considering only values that fall within twice the localization precision (2σ) used in the simulation.

The fit is applied directly to the absolute errors $\|\mathbf{r}_{\text{shift,est.}} - \mathbf{r}_{\text{shift}}\|$ with the distribution function

$f(r) = \frac{r}{\sigma^2} e^{-r^2/2\sigma^2}$. Similarly, in the 3D case, x- and z-precision are evaluated separately by fitting

1D Gaussian functions to the x- and z-errors, respectively. Values that fall outside of the 2σ window are reported as failures of the algorithm, and contribute to the failure rates reported in figures. Since failures can return values with large errors, they can have an outsized impact on simpler precision metrics, such as the root mean square error (RMSE). Computation time was

assessed in MATLAB using the built-in tic and toc functions. For simulated data, computation time was averaged over 500 simulations for each condition. For experimental data, computation time includes the $N(N-1)/2$ shift estimates and the error estimates for each shift estimate in the MS case. Normalized residual degrees of freedom (Normalized DOF) of the linear least squares algorithm are calculated by the ratio of shifts that are used for the final least squares minimization step (after removal of outliers) to the number of time points at which drift is estimated (i.e. $N-1$). This serves as a diagnostic for how much redundancy is included in the linear least squares minimization step.

4.4.8 Evaluating the resolution of drift-corrected datasets

Resolutions of the final reconstructed images were compared using Fourier Ring Correlation (80, 81). Specifically, we used code adapted from the supplementary software of (81). To compute the x-y resolution, nearby localizations belonging to adjacent camera frames were grouped together, with the position taken to be the average of the relevant coordinates. The FRC curves were produced by dividing the dataset into blocks of 500 frames and allocating an equal number of blocks randomly to each of the two sets. The pixel size was taken to be 5nm. For the B cell dataset, the Fourier Ring approach was applied to the xy and xz projections in turn, also using a pixel size of 5nm in each case.

4.4.9 Preparation of cellular samples for imaging

Mouse primary neurons were isolated from P0 mouse pups that were decapitated and brains were isolated into ice cold, filtered dissection buffer (6.85 mM sodium chloride, 0.27mM potassium chloride, 0.0085 mM sodium phosphate dibasic anhydrous, 0.011mM potassium phosphate monobasic anhydrous, 33.3 mM D-glucose, 43.8 mM sucrose, 0.277 mM HEPES, pH 7.4) as described in (290). After removing the cerebellum and the meninges, cortices were dissected out, placed into a microcentrifuge tube, and cut into small pieces with dissection forceps. Cortices were incubated in 50 μ L papain (2mg/mL; BrainBits) and 10 μ L DNase I (1mg/mL; Worthington Biochemical) for 30min at 37 °C. 500 μ L BrainPhys Neuronal Medium (Stemcell Technologies) and 10 μ L additional DNase I were added, and cortices were triturated using P1000 and P200 pipet tips. Triturated cortices were centrifuged at 1000rpm for 5min. After discarding the supernatants, the pellets were triturated and centrifuged three more times until the supernatant

remained clear and neuronal pellets were visible. Pelleted neurons were resuspended in BrainPhys Neuronal medium with SM1 supplement as previously described (291), then plated onto 35mm #1.5 glass-bottom dishes (MatTek Life Sciences) coated with polyethylenimine (100 µg/ml; Polysciences). Neurons were incubated in 5% CO₂ at 37 °C, and 1mL of media was replaced every four days.

On day 10 of culture (days *in vitro* 10), neurons were rinsed with sterile Hank's Balanced Salt Solution, then incubated for 1min with pre-warmed 2% PFA (Electron Microscopy Sciences) in Phosphate Buffered Saline (PBS). The neurons were then incubated in 0.4% Triton X-100 (Millipore Sigma) in PBS for 3min, and fixed for 30min with 2% PFA in PBS. Neurons were then washed with PBS five times, incubated in blocking buffer containing 5% Normal Donkey Serum and 5% Normal Goat Serum (Jackson Laboratories) for 30min, then labeled with Nup210 polyclonal antibody diluted in blocking buffer (1:200; Bethyl laboratories A301-795A) overnight in 4 °C. The following day, neurons were washed three times in PBS then stained with Goat-anti-rabbit Alexafluor 647 secondary antibody (1:1000; ThermoFisher) for an hour, washed three times with PBS, then imaged.

CH27 B cells (292) were cultured, allowed to adhere to 35mm #1.5 glass-bottom dishes (MatTek Life Sciences) overnight, then incubated in Alexa647 conjugated fAb prior to fixation in 4% PFA and 0.1% glutaraldehyde (Electron Microscopy Sciences), as described previously (83). The labeled fAb antibody was prepared by conjugating an Alexa647 NHS ester (ThermoFisher) to an unconjugated fAb (Goat Anti-Mouse IgM, µ chain specific; Jackson Immunoresearch) using established protocols (83).

4.4.10 Single molecule imaging and localization

Imaging was performed using an Olympus IX83-XDC inverted microscope. TIRF laser angles were achieved using a 60X UAPO TIRF objective (NA = 1.49), and active Z-drift correction (ZDC) (Olympus America) as described previously. The ZDC was not used for collection of 3D datasets. Alexa 647 was excited using a 647 nm solid state laser (OBIS, 150 mW, Coherent) coupled in free-space through the back aperture of the microscope. Fluorescence emission was detected on an EMCCD camera (Ultra 897, Andor) after passing through a 2x expander. Imaging in 3D was accomplished using a SPINDLE module equipped with a DH-1 phase mask (DoubleHelix LLC).

Single molecule positions were localized in individual image frames using custom software written in Matlab. Peaks were segmented using a standard wavelet algorithm (55) and segmented peaks were then fit on GPUs using previously described algorithms for 2D (58) or 3D localizations (293). After localization, points were culled to remove outliers prior to drift correction. Images were rendered by generating 2D histograms from localizations followed by convolution with a Gaussian for display purposes.

4.5 Supplemental Figures

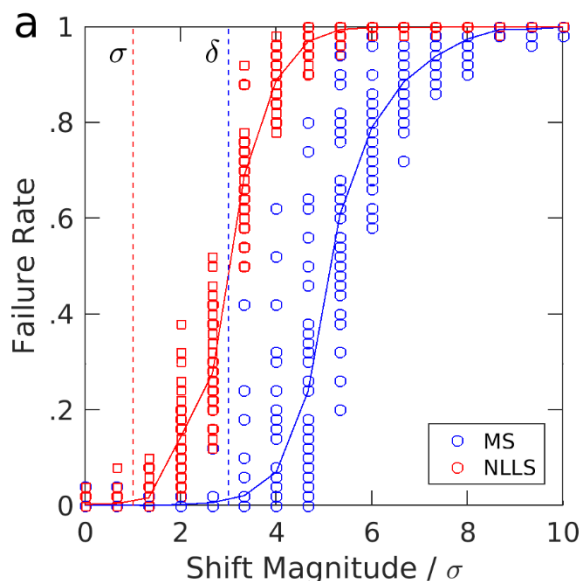


Figure 4.5 Evaluating the mean shift (MS) and NLLS algorithms with the start-point at the origin.

Simulations and shift determination approaches are described in Methods. **a)** Shifts between 0 and 10 times the localization precision (σ) are applied in a random direction, and the MS and NLLS algorithms are applied to determine these displacements. The observation window (δ) for the MS algorithm has an extent of 3σ as indicated by the dashed blue line. The “failure rate” is the fraction of simulations whose error exceeds 2σ . Each point represents a given combination of fluorophore density and number of localizations per molecule, averaged over 50 independent trials. Densities range from 5 to 100 molecules per μm^2 , and localizations per molecule range from .05 to .2.

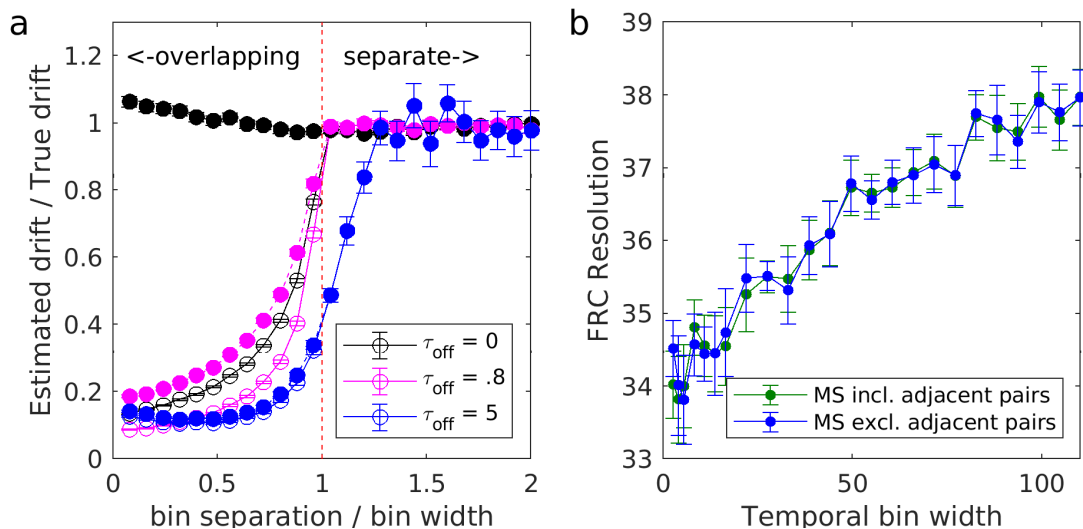


Figure 4.6 Drift estimates between overlapping temporal bins are prone to bias.

a) Mean shift (MS) shift estimates for pairs of bins that are overlapping or separated by short times. Simulations are similar to those for Figure 4.2 with a cell of radius $10 \mu\text{m}$ and randomly distributed fluorophores at a density of $20 \text{ per } \mu\text{m}^2$, but with explicit blinking kinetics, modeled as a simple two state (fluorescent/dark) system, with activation and deactivation time constants τ_{on} and τ_{off} for the dark \rightarrow fluorescent and fluorescent \rightarrow dark transitions, respectively. A constant drift rate of $.1 \text{ nm per frame}$ is applied. The drift estimates shown here are for 50 frame temporal bins, with bin starts separated by the bin separation times as shown. For each τ_{off} , τ_{on} is adjusted so that the average number of localizations per fluorophore is approximately 0.15. Open circles represent drift estimates using all pairs of simulated points, including the trivial 0 displacements between points that appear in the overlap of the two temporal bins. Filled circles represent drift estimates using all pairs of distinct point, i.e. excluding the trivial 0 displacements between points that appear in the overlap of the two temporal bins. Note that overlapping bins are subject to substantial bias even for quite short τ_{off} (magenta points), and that even non-overlapping bins may be subject to bias when τ_{off} is long (blue points). b) FRC resolutions for the nuclear pore complex data of Figure 4.3, with MS drift corrections including or excluding drift estimates for adjacent pairs of temporal bins in the linear least squares fit.

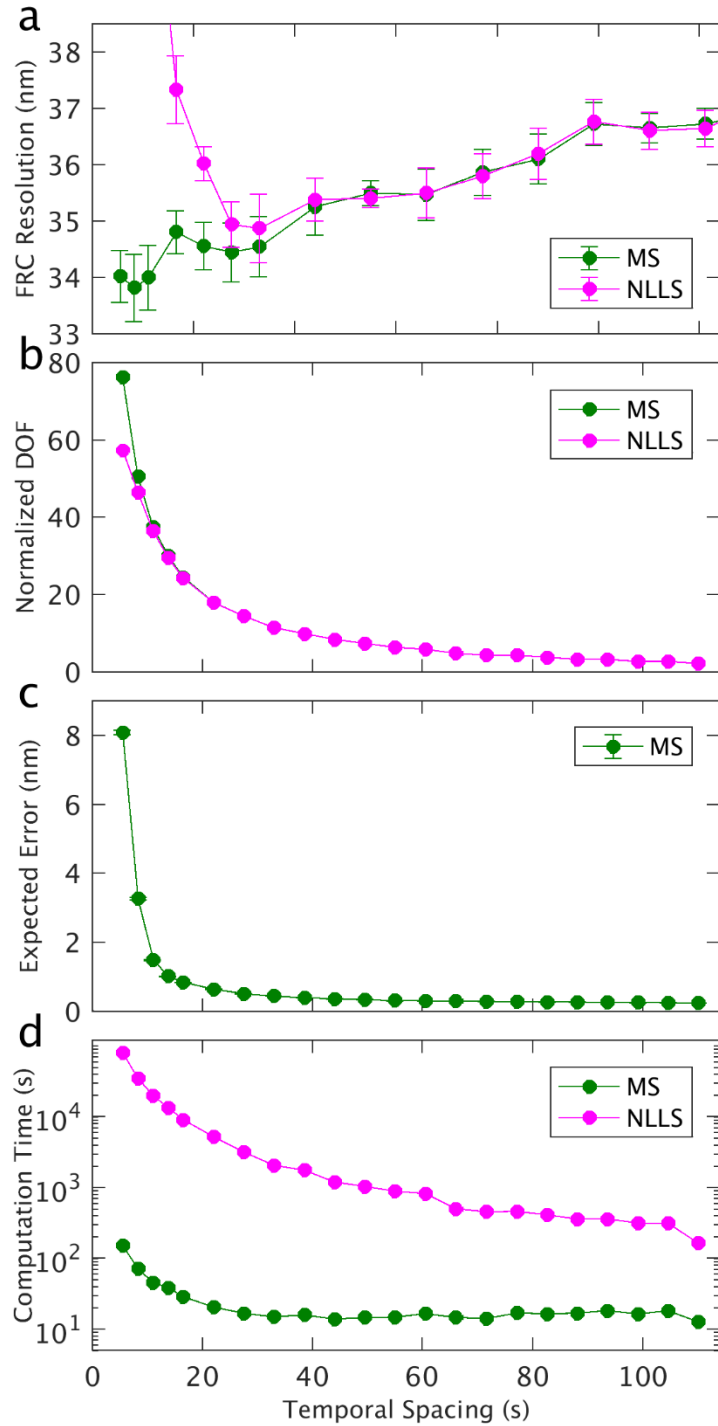


Figure 4.7 Drift correction diagnostics for the nuclear pore complex dataset of Figure 4.3.

a) FRC resolutions. Error bars are given by the standard deviation over 20 trials. b) The number of degrees of freedom (DOF) after removing outliers (normalized by the number of parameters) for the redundant least square minimization calculation. c) RMSE of the expected errors for the mean shift method. Error bars are given by the standard error of the mean. d) Total computation time.

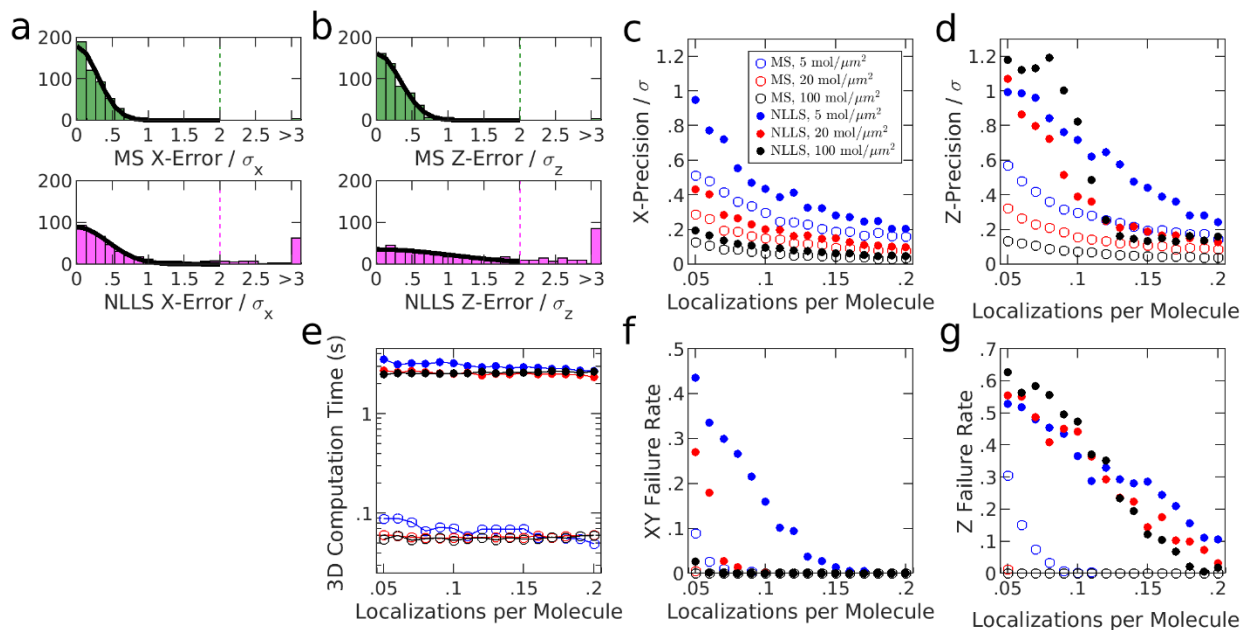


Figure 4.8 Evaluating the mean shift (MS) algorithm on 3D simulated data, compared to the NLLS approach.

Simulations and shift determination approaches are described in methods. a,b) Histograms of x-errors (a) and z-errors (b) for the MS and NLLS approaches for the 0.01 molecules/ μm^3 and 0.05 localizations per molecule condition. The precision of each method is evaluated for each condition as the standard deviation of a Gaussian fit to the central peak of the histogram. The “failure rate” is the fraction of simulations whose error exceeds twice the localization precision σ , indicated as a dashed line. Three densities are shown in the plots: 5, 20, and 100 molecules/ μm^2 . c,d) Comparison of the lateral (c) and axial (d) precision of each approach, plotted versus the number of localizations per molecule. e) Computation times of the two approaches over the same conditions shown in c,d). f,g) Comparison of the failure rate of each approach in lateral (f) and axial (g) directions under the same range of conditions.

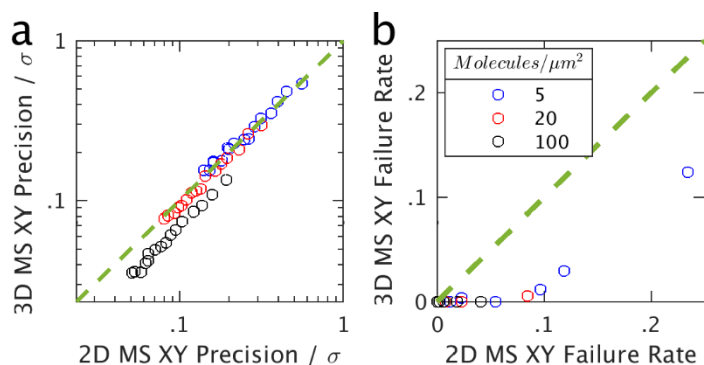


Figure 4.9 2D projections of 3D data degrade mean shift (MS) shift estimation performance.

Lateral (x-y) precision (a) and failure rate (b) when MS shift is determined in 3D or in 2D after projecting the localizations into the x-y plane. The points shown each summarize 500 replicates of one simulation condition, with fluorophore density as shown in the legend, and localizations per molecule ranging from .05 to .2. The simulated 3D data used here is constructed such that its 2D projection is identical to the 2D data used for Figure 4.2.

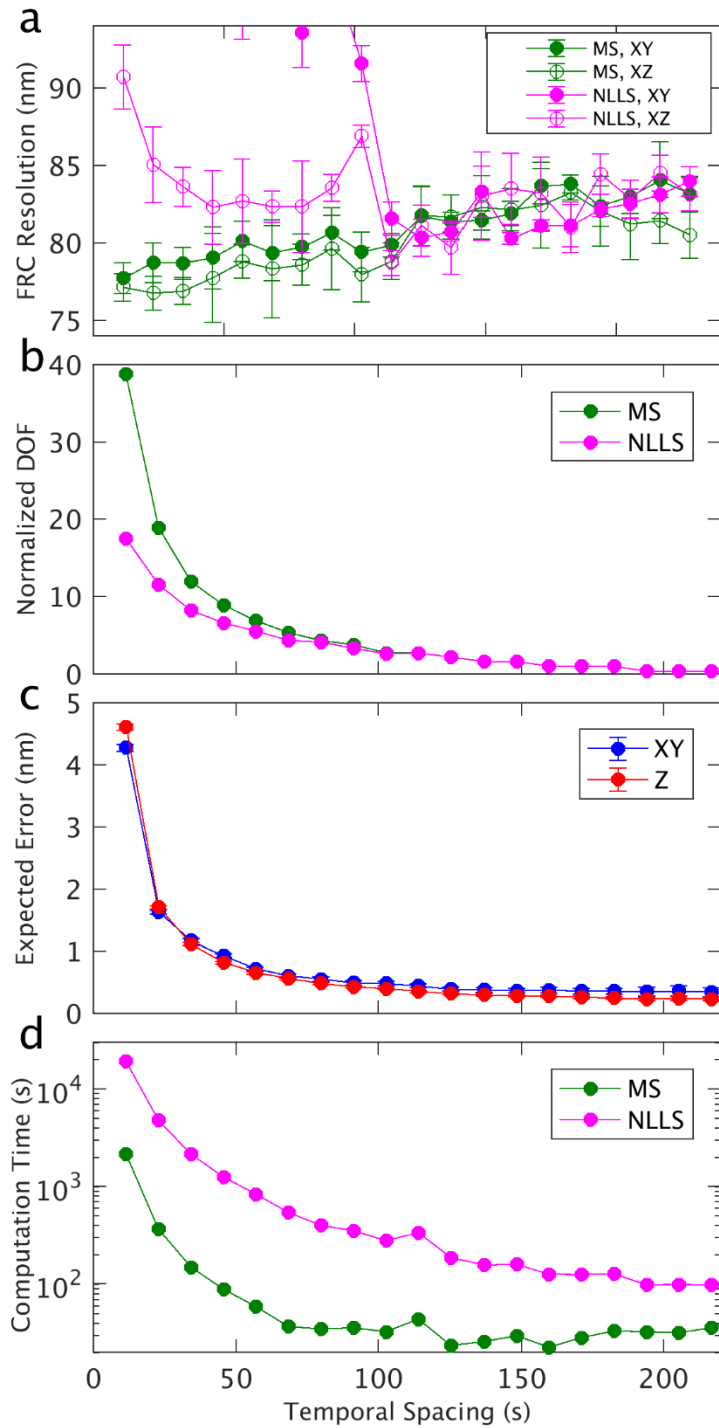


Figure 4.10 Drift correction diagnostics for the 3D B cell dataset of Figure 4.4.

a) FRC resolutions. Error bars are given by the standard deviation over five trials. b) The number of degrees of freedom (DOF) after removing outliers (normalized by the number of parameters) for the least square minimization calculation. c) RMSE of the lateral (x-y) and axial (z) expected errors for the mean shift calculation of pairwise shifts. Error bars are given by the standard error of the mean. d) Total computation time.

Chapter 5 Estimating the Localization Spread Function of Static Single Molecule Localization Microscopy Images

Single molecule localization microscopy drastically reduces the position uncertainty of observed molecules in biological samples. To be confident in the interpretation of detailed features of these images, one must answer the question: how much uncertainty remains? In this chapter, I propose a method that takes advantage of the temporal correlations present in most SMLM techniques to statistically separate the distribution of molecules in the sample from the distribution of localizations around each molecule. This chapter is a collaborative work. The ideas, derivations of the validity and limitations, analysis code and text were mostly developed by myself and Sarah L Veatch, with contributions from Frank J Fazekas. Sumin Kim, Jennifer C Flanagan-Natoli and Emily Sumrall provided the biological samples and in some cases the SMLM datasets that are used for demonstration of the method. A version of this chapter has been published in *Biophysical Journal* (294).

5.1 Abstract

Single molecule localization microscopy (SMLM) permits the visualization of cellular structures an order of magnitude smaller than the diffraction limit of visible light, and an accurate, objective evaluation of the resolution of an SMLM dataset is an essential aspect of the image processing and analysis pipeline. Here we present a simple method to estimate the localization spread function (LSF) of a static SMLM dataset directly from acquired localizations, exploiting the correlated dynamics of individual emitters and properties of the pair autocorrelation function evaluated in both time and space. The method is demonstrated on simulated localizations, DNA origami rulers, and cellular structures labelled by dye-conjugated antibodies, DNA-PAINT, or fluorescent fusion proteins. We show that experimentally obtained images have LSFs that are broader than expected from the localization precision alone, due to additional uncertainty accrued when localizing molecules imaged over time.

5.2 Introduction

Single molecule localization microscopy (SMLM) is a powerful tool to image structures in cells with dimensions ranging between tens of nanometers to tens of microns. Methods such as (d)STORM (38, 40), (F)PALM (37, 39), and PAINT (47) exploit the stochastic blinking of single fluorophores to localize emitting molecules with a localization precision much smaller than the diffraction limit of visible light, by imaging only a small subset of probes in any given image frame. These samples are then imaged over time, and acquired localizations are typically assembled into a single reconstructed super-resolved image.

Assessing the quality of reconstructed images can be challenging as numerous factors can contribute. These factors can include the labeling density, the types of structures being imaged, the brightness and blinking dynamics of the fluorophore, the finite size of labeling antibodies, motions of the stage or labeled molecules during acquisition, and the analytical methods used in post-processing. One important measure of the quality of a measurement is the localization precision of single fluorophores, which is influenced by many of the factors listed above. Many localization algorithms directly return estimates of the localization precision of single fits, and similar information can be extracted directly from the localizations themselves through the use of pair-correlation functions or nearest neighbor analyses that extract the distribution of positions of molecules detected in adjacent frames (82, 87). However these methods are not sensitive to errors introduced on time-scales longer than a few image frames. Other metrics of image quality have been developed that integrate both precision and spatial sampling. One widely used method, called Fourier Ring Correlation (FRC) (80, 81), effectively captures the impact of factors that degrade quality over the entire span of image acquisition. The FRC curve depends on the types of structures imaged, how well they are sampled, and the specific regions of interest used. The resolution value it reports indicates the length-scale below which the signal to noise ratio falls below a specified limit. This value is useful for comparing imaging conditions for a particular sample, but hard to compare across sample types, making its use highly context dependent.

Here we present a simple method to estimate the average localization spread function (LSF) of a SMLM dataset as a way to evaluate its resolution. The LSF is analogous to the point spread function (PSF) of a conventional microscopy measurement and can aid the interpretation of acquired images. The LSF reports on how accurately distances can be measured between labeled objects within images, and can be used when constraining the structure of a multi-protein complex

or when estimating the statistical co-distribution of labeled components. The LSF can also be used to evaluate optimizations all along the image acquisition and processing pipelines. The method presented exploits temporal correlations in the blinking dynamics of single fluorophores commonly used for localization microscopy (295–297). This enables the method to report on errors accumulated over time that do not typically impact the accuracy of single molecule fitting but impact how accurately a molecule’s position is determined relative to others (298, 299). Here, we derive a method to isolate the LSF directly from acquired localizations, validate it through simulation and images of DNA origami rulers, and apply it to several images of labeled structures in cells. In all of the experimental examples interrogated, the LSF width increases over time-scales at and below that of drift correction.

5.3 Materials and Methods

5.3.1 Simulations

Simulations mimicking DNA origami rods were accomplished by randomly placing pairs of fluorophores positioned 50 nm apart within a 40 μm by 40 μm region of interest at an average density of 1 pair per μm^2 . 20,000 individual image frames with an effective frame time of 0.1 sec were simulated by sampling a subset of molecular positions with a localization precision of 10 nm in each lateral dimension. The dynamics of individual fluorophores were governed by a continuous time Markov process involving five states: one on state (1), three dark states (0, 01, 02), and an irreversible bleached state (B), following the procedure described previously (297, 300). The on state was accessible from any of the dark states, while dark state 0 was accessible only from the on state, and dark states 01 and 02 were accessible only from the previous dark states, 0 and 01, respectively. We used the following parameters (using the notation described in (297, 300)) to capture essential elements of our experimental observations: $\lambda(0 \rightarrow 1) = 1.2$ Hz; $\lambda(0 \rightarrow 01) = 0.05$ Hz; $\lambda(01 \rightarrow 02) = 0.0033$ Hz; $\lambda(01 \rightarrow 1) = 0.02$ Hz; $\lambda(02 \rightarrow 1) = 0.0005$ Hz; $\lambda(1 \rightarrow 0) = 5$ Hz; $\mu(1 \rightarrow B) = 0.05$ Hz. The continuous time Markov process was simulated in MATLAB (The MathWorks) using File Exchange code “simCTMC.m” (301). When present, drift was applied to all molecular positions with a constant rate of 0.3 nm/sec in the x direction along with diffusive drift characterized by a diffusion coefficient of $D=2.5$ nm²/sec. Drift was corrected using the mean shift algorithm described previously (271) using 1000 frames per alignment (10 s).

5.3.2 Experimental sample preparation:

DNA origami “gatta-STORM” nanorulers were purchased from Gattaquant GMBH (Grafelfing, Germany) and a sample was prepared following manufacturer’s instructions. Briefly, biotinylated Bovine Serum Albumin (biotin-BSA; ThermoFisher; 1 mg/ml) was absorbed to a clean 35 mm #1.5 glass-bottom dish (MatTek well; MatTek Life Sciences) for 5min then washed. Streptavidin was then applied (1 mg/ml) for 5min, then washed with a solution of phosphate buffered saline (PBS) plus 10 mM MgCl₂. A solution containing the biotinylated DNA origami was then applied. Samples were then washed and imaged in an imaging buffer supplemented with 10 mM MgCl₂. “gatta-PAINT” 80RG nanorulers in a sealed sample chamber were purchased from Gattaquant GMBH and imaged in the Atto655 color channel following manufacturer recommendations.

Mouse primary neurons were isolated from P0 mouse pups as described previously and cultured on MatTek wells (271). On day 10 of culture (days *in vitro* 10), neurons were rinsed with sterile Hank’s Balanced Salt Solution and fixed for 10min with pre-warmed 4% PFA (Electron Microscopy Sciences) in Phosphate Buffered Saline (PBS). The fixed neurons were rinsed three times with PBS and permeabilized in 0.2% Triton X-100 (Millipore Sigma) in PBS for 5min. Neurons were then incubated in blocking buffer containing 5% BSA for 30 min, and labeled with Nup210 polyclonal antibody diluted in PBS (1:200; Bethyl laboratories A301-795A) overnight in 4 °C. The following day, neurons were washed three times in PBS and stained with goat-anti-rabbit AlexaFluor 647 Fab Fragment (1:800; Jackson ImmunoResearch 111-607-003) for an hour, washed three times with PBS, then imaged.

CH27 B-cells (mouse, Millipore Cat# SCC115, RRID:CVCL_7178), a lymphoma-derived cell line (292) were acquired from Neetu Gupta (Cleveland Clinic). CH27 Cells were maintained in culture as previously described (53). Cells were adhered to MatTek wells coated with VCAM following procedures described previously (302). Briefly 0.1 mg/ml IgG, Fc γ -specific was adsorbed to a plasma cleaned well for 30 min at room temperature. Wells were rinsed with PBS, then nonspecific binding was blocked with 2% BSA at room temperature for 10 minutes, followed by incubation with 0.01 mg/mL recombinant human VCAM-1/CD106 Fc chimera protein (R&D Systems) and 0.01 mg/mL ChromPure Human IgG, Fc fragment (Jackson Immunoresearch) for 1 hour at room temperature or overnight at 4°C. VCAM-1 coated dishes were stored up to 1 week in VCAM-1 and Fc at 4°C. Immediately prior to plating, dishes were blocked at room temperature

in 2% goat serum (Gibco) for 10 min, then cells were allowed to adhere for 15 min in media prior to chemical fixation in 2% PFA and 0.2% glutaraldehyde (Electron Microscopy Sciences). F-Actin was stained by permeabilizing cells with 0.1% Triton-X-100 prior to incubation with 3.3 μ M phalloidin-AlexaFluor647 (Invitrogen) for at least 15 min. Phalloidin stained cells were imaged immediately after removing label. Cells transiently expressing Clathrin-GFP were permeabilized after fixation with 0.1% Triton-X-100 followed by labeling with a single domain anti GFP antibody (MASSIVE-TAG-Q ANTI-GFP) from Massive Photonics GMBH (Grafelfing, Germany) for 1 h at room temperature, then imaged in 0.5 nM of imaging strand in the imaging buffer supplied by the manufacturer.

Cells expressing the membrane label Src15-mEos3.2 were prepared by transiently transfecting 10^6 cells with a 1 μ g of plasmid encoding Src15-mEos3.2 (N'-MGSSKSKPKDPSQRRNNNNGPVAT-[mEos3.2]-C') which was derived from a GFP tagged version by replacing GFP with mEos3.2 (303, 304). Transfection was accomplished by Lonza Nucleofector electroporation (Lonza, Basel, Switzerland) with program CA-137 and cells were grown in flasks overnight prior to plating and fixation as described above.

5.3.3 Single molecule imaging and localization

Imaging was performed using an Olympus IX83-XDC inverted microscope. TIRF laser angles were achieved using a 100X UAPO TIRF objective (NA = 1.50), and active Z-drift correction (ZDC) (Olympus America). AlexaFluor647 was excited using a 647 nm solid state laser (OBIS, 150 mW, Coherent) and mEos3.2 was excited using a 561 nm solid state laser (Sapphire 561 LP, Coherent), both coupled in free-space through the back aperture of the microscope. Fluorescence emission was detected on an EMCCD camera (Ultra 897, Andor). Samples containing AlexaFluor647 were imaged in a buffer containing 100 mM Tris, 10 mM NaCl, 550 mM glucose, 1% (v/v) beta-mercaptoethanol, 500 μ g/ml glucose oxidase (Sigma) and 40 μ g/ml catalase (Sigma), with 10 mM $MgCl_2$ for the DNA origami sample. Samples with mEos3.2 or DNA PAINT Atto655 probes were imaged in imaging buffer from Massive Photonics GMBH. Single molecule positions were localized in individual image frames using custom software written in MATLAB. In most cases, peaks were segmented and fit from background corrected images, where the background was estimated as the median signal over 500 acquisition frames. Peaks were segmented using a standard wavelet algorithm (55) and segmented peaks were

then fit as single emitters on GPUs using previously described algorithms for 2D (58), or as multi-emitters on a CPU using the ThunderSTORM ImageJ plugin (305). After localization, points were culled to remove outliers prior to drift correction (271). Images were rendered by generating 2D histograms from localizations followed by convolution with a Gaussian for display purposes. Rendering parameters are included in captions and typically images showing larger regions are reconstructed with large pixels and Gaussian filters (10-50 nm) while small regions are rendered with small pixels and Gaussian filters (1 nm and 4-10 nm respectively). For the nano-ruler samples, localizations were assigned to single fluorophores using a home-built implementation of DBSCAN (306), with $\varepsilon = 12$ nm and $\text{minPts} = 15$.

5.3.4 Evaluation of space-time autocorrelations.

Space-time autocorrelations were tabulated by first tabulating space- and time-displacements between all pairs of localizations within a specified region of interest (ROI) detected in a given dataset. This was accomplished using a crosspairs function based on the one from the R package spatstat (289), but used here as a C routine with a MATLAB interface, as described previously (271). Lists of displacements were converted into space-time autocorrelation functions by binning in both time and space within the C routine for improved performance, followed by a normalization implemented in MATLAB that produces a value $g(r, \tau) = 1$ when localizations are randomly distributed in both space and time within the specified ROI. A derivation of the form of this normalization and an explanation of how it is computed are presented in Supplementary Note 1 (Section 5.7).

5.3.5 Estimation of $g_{LSF}(r, \tau)$ and $\sigma_{xy}(\tau)$

The core computations of the effective PSF estimation are gathered in a single MATLAB function. First, $g(r_i, \tau_j)$ is computed as described above, for a range of distance and time separation values. By default, $r_i, i = 1, \dots, N_r$ represent bins with bin edges from 0 to 250 nm with equal spacing of 5 nm, resulting in bin centers ranging between 2.5 to 247.5 nm, and $\tau_j, j = 1, \dots, N_\tau$ represent bins with edges that are log-spaced. The lower edge of the final time-separation bin is determined by identifying the lowest τ that satisfies $g(r < 25 \text{ nm}, \tau)/g(r < 25 \text{ nm}, \frac{3}{4}T < \tau \leq T) \leq 1.5$, where T is the total duration of image acquisition. The τ_{max} reported in figures is the bin center of this final time-separation bin.

Then, $\Delta g(r, \tau_j) = g(r, \tau_j) - g(r, \tau_{max})$ are computed for each τ , normalized by their first spatial points ($r < 5$ nm), and fitted to a Gaussian of the form $A \exp(-r^2/4\sigma_{xy,j}^2)$, using MATLAB's nonlinear least squares fitting routine. $\sigma_{xy,j}$ is reported as the estimate of $\sigma_{xy}(\tau_j)$. Bootstrapped standard errors are determined by choosing eight subsamples of the points, each containing one quarter as many points as the full dataset, and estimating $\sigma_{xy}(\tau_j)_k$ for each subsample k , in the same way as for the full dataset. The standard error is reported as $\frac{1}{2}$ std. dev. ($\sigma_{xy}(\tau_j)_k$), where the $\frac{1}{2}$ accounts for the overestimate of errors due to using fewer points by a factor of four.

5.3.6 Estimating $g_{LSF}(r, \tau)$ by grouping localizations with molecules

In simulations and DNA origami samples, localizations imaged at x_i, y_i, τ_i are associated with the molecules that produced them. For the nano-ruler samples, localizations were assigned to single fluorophores using a home-built implementation of DBSCAN (306), with $\epsilon = 12$ nm and $\text{minPts} = 15$.

We tabulate displacements between all pairs associated with the same molecule $\Delta r_{ij} = \sqrt{\Delta x_{ij}^2 + \Delta y_{ij}^2}$ and Δt_{ij} . The list of all pairs is binned into two dimensional histograms following the same r and τ bin-edges as described for computing $g(r, \tau)$ above and are normalized by the number of pairs contributing to each bin. Distributions at each τ bin are fit to the same Gaussian form as applied to the $g_{LSF}(r, \tau)$ estimated from $\Delta g(r, \tau)$.

5.3.7 Measuring distances between distinct molecules on the same ruler.

In simulations, all localizations were associated with the molecules and rulers that produced them. In DNA origami samples, DBSCAN segmented molecules on the same ruler were identified as segments whose average localization position was within 10 nm of the expected displacement between probes on rulers specified by the manufacturer (40 – 60 nm for Figure 5.3 and 70 - 90 nm in Figure 5.4). We then tabulate displacements between all pairs of localizations associated with different molecules on the same ruler $\Delta r_{ij} = \sqrt{\Delta x_{ij}^2 + \Delta y_{ij}^2}$.

5.3.8 Determining the resolution with Fourier Ring Correlation

The resolution of each dataset was assessed with Fourier Ring Correlation (FRC) (81). To produce the FRC curves, localizations were divided into consecutive blocks of 500 frames, and these blocks were randomly placed into one of two statistically independent subsets. For the simulated and experimental DNA origami datasets, as well as the nuclear pore complex dataset, the pixel size for the FRC calculation was taken to be 5 nm, and square regions 10 μm on a side were used as a mask. For the actin and Src15 datasets, the pixel size was 10 nm, with the mask 20 μm on each side. 20 randomly determined repetitions of the calculation were performed for each dataset.

5.3.9 Determining localization precision using Nearest Neighbor distributions

Nearest Neighbors were identified in adjacent image frames using the crosspairs algorithm using a time-interval of 1 frame and a distance cutoff of 100 nm. The closest nearest neighbor was identified for each molecule and included in the distribution $P(r_{NN})$. This distribution was then fit using the MATLAB function fit to extract the localization precision (σ_{NN}) using the functional form (82):

$$P(r_{NN}) = \frac{Ar_{NN}}{2\sigma_{NN}^2} \exp\left(-\frac{r_{NN}^2}{4\sigma_{NN}^2}\right) + \frac{B}{\sqrt{2\pi W^2}} \exp\left(-\frac{(r_{NN} - r_o)^2}{2W^2}\right) + C.$$

The second and third term correct for pairs of localizations not originating from the same molecule.

5.4 Results

5.4.1 Derivation of the estimated LSF

The spatial autocorrelation function describing a distribution of static molecules is given by $g_{molecules}(r)$ and is tabulated as described in Methods. This function can be divided into two components:

$$g_{molecules}(\vec{r}) = \frac{1}{\rho} \delta(\vec{r}) + g_p(\vec{r}). \quad (1)$$

The first term in Eqn. 1 comes from counting single emitters and is a Dirac delta function ($\delta(\vec{r})$) with magnitude equal to the inverse average density of molecules (ρ) over the region of interest (ROI). The second term in Eqn. 1 comes from correlations between distinct pairs of molecules and reports on the sample-dependent detailed structure present in the image. In the special case of complete spatial randomness, $g_p(\vec{r}) = 1$. In SMLM, single emitters labeling molecules have dynamics governed by the probe photophysics, which can be described with the

temporal autocorrelation function $g_e(\tau)$. Probes can remain on for multiple sequential image frames and can blink on again at a later time before eventually bleaching irreversibly (295–297). As a result, $g_e(\tau)$ is highly correlated (>1) at short time-intervals and decays sharply on time-scales describing the average on-time of fluorophores. This function continues to decay slowly at long τ , both because some probes tend to flicker over medium to long time-scales and because some fluorophores eventually bleach. Including $g_e(\tau)$ produces the following spatio-temporal autocorrelation function for the emitting molecules:

$$g_{emitters}(\vec{r}, \tau) = \frac{1}{\rho} \delta(\vec{r}) g_e(\tau) + g_p(\vec{r}). \quad (2)$$

In other words, the central peak due to the same fluorophore being “on” at different times inherits the dynamics of the probe, while the contributions from pairs of different molecules remains time-independent. Eqn. 2 assumes the blinking statistics of fluorophores labeling different molecules are uncorrelated, which is why $g_e(\tau)$ multiplies only the first term.

When fluorophores are localized with finite spatial resolution, the distribution of localization errors can be described as a probability density function that characterizes the resolution of the image. This distribution induces a characteristic blurring of the true locations of the molecules, just as a conventional microscope can be thought of as convolving a true image with a point spread function (PSF). By analogy, we instead call this distribution the localization spread function or LSF. The autocorrelation function of localizations, $g(\vec{r}, \tau)$, is the autocorrelation of the emitters, $g_{emitters}(\vec{r}, \tau)$, blurred (convolved) by the autocorrelation of the LSF, or $g_{LSF}(\vec{r}, \tau)$. Including this factor, $g(\vec{r}, \tau)$ becomes:

$$g(\vec{r}, \tau) = \frac{1}{\rho} g_{LSF}(\vec{r}, \tau) g_e(\tau) + g_{LSF}(\vec{r}, \tau) * g_p(\vec{r}), \quad (3)$$

where $*$ indicates a convolution. The first term in Eqn. 3 describes multiple observations of the same molecule and is exactly proportional to the LSF at time-interval τ .

The goal of subsequent steps of this derivation is to isolate $g_{LSF}(\vec{r}, \tau)$ from the first term of Eqn. 3 by comparing $g(\vec{r}, \tau)$ tabulated from pairs of localizations acquired at different time-intervals τ . In particular, we choose a long time-interval τ_{max} and consider differences:

$$\Delta g(\vec{r}, \tau) = g(\vec{r}, \tau) - g(\vec{r}, \tau_{max}). \quad (4)$$

First consider the simple case where the LSF is independent of τ , $g_{LSF}(\vec{r}, \tau) = g_{LSF}(\vec{r})$. In this limit, the second term of Eqn. 3 is independent of τ , so $\Delta g(\vec{r}, \tau)$ is exactly proportional to $g_{LSF}(\vec{r})$:

$$\Delta g(\vec{r}, \tau) = \frac{1}{\rho} (g_e(\tau) - g_e(\tau_{max})) g_{LSF}(\vec{r}) \propto g_{LSF}(\vec{r}, \tau). \quad (5)$$

Thus, the difference $\Delta g(\vec{r}, \tau)$ can be taken as a direct measurement of $g_{LSF}(\vec{r}, \tau)$.

In practice, the above assumption does not hold exactly, so we must consider the effects of time-varying $g_{LSF}(\vec{r}, \tau)$. In this more general case, the above equality becomes an approximation. However, under reasonable experimental conditions, the approximation often remains quite accurate. In the following, we discuss the potential sources of error, and relevant limits under which the errors become negligible.

The first source of error in Eqn. 5 under time-varying $g_{LSF}(\vec{r}, \tau)$ arises because the second term of Eqn. 3 is no longer independent of τ . The resulting error is given by

$$g_{LSF}(\vec{r}, \tau) * g_p(\vec{r}) - g_{LSF}(\vec{r}, \tau_{max}) * g_p(\vec{r})$$

Note that $g_{LSF}(\vec{r}, \tau)$ is a PDF, so the magnitude of this error can be at most $g_p(\vec{r})$. As a result, the approximation is likely to be valid when $\frac{1}{\rho} (g_e(\tau) - g_e(\tau_{max})) \gg g_p(\vec{r})$. In practice, we find that even samples with relatively strong structure satisfy this assumption for short τ , where $g_e(\tau)$ decays rapidly. In addition, this source of error can be negligible if $g_p(\vec{r})$ is nearly constant, as is the case in a sample with weak interactions between labeled molecules. Similarly, if the LSF only broadens slightly, so $g_{LSF}(\vec{r}, \tau) \approx g_{LSF}(\vec{r}, \tau_{max})$ the error will also be negligible. This condition often holds for accurately drift-corrected images, where we find the width of the LSF to be within a few nm of the localization precision even at τ_{max} .

A second source of error in Eqn. 5 under time-varying $g_{LSF}(\vec{r}, \tau)$ comes from the first term of Eqn. 3. In particular, assuming that the first source of error is negligible, Eqn. 4 yields:

$$\Delta g(\vec{r}, \tau) \approx \frac{1}{\rho} (g_{LSF}(\vec{r}, \tau) g_e(\tau) - g_{LSF}(\vec{r}, \tau_{max}) g_e(\tau_{max})) \quad (6)$$

In principle, Eqn. 6 could be used to extract $g_{LSF}(\vec{r}, \tau)$ and $g_{LSF}(\vec{r}, \tau_{max})$ through simultaneous fitting. In practice, we make the further approximation that

$$(g_{LSF}(\vec{r}, \tau) g_e(\tau) - g_{LSF}(\vec{r}, \tau_{max}) g_e(\tau_{max})) \approx g_{LSF}(\vec{r}, \tau) (g_e(\tau) - g_e(\tau_{max}))$$

yielding the simple relation:

$$\Delta g(\vec{r}, \tau) \propto g_{LSF}(\vec{r}, \tau) \quad (7)$$

This applies in the limit of $g_e(\tau) \gg g_e(\tau_{max})$ or $g_{LSF}(\vec{r}, \tau) \approx g_{LSF}(\vec{r}, \tau_{max})$, but introduces some practical limitations that are discussed in detail in the next section.

Eqn. 7 can be used to estimate the full $g_{LSF}(\vec{r}, \tau)$ from acquired localizations. To summarize the LSF using a single number, we further assume that $g_{LSF}(\vec{r}, \tau)$ takes on a Gaussian form:

$$g_{LSF}(\vec{r}, \tau) \propto \exp\left(-\frac{x^2}{4\sigma_x^2(\tau)} - \frac{y^2}{4\sigma_y^2(\tau)}\right),$$

where $\sigma_x(\tau)$ is the standard deviation in the x direction of the distance between the true position of the molecule at time t and a localization at time $t + \tau$. The extra factor of 2 in the denominator accounts for the fact that g_{LSF} reports on the distribution of distances between pairs of localizations, resulting in twice the variance compared to the error in a single localization relative to the true location of the corresponding emitter. Typically, the LSF is isotropic in the lateral dimensions, so we take $\sigma_{xy} := \sigma_x = \sigma_y$ and compute angularly averaged correlation functions resulting in:

$$g_{LSF}(r, \tau) \propto \exp\left(-\frac{r^2}{4\sigma_{xy}(\tau)^2}\right) \quad (8)$$

It is convenient to also define the mean squared displacement $\sigma_r^2(\tau) = \sigma_x^2(\tau) + \sigma_y^2(\tau) = 2\sigma_{xy}^2(\tau)$, which accounts for errors in both dimensions. When localizations are acquired in three dimensions, the axial resolution often differs from the lateral resolution, and this component can be considered independently:

$$g_{LSF}(z, \tau) \propto \exp\left(-\frac{z^2}{4\sigma_z(\tau)^2}\right)$$

We have implemented this method as MATLAB code, which is available online (307).

5.4.2 Practical limitations of the LSF estimate.

The derivation above mathematically demonstrates why it is possible to simply isolate the autocorrelation of the LSF from the full autocorrelation of the image by tabulating the differences between measured autocorrelations obtained at different time-intervals (Eqn. 7). This simple method works because repeated observations of the same molecule are typically correlated in time while repeated observations of pairs of distinct molecules are uncorrelated in time. Because of this, simply subtracting observations at different time-intervals results in isolating the average contribution from multiple observations of the same molecule. As mentioned above, this simple view requires a few assumptions that limit the applicability of this approach.

The main assumption used to arrive at Eqn. 7 is that individual molecules on average produce localizations that are correlated in time. These correlations are expected to extend out to some finite time-interval beyond which the method no longer applies because repeated observations become uncorrelated. Past studies document surprisingly long correlation times for many (d)STORM and (f)PALM probes (295–297) under a range of imaging conditions, suggesting that Eqn. 7 should apply even at extended time-intervals for images generated using these methods and probes. PAINT probes, which produce localizations through binding and unbinding of a probe fluorophore to a target molecule, only produce temporal correlations up until the off-rate of the specific binding interaction, since the binding of new probes from solution does not depend on the history of probe binding to a specific site (47, 308). These reduced correlations contribute to the more uniform appearance of SMLM images acquired using PAINT, and will also limit the applicability of this method. Conveniently, plots of $g(r < \sigma_r, \tau)$ capture the time-interval dependence of correlated observations from single molecules ($g_e(\tau)$) up to a numerical offset, and examples showing this decay for several experiments with different fluorophores are shown in Supplemental Figure 5.9 (see Section 5.6 for Supplemental Figures). These curves can be used to guide the range of τ over which Eqn. 7 is expected to apply.

For the case where $g_{LSF}(\vec{r}, \tau) = g_{LSF}(\vec{r}, \tau_{max})$, Eqn. 7 will correctly estimate $g_{LSF}(\vec{r}, \tau)$ as long as $g_e(\tau) > g_e(\tau_{max})$, since the two LSFs will have the same shape in space and subtraction will not lead to distortion even when both components have similar amplitudes ($\Delta g_e(\tau) \ll g_e(\tau)$). When $g_{LSF}(\vec{r}, \tau)$ does not have exactly the same shape as $g_{LSF}(\vec{r}, \tau_{max})$, distortions can arise for small $\Delta g_e(\tau)/g_e(\tau_{max})$ that can lead to systematic errors in estimates of $g_{LSF}(\vec{r}, \tau)$. These systematic distortions are demonstrated in Supplemental Figure 5.10 for the example of a Gaussian LSF with a standard deviation that varies with τ . For cases where $g_{LSF}(\vec{r}, \tau)$ broadens slightly with increasing time-interval, our approach will produce a systematically narrow estimate of $g_{LSF}(\vec{r}, \tau)$ for small $\Delta g_e(\tau)/g_e(\tau_{max})$. This occurs, for example, when labeled molecules diffuse over length-scales comparable to the localization precision during the acquisition time, and is demonstrated on the simulated example of Supplemental Figure 5.11. For the purposes of this report, we do not include estimates $g_{LSF}(\vec{r}, \tau)$ that may be subject to this systematic bias, using a cutoff of $\frac{\Delta g(r < 5 \text{ nm}, \tau)}{g(r < 5 \text{ nm}, \tau_{max})} > \frac{1}{2}$. In principle, a user could extend the applicability of this method to larger τ by instead fitting to Eqn. 6 which independently models

$g_{LSF}(\vec{r}, \tau)$ and $g_{LSF}(\vec{r}, \tau_{max})$ or could instead tabulate $\Delta g(\vec{r}, \tau)$ for closely spaced τ where changes in $g_{LSF}(\vec{r}, \tau)$ are expected to be more subtle.

The decay of $g_e(\tau)$ means that the signal to noise ratio of $\Delta g(\vec{r}, \tau)$ will degrade at longer time-intervals, over which fewer correlated pairs are observed. To increase statistical stability, we group time-intervals into increasingly large disjoint τ bins and estimate $g(r, \tau)$ as a weighted average. In this report, τ bin edges are log-spaced to account for the exponential decay inherent in $g_e(\tau)$ and we report the average τ value of the bin, but occasionally show the full range of τ values included in the bin. We typically use the back quarter of the dataset to initially tabulate $g(\vec{r}, \tau_{max})$, meaning that $\frac{3}{4} < \tau_{max}/T_{max} < 1$, where T_{max} is the total acquisition time. We then identify the cutoff by finding the τ where $\Delta g(r < 5 \text{ nm}, \tau) / g(r < 5 \text{ nm}, \tau_{max})$ first falls below 0.5. We then recalculate $g(\vec{r}, \tau_{max})$ using this cutoff as the low τ edge of the τ_{max} bin. Statistical confidence is estimated through bootstrapping and we estimate the statistical power of $\Delta g(\vec{r}, \tau)$ directly from $g(\vec{r}, \tau)$, as described in Materials and Methods.

5.4.3 Validation through Simulation

To validate this approach, we generated simulated datasets of DNA origami nanorulers in which fluorophores are separated by a fixed distance of 50 nm. Fluorophore blinking was subject to a photophysical model based on (297, 300). Briefly, fluorophores could exist in an “on” state, one of three dark states, or a bleached state. Transitions between states were governed by a continuous time Markov process, with transition rates roughly based on those measured in (297) but modified to reflect the experimental conditions used to obtain experimental images in this work. Nanorulers were placed randomly and uniformly with an average density of $1/\mu\text{m}^2$ across a $40 \mu\text{m}$ by $40 \mu\text{m}$ field of view with the molecules having a localization precision of 10 nm in each lateral dimension. 20,000 image frames were simulated with a frame time of 0.1 s. Figure 5.1a illustrates a small field of view containing 3 nanorulers, both as a reconstructed image and with localizations colored by time. An image showing a larger subset of the field of view is shown as Supplemental Figure 5.12.

Simulated localizations were subjected to a spatiotemporal auto-correlation analysis as described in Methods and representative plots of the spatial component of $g(r, \tau)$ are shown in Figure 5.1b. This family of curves contains two major features: an initial peak at short displacements ($r < 40 \text{ nm}$) arising from multiple localizations from the same molecule, and a

second feature at wider radii ($40 \text{ nm} < r < 100 \text{ nm}$) arising from displacements between localizations from different molecules on the same ruler. The amplitude of the initial peak decreases with increasing τ , while the second feature is largely independent of τ . The τ dependent component is isolated by subtracting $g(r, \tau)$ at long τ from those arising from shorter τ to obtain $\Delta g(r, \tau)$ as shown in Figure 5.1c. In this simulation, there are no τ dependent effects that would impact resolution, resulting in $\Delta g(r, \tau)$ having the same width for all τ . This is summarized by fitting $\Delta g(r, \tau)$ to the Gaussian function of Eqn. 8 to extract the LSF width, σ_{xy} , reported in Figure 5.1d. Representative $\Delta g(r, \tau)$ with Gaussian fits are shown in Supplemental Figure 5.13. In this simulated example, we can associate all localizations with the molecules that produced them and can directly compute $g_{LSF}(r, \tau)$ from the relative positions of localizations originating from the same molecule as described in Methods. This is followed by fitting to Eqn. 8 to obtain σ_{xy} which is also reported in Figure 5.1d. Lastly, we tabulate displacements between all localizations originating from distinct molecules on the same ruler. The distribution of these displacements is shown in Figure 5.1e, and its properties are described by simulation parameters. The line in Figure 5.1e has a Gaussian shape with the form: $P(r) \propto \exp\left(-\frac{(r-\langle r \rangle)^2}{4\sigma_{xy}^2}\right)$ where σ_{xy} is the localization precision (10 nm) and $\langle r \rangle$ is the average displacement between localizations originating at the ruler endpoints (52 nm). The slight bias in $\langle r \rangle$ towards a value larger than the actual separation between molecules (50 nm) arises from the components of localizations that fall perpendicular to the ruler axis and always contribute positive values to the measured displacements ($\sqrt{50^2 + 2\sigma_{xy}^2} = 52 \text{ nm}$).

The simulation of Figure 5.1 does not contain any factors that degrade image resolution over time. In Figure 5.2, the same simulation is subjected to a directed drift in the x-direction as well as diffusive drift in both x and y. Drift is corrected using a mean shift algorithm (271) that works by evenly dividing localizations into time-bins, then finding the displacement that minimizes the mean distance between localizations across all time-bins. The applied drift and the calculated drift correction are shown in Figure 5.2a along with the resulting image reconstruction. $g(r, \tau)$ curves over different windows in time-interval (Figure 5.2b.) closely resemble those shown in the static simulation, but $\Delta g(r, \tau)$ now broadens with increasing τ (Figure 5.2c.). This broadening reflects a degradation of the LSF beyond the localization precision at all but the shortest time-intervals and plateaus near the timescale of drift correction (Figure 5.2d.).

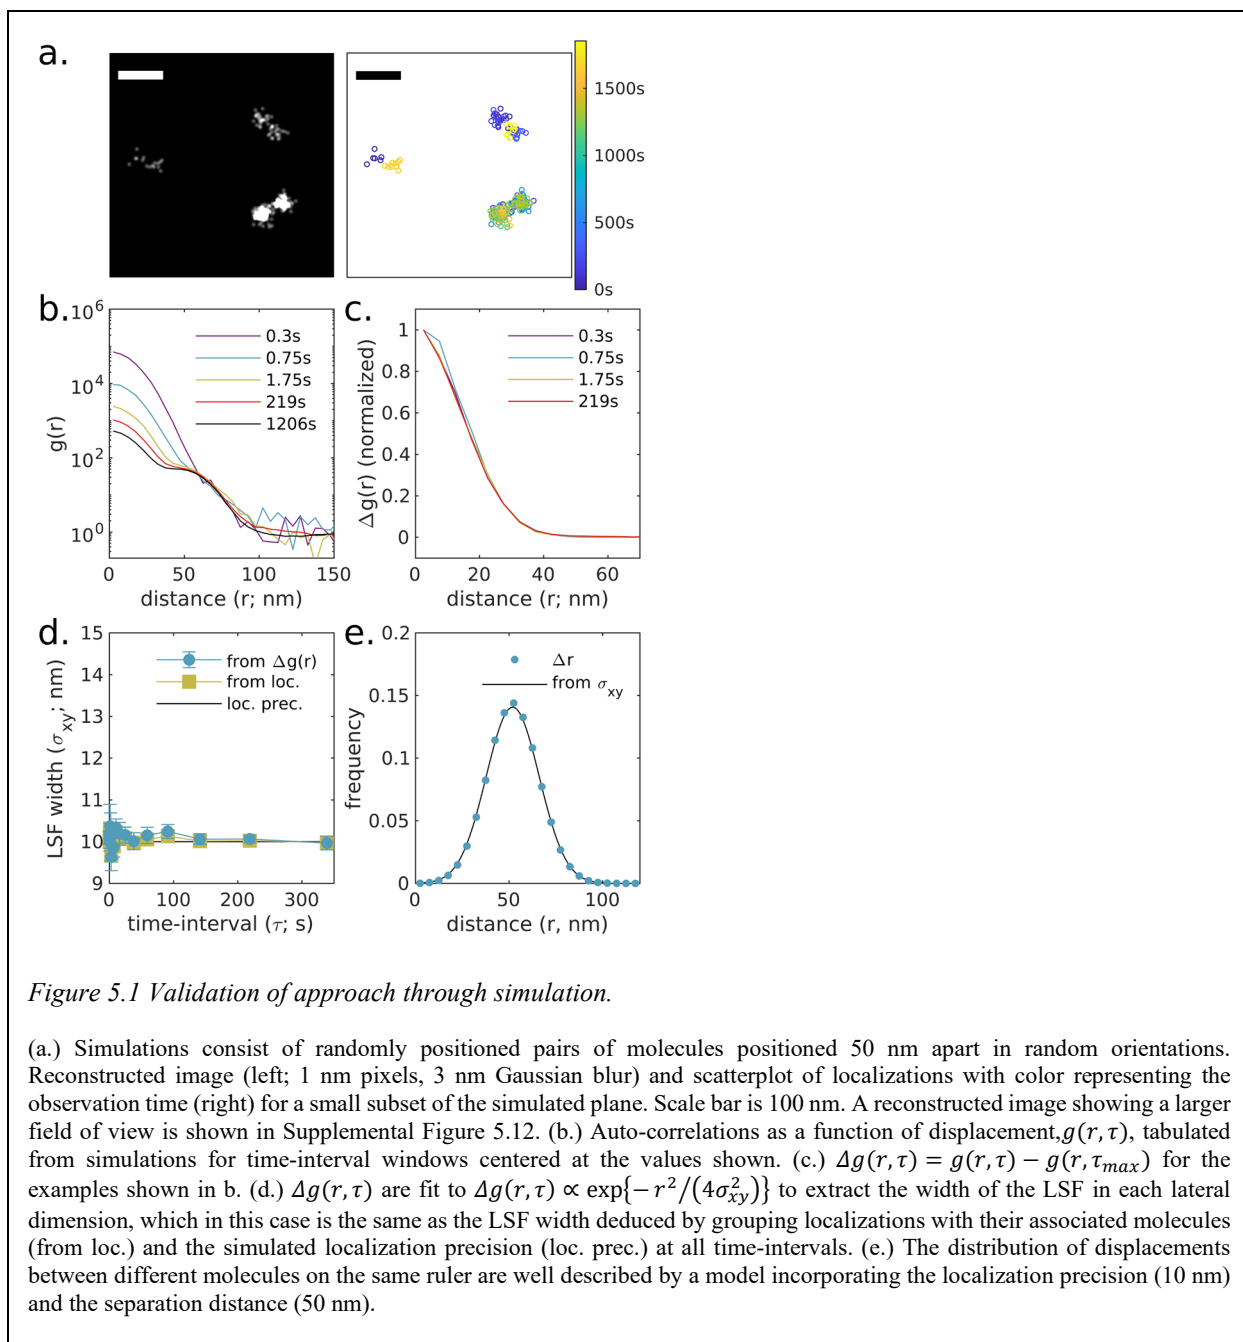


Figure 5.1 Validation of approach through simulation.

(a.) Simulations consist of randomly positioned pairs of molecules positioned 50 nm apart in random orientations. Reconstructed image (left; 1 nm pixels, 3 nm Gaussian blur) and scatterplot of localizations with color representing the observation time (right) for a small subset of the simulated plane. Scale bar is 100 nm. A reconstructed image showing a larger field of view is shown in Supplemental Figure 5.12. (b.) Auto-correlations as a function of displacement, $g(r, \tau)$, tabulated from simulations for time-interval windows centered at the values shown. (c.) $\Delta g(r, \tau) = g(r, \tau) - g(r, \tau_{max})$ for the examples shown in b. (d.) $\Delta g(r, \tau)$ are fit to $\Delta g(r, \tau) \propto \exp\{-r^2/(4\sigma_{xy}^2)\}$ to extract the width of the LSF in each lateral dimension, which in this case is the same as the LSF width deduced by grouping localizations with their associated molecules (from loc.) and the simulated localization precision (loc. prec.) at all time-intervals. (e.) The distribution of displacements between different molecules on the same ruler are well described by a model incorporating the localization precision (10 nm) and the separation distance (50 nm).

Here, the measured LSF width reaches a local maximum at a time separation somewhat smaller than the drift correction timescale, which we attribute to the drift correction algorithm itself as it is also apparent in the LSF width obtained by associating localizations with their originating molecules. Deviations from the expected distribution of pairwise distances between localizations of molecules from opposite ends of the same nanoruler also exceed those of the static case (Figure 5.2e) and are better described by a model that incorporates the measured image resolution,

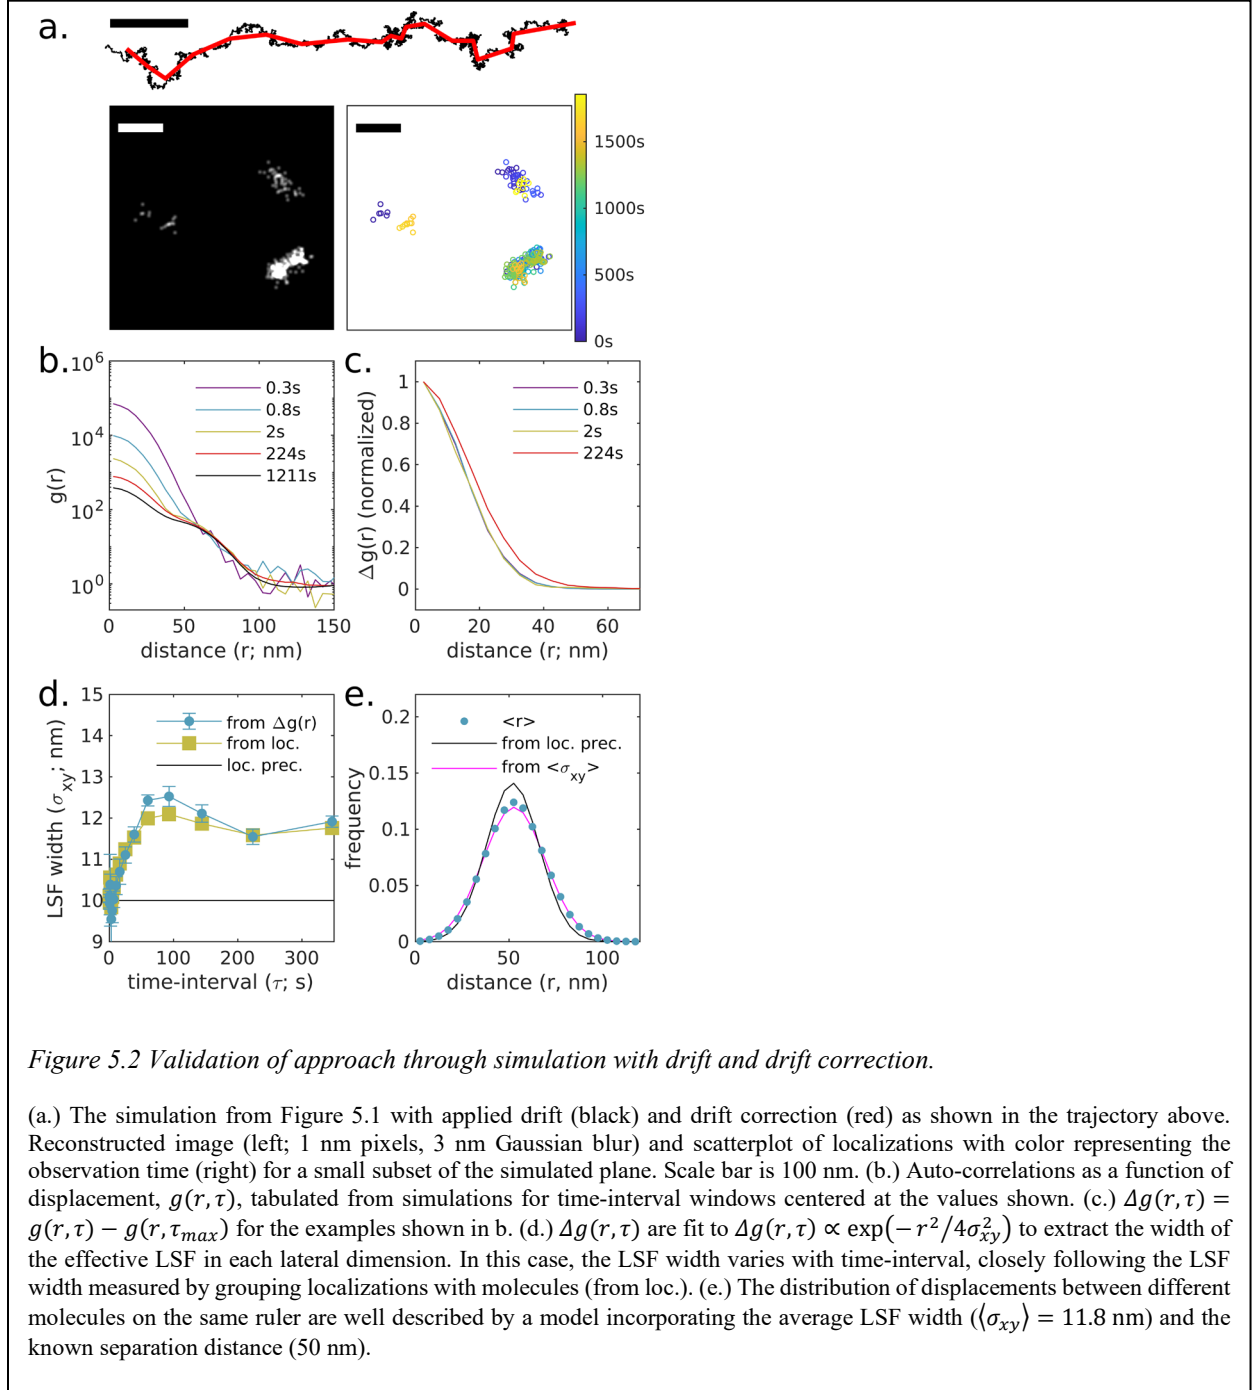


Figure 5.2 Validation of approach through simulation with drift and drift correction.

(a.) The simulation from Figure 5.1 with applied drift (black) and drift correction (red) as shown in the trajectory above. Reconstructed image (left; 1 nm pixels, 3 nm Gaussian blur) and scatterplot of localizations with color representing the observation time (right) for a small subset of the simulated plane. Scale bar is 100 nm. (b.) Auto-correlations as a function of displacement, $g(r, \tau)$, tabulated from simulations for time-interval windows centered at the values shown. (c.) $\Delta g(r, \tau) = g(r, \tau) - g(r, \tau_{max})$ for the examples shown in b. (d.) $\Delta g(r, \tau)$ are fit to $\Delta g(r, \tau) \propto \exp(-r^2/4\sigma_{xy}^2)$ to extract the width of the effective LSF in each lateral dimension. In this case, the LSF width varies with time-interval, closely following the LSF width measured by grouping localizations with molecules (from loc.). (e.) The distribution of displacements between different molecules on the same ruler are well described by a model incorporating the average LSF width ($\langle \sigma_{xy} \rangle = 11.8$ nm) and the known separation distance (50 nm).

$\langle \sigma_{xy} \rangle = 11.8$ nm, which is determined by averaging over estimated $\sigma_{xy}(\tau)$ weighted by the number of pairs associated with each time-interval window.

Observing a plateau in plots of $\sigma_{xy}(\tau)$ is a good indicator that the LSF estimator is generating reliable estimates, since drift correction is designed to stabilize localization error on long time-scales. Supplemental Figure 5.11 shows an example of the same simulation with drift

and drift correction, but where individual molecules are also allowed to diffuse slowly such that $g_{LSF}(r)$ broadens substantially with τ in a way that is not accounted for through drift correction. In that case, $\sigma_{xy}(\tau)$ increases with τ and is underestimated by $\Delta g(r, \tau)$. This is a case where the approximations needed to estimate $g_{LSF}(r)$ as $\Delta g(r, \tau)$ are not appropriate.

The simulated blinking dynamics of fluorophores includes a chance of photo-bleaching, or an irreversible transition into a dark state. In this example and in our experience in general, we find that this analysis is largely independent of photo-bleaching rates. This is because the normalization used to tabulate $g(r, \tau)$ accounts for any systematic reduction in the number of localizations over time. Photo-bleaching does reduce the number of pairs observed at large separation times, and therefore the statistical performance of the measurement.

5.4.4 Estimating the LSF of DNA origami datasets

Figure 5.3 demonstrates this approach on an experimental dataset of DNA origami nanorulers that resemble the simulated rulers with AlexaFluor647 labeling sites separated by 50 nm. Figure 5.3a shows a small subset of the field of view of the acquired image, that was reconstructed from 29,000 image frames acquired over 53 min at a frame rate of 0.1 s, with a total of over 126,000 individual localizations. In post-processing, a drift correction was applied with a time-window width of 25 s or 250 image frames. As in the simulated case, $g(r, \tau)$ decays at short r with increasing τ (Figure 5.3c), and $\Delta g(r, \tau)$ is roughly Gaussian (Figure 5.3d). Fitting $\Delta g(r, \tau)$ yields the resolution $\sigma_{xy}(\tau)$. As in the simulated example, the estimated LSF width is lowest at short time-intervals (6.5 ± 0.1 nm) and plateaus at time-scales somewhat shorter than the frequency of the applied drift correction.

Since the localization clouds from individual Alexa647 molecules were visually distinct, we applied a DBSCAN segmentation algorithm (306) to associate localizations with individual molecules. From this segmentation, we tabulated the LSF width within segmented molecules (Figure 5.3d) and find good general agreement with estimates of the LSF width obtained from $\Delta g(r, \tau)$ at short τ , further validating this approach. At longer τ , the LSF width differs somewhat by these two methods. We attribute this to inaccurate segmenting of localizations to molecules by the DBSCAN algorithm, which is expected to become more prominent at longer τ due to reduced temporal correlations of emitters. The segmented localizations are also used to tabulate the distribution of pairwise distances between different molecules on the same origami (Figure 5.3e).

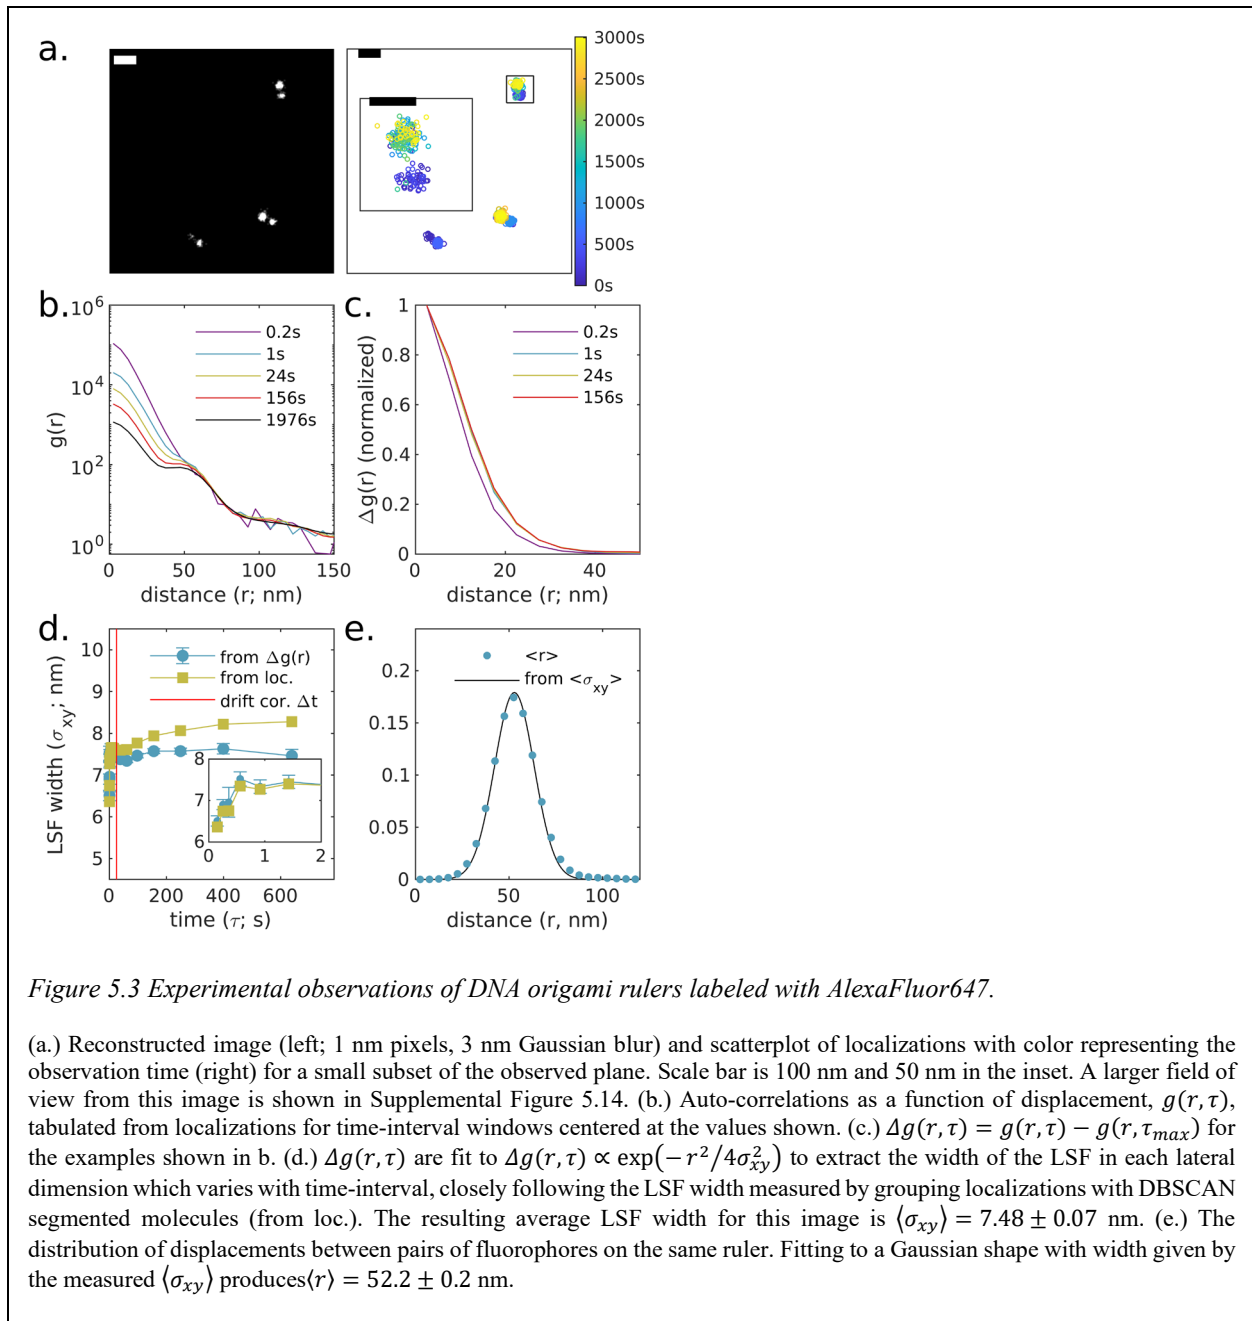
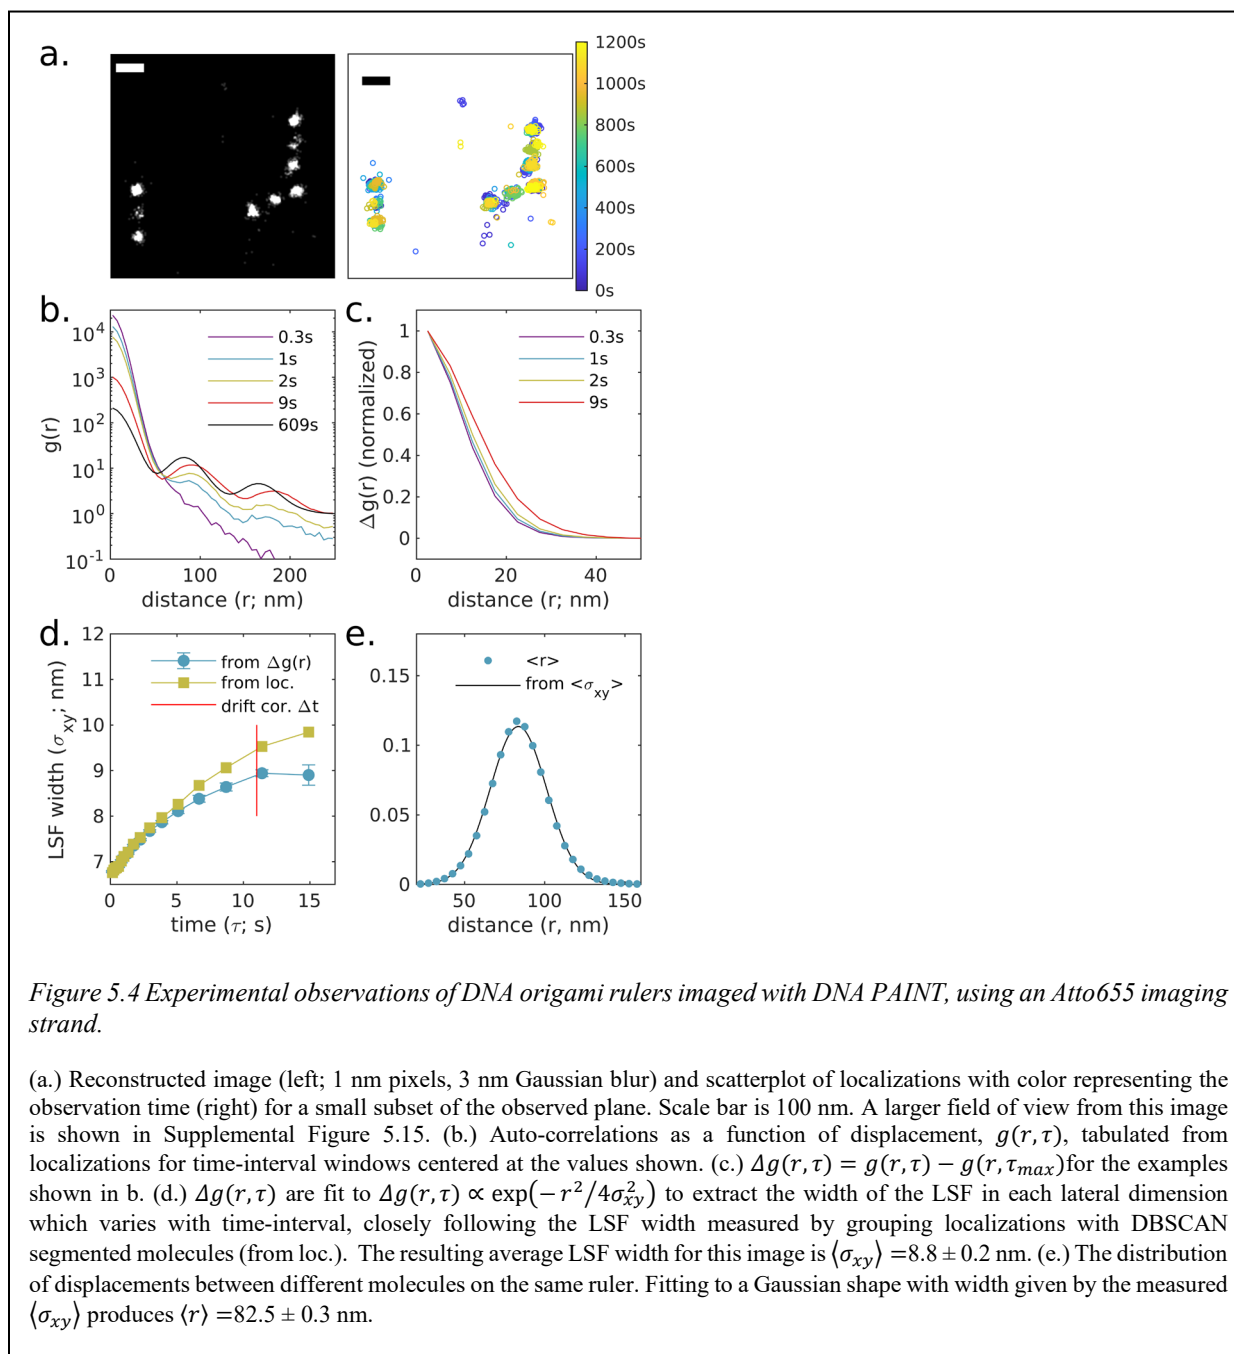


Figure 5.3 Experimental observations of DNA origami rulers labeled with AlexaFluor647.

(a.) Reconstructed image (left; 1 nm pixels, 3 nm Gaussian blur) and scatterplot of localizations with color representing the observation time (right) for a small subset of the observed plane. Scale bar is 100 nm and 50 nm in the inset. A larger field of view from this image is shown in Supplemental Figure 5.14. (b.) Auto-correlations as a function of displacement, $g(r, \tau)$, tabulated from localizations for time-interval windows centered at the values shown. (c.) $\Delta g(r, \tau) = g(r, \tau) - g(r, \tau_{max})$ for the examples shown in b. (d.) $\Delta g(r, \tau)$ are fit to $\Delta g(r, \tau) \propto \exp(-r^2/4\sigma_{xy}^2)$ to extract the width of the LSF in each lateral dimension which varies with time-interval, closely following the LSF width measured by grouping localizations with DBSCAN segmented molecules (from loc.). The resulting average LSF width for this image is $\langle \sigma_{xy} \rangle = 7.48 \pm 0.07$ nm. (e.) The distribution of displacements between pairs of fluorophores on the same ruler. Fitting to a Gaussian shape with width given by the measured $\langle \sigma_{xy} \rangle$ produces $\langle r \rangle = 52.2 \pm 0.2$ nm.

This distribution is well described by a model applying the measured $\langle \sigma_{xy}(\tau) \rangle = 7.5$ nm with $\langle r \rangle = 52.2 \pm 0.2$ nm, where the error is dominated by uncertainty in the sample magnification at the camera. This yields a separation distance of 51.1 ± 0.2 nm between labels on individual rulers, which is within the manufacturer's specifications.

We have conducted this same analysis on a similar DNA origami sample that was imaged using DNA PAINT, this time using rulers containing 3 collinear docking sites separated by 80 nm and summarized in Figure 5.4. In contrast to the dSTORM fluorophores of Figure 5.3, molecules



imaged by DNA PAINT do not exhibit long time-scale correlations, limiting the applicability of this method. The DNA PAINT probes used for this image do remain correlated over time-scales relevant for drift-correction (~ 15 s), which is long enough to provide a useful estimate of image resolution. In this example, drift correction was applied with a time-window width of 11 s or 110 image frames. Pairwise distances between labels on the center and ends of the origami were measured after applying DBSCAN to segment localizations from distinct docking sites and are

well described by a model applying the measured $\langle\sigma_{xy}(\tau)\rangle = 8.7$ nm, with $\langle r \rangle = 82.5 \pm 0.3$ nm. This yields a separation distance of 81.6 ± 0.3 nm between the center and endpoint labels on individual rulers. Since temporal correlations of the PAINT probes used in this example only extend for a small fraction of the acquisition time (20 min), the average σ_{xy} is given primarily by the value determined in the largest time-interval bin ($\tau = 15$ s). This is because the vast majority of pairs are detected at time-intervals that are grouped into τ_{max} , where we do not estimate σ_{xy} but instead apply the value estimated at the previous time-window bin. The good agreement between the model and measured distributions in Figure 5.3e validates this approach, at least for this specific example where drift correction was accomplished on a shorter time-scale.

5.4.5 Estimating the LSF from datasets of labeled structures in chemically fixed cells

We next apply this method to image labeled structures in cells. Figure 5.5 shows the method applied to nuclear pore complexes (NPCs) within the nuclear envelope of chemically fixed primary mouse neurons. In these images, a protein component of NPCs, Nup210, was labeled with a conventional primary antibody and a Fab secondary directly conjugated to AlexaFluor647. 12500 images were acquired over 23 min with an integration time of 0.1 s and a total of 178873 localizations detected within the masked ROI at the nuclear envelope. Drift correction was accomplished with a time-window of 8.3 s or 83 image frames. Reconstructed images of the entire nucleus and single pores are shown in Figure 5.5a. along with a scatter plot demonstrating that individual NPC subunits are sampled at times throughout the observation. $g(r, \tau)$ (Figure 5.5b.) curves extend to beyond 100 nm reflecting the extended structure of individual labeled NPCs, but extended structure is effectively removed by examining $\Delta g(r, \tau)$ (Figure 5.5c.). Fitting $\Delta g(r, \tau)$ to a Gaussian shape quantifies image LSF width over time, which is smallest at short τ (8.3 ± 0.3 nm) and increases at larger time-intervals. We estimate $\langle\sigma_{xy}\rangle$ to be 10.9 ± 0.8 nm.

Figure 5.6 shows a similar class of cellular structure imaged using DNA PAINT. In this example, clathrin-GFP is transiently expressed in CH27 cells then labeled post fixation with an anti GFP nanobody conjugated to an ssDNA docking strand. Cells are then imaged in the presence of a complementary imaging strand labeled with Atto 655. Similar to the origami DNA PAINT sample of Figure 5.4, temporal correlations from single molecules remain for short to medium time-scales (~ 8 s), allowing for accurate estimation of LSF broadening due to drift and drift correction. Here we estimate the average LSF width $\langle\sigma_{xy}\rangle$ to be 11.6 ± 0.3 nm.

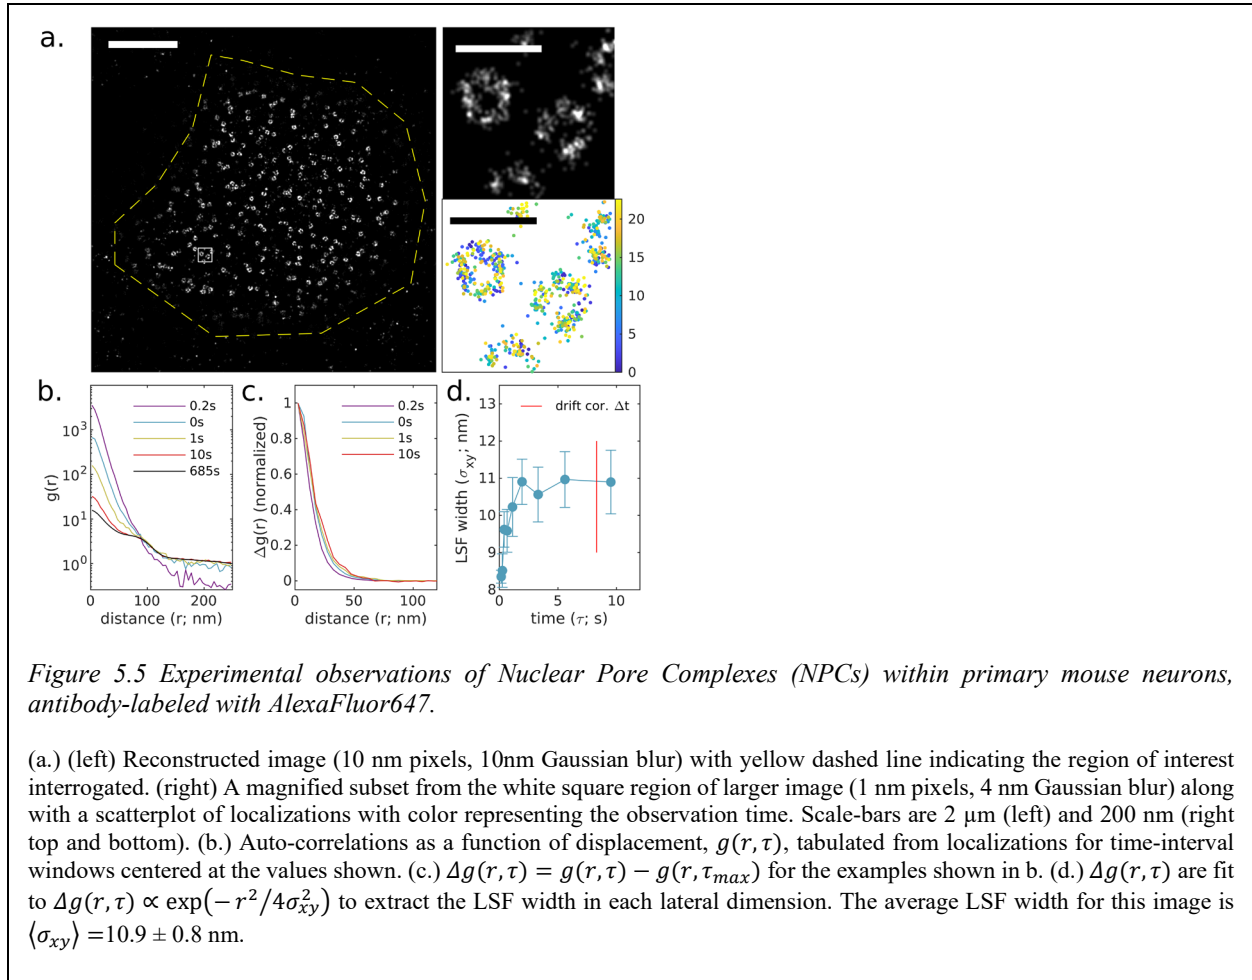
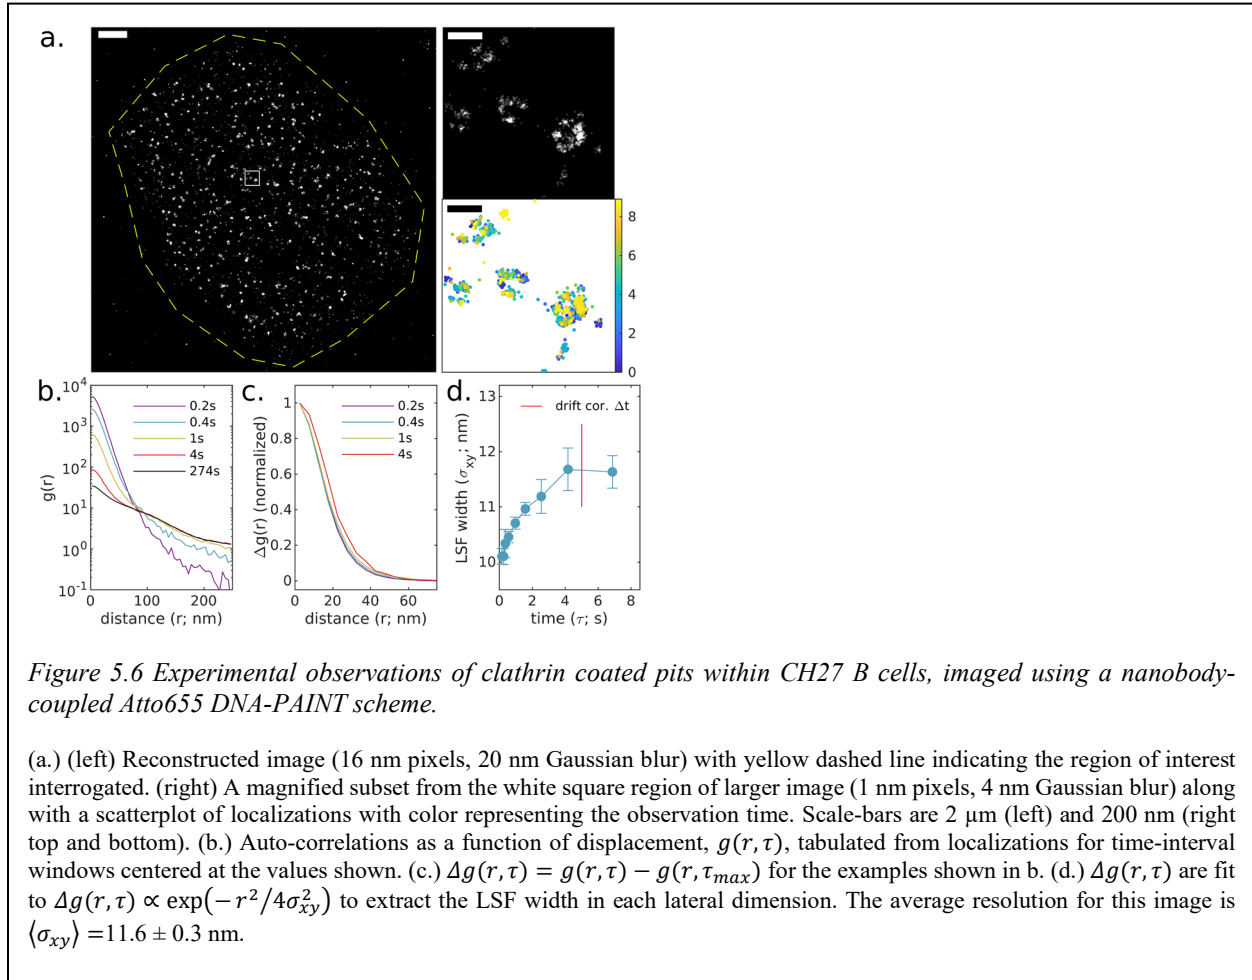


Figure 5.7 shows the method applied to an image of F-actin staining by phalloidin-AlexaFluor647 in chemically fixed CH27 B cells adhered to a glass surface decorated with VCAM. For this sample, 5000 images were acquired over 4.9 min with an integration time of 0.05 s and a total of 302681 localizations within the masked ROI. Drift correction was accomplished with a time-window of 2.5 s or 50 image frames. Unlike Figure 5.5 and Figure 5.6 where labels decorate isolated structures scattered over a surface, this reconstructed image of F-actin is more space filling, making up a web of fibers that extend across the entire ventral cell surface (Figure 5.7a.). This extended structure can be detected in $g(r, \tau)$ (Figure 5.7b.) as increased intensity in the tail that extends to large separation distances for curves generated at all τ . This large-scale structure is effectively removed in $\Delta g(r, \tau)$ (Figure 5.7c) allowing for a determination of the LSF width over a range of time-scales as shown in Figure 5.7d. In this example, the ROI was drawn within the cell boundary to minimize the intensity of $g_p(r)$ which allows for accurate estimation of $g_{LSF}(r, \tau)$ out to longer time-intervals. This is because the amplitude of $g(r, \tau_{max})$ includes



contributions from $g_p(r)$, while the amplitude of $\Delta g(r, \tau)$ only depends on $g_e(\tau)$, therefore $\frac{\Delta g(r < 25 \text{ nm}, \tau)}{g(r < 25 \text{ nm}, \tau_{max})}$ will remain larger than the cutoff for a wider range of τ . The estimate for $\langle \sigma_{xy} \rangle$ is 11.8 ± 1.5 nm.

As a final demonstration, Figure 5.8 shows the method applied to an image of Src15-mEos3.2, a myristoylated peptide bound to the inner leaflet of the plasma membrane and directly conjugated to the photo-switchable fluorescent protein mEos3.2. This peptide uniformly decorates the ventral surface of a chemically fixed CH27 B cell adhered to a glass surface decorated with VCAM, as seen in the reconstructed image of Figure 5.8a. For this sample, 7000 images were acquired over 12.7 min with an integration time of 0.1 s and a total of 240,503 localizations. Drift correction was accomplished with a time-window of 12.5 s or 125 image frames. mEos3.2 exhibits different blinking dynamics than AlexaFluor647, with some probes exhibiting correlated blinking on long time-scales. This can be seen in plots of $g(r, \tau)$ that take long timescales to decay

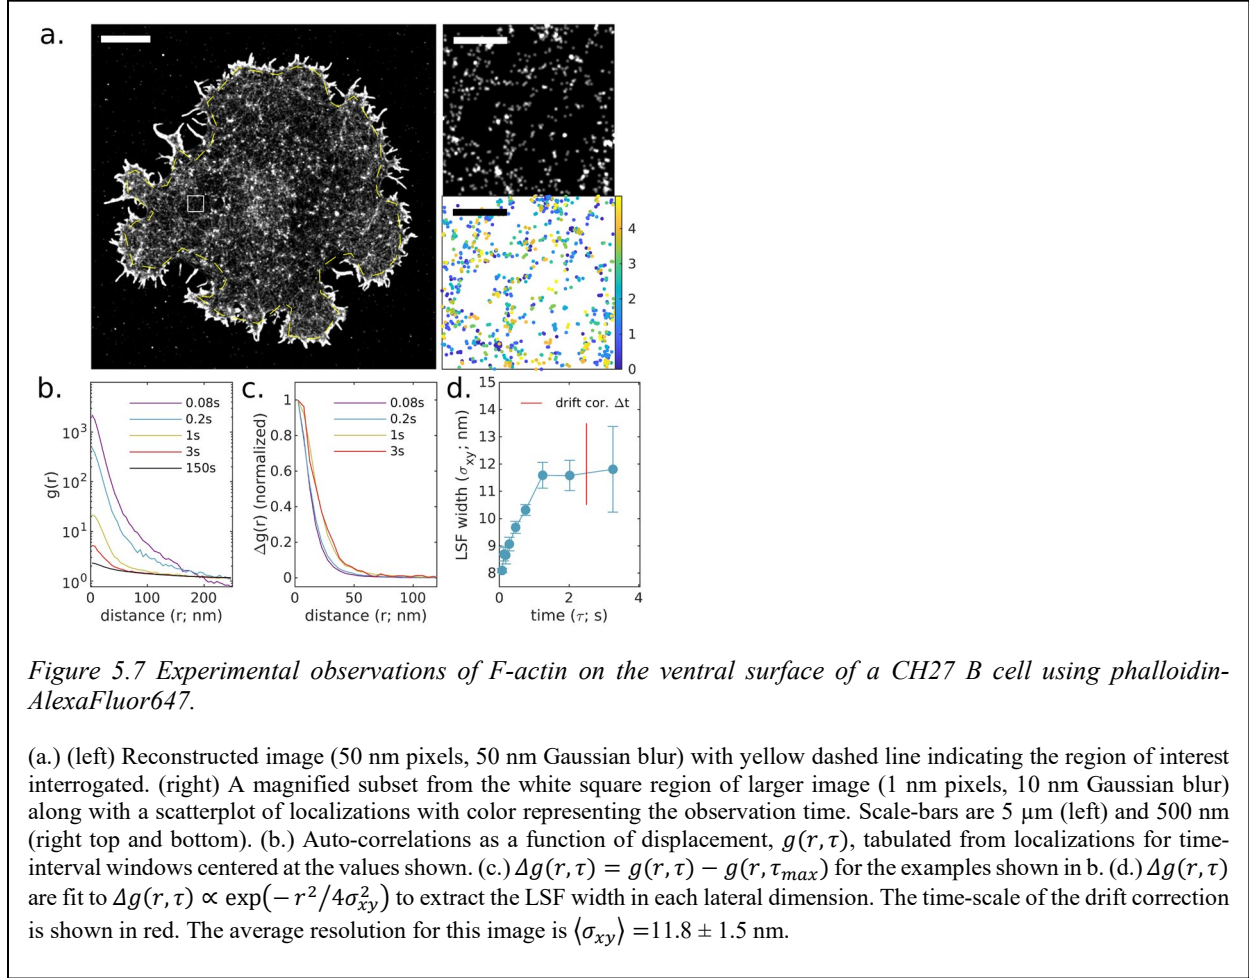


Figure 5.7 Experimental observations of F-actin on the ventral surface of a CH27 B cell using phalloidin-AlexaFluor647.

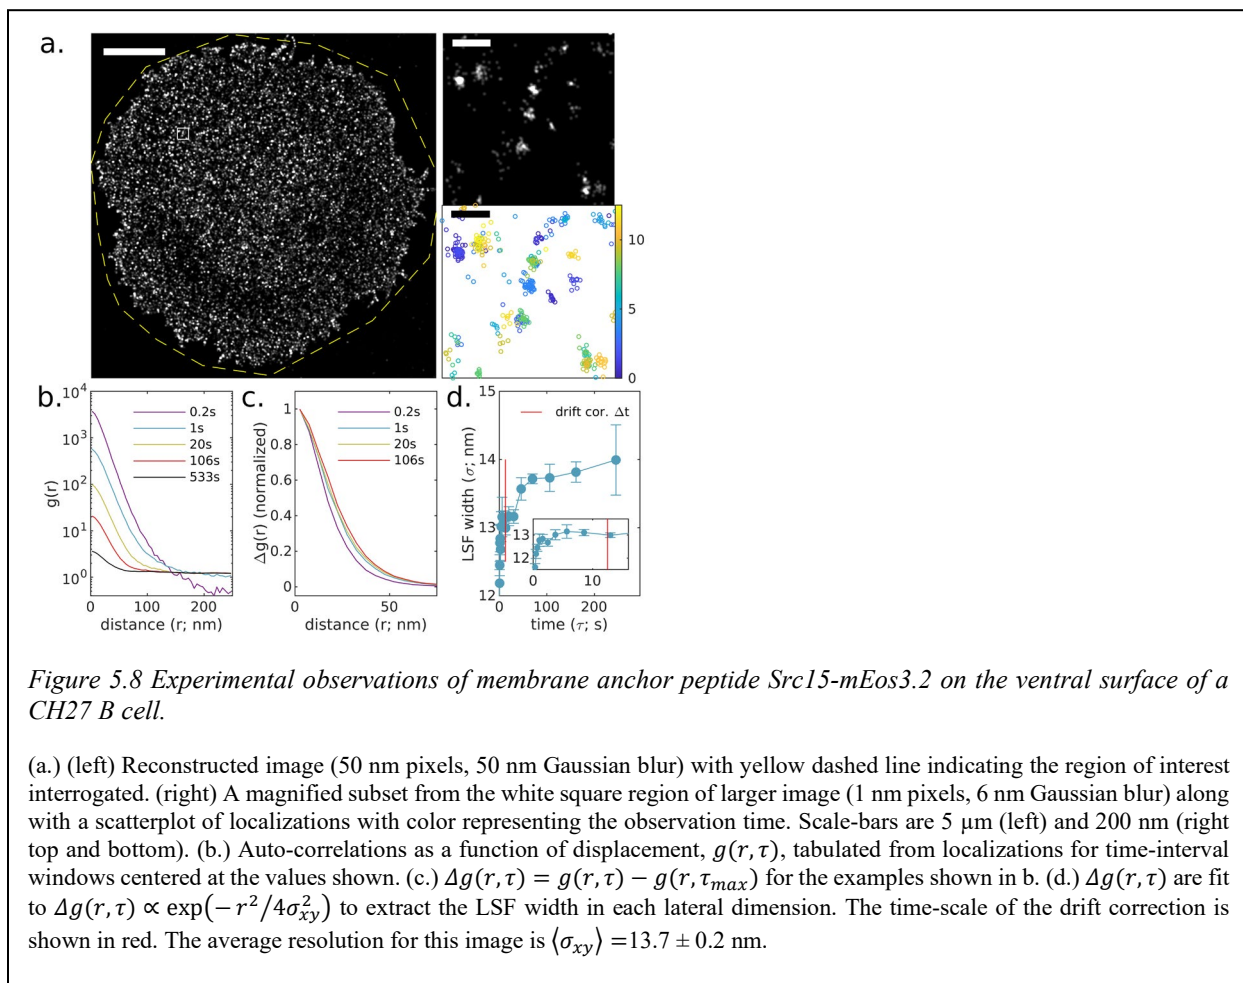
(a.) (left) Reconstructed image (50 nm pixels, 50 nm Gaussian blur) with yellow dashed line indicating the region of interest interrogated. (right) A magnified subset from the white square region of larger image (1 nm pixels, 10 nm Gaussian blur) along with a scatterplot of localizations with color representing the observation time. Scale-bars are 5 μm (left) and 500 nm (right top and bottom). (b.) Auto-correlations as a function of displacement, $g(r, \tau)$, tabulated from localizations for time-interval windows centered at the values shown. (c.) $\Delta g(r, \tau) = g(r, \tau) - g(r, \tau_{max})$ for the examples shown in b. (d.) $\Delta g(r, \tau)$ are fit to $\Delta g(r, \tau) \propto \exp(-r^2/4\sigma_{xy}^2)$ to extract the LSF width in each lateral dimension. The time-scale of the drift correction is shown in red. The average resolution for this image is $\langle \sigma_{xy} \rangle = 11.8 \pm 1.5$ nm.

(Figure 5.8b). Again, $\Delta g(r, \tau)$ curves isolate the initial peak, allowing for the quantification of the LSF width. In this example, the slow decay of $g(r < 25 \text{ nm}, \tau)$ with τ allows for estimation of σ_{xy} out to large time-intervals. The estimate for $\langle \sigma_{xy} \rangle$ is 13.7 ± 0.2 nm.

5.4.6 Comparison to other measures of image resolution

The datasets interrogated in Figures 1-8 were also subjected to other methods that report on image resolution and results are summarized in Table 1. These include the Fourier Ring Correlation (FRC (80, 81)), errors returned directly from fitting localizations (Cramer-Rao lower bound or CRLB of the variance of a maximum likelihood estimator (58, 56)), and using the nearest neighbor distribution in adjacent frames (Nearest neighbor based analysis, or NeNA (82)).

The FRC method involves reconstructing images and quantifying them in Fourier space, identifying the highest frequency signals that exceed some predetermined noise threshold, as illustrated in Supplemental Figure 5.16. As a result, the resolution values returned by the FRC



algorithm depend on the localization accuracy of single emitters but also the structure present in images and the spatial sampling of that structure. Because of this, the FRC resolutions reported in Table 1 do not trend systematically with LSF widths extracted through $\Delta g(r, \tau)$, as these report on the localization accuracy alone. The FRC method is sensitive to factors that erode the LSF over time, as indicated by the larger value generated from the simulated dataset with drift and imperfect drift correction (Figure 5.2; 35 nm) as compared to the simulation without drift (Figure 5.1; 30 nm). The FRC is also highly dependent on the types of structures imaged, returning a very large value for the dataset containing f-actin localizations (Figure 5.7; 285 nm) even though the LSF width estimate from $\Delta g(r, \tau)$ is similar to the other samples imaged (11.8 nm). This is because this dataset contains the most space-filling structure so there is less signal at high spatial frequencies, and because this structure is less spatially sampled than the other datasets shown. We also note that long blinking time correlations can lead to artificially low FRC resolution values

(81), whereas the LSF width estimate is dependent on these time correlations to achieve accurate results.

The fitting algorithms used return estimated localization precisions for each fit, known as the Cramer-Rao lower bound, or CRLB. These error estimates are obtained by applying a model that incorporates photon counting statistics, the Gaussian shape of the point spread function, and specifics of the camera such as its gain and offset. The accuracy of the error estimates depends in turn on the accuracy of the assumptions used to build the fitting model. While it is possible to achieve accurate error estimates with CRLBs (58), we find that factors such as imperfect gain calibration, read noise, and non-uniform background fluorescence frequently lead to imperfect error estimates in realistic experimental conditions. Our typical imaging processing pipeline involves a pre-processing step in which the image background is estimated then subtracted from raw image frames. The mean of the background on each fitting region is added back to that region, to approximately reproduce the appropriate counting statistics. Since this background subtraction is not incorporated into the model used to estimate errors, the values returned by the fitting algorithm are inaccurate. The values reported in Table 1 are obtained by re-fitting the dataset without background subtraction, leading to somewhat different values of $\sigma_{xy}(\tau)$ as shown in Supplemental Figure 5.17. For each dataset, there is a broad distribution of errors peaked at a value close to the LSF width estimated from $\Delta g(r, \tau)$ at the shortest time-interval interrogated ($\sigma_{xy}(\tau_{frame})$). This distribution extends asymmetrically to larger errors, biasing the average error to larger values. We speculate that the average error from fits differs from that estimated by $\Delta g(r, \tau)$ at short τ due to simplifying or inaccurate assumptions in the model employed by the fitting algorithm. It is also possible that larger errors originate from localizations that are not correlated in time and therefore do not contribute to the estimate from $\Delta g(r, \tau_{frame})$.

A third method uses the distribution of nearest neighbor distances between probes imaged in adjacent frames to estimate the average localization precision through fitting (Supplemental Figure 5.18) (82). This method makes the often valid assumption that the vast majority of neighbor localizations in adjacent frames arise from the same labeled molecule, therefore this distribution reports on the accuracy of localization at short time intervals. As expected the localization precision reported by the nearest neighbor distribution method is in good general agreement for the LSF width estimated from $\Delta g(r, \tau)$ at the shortest time-intervals interrogated ($\sigma_{xy}(\tau_{frame})$), especially for datasets of well-separated molecules (Figs 1-4). For images of cellular structures,

we find a slight bias of the NeNA estimate towards smaller values than those estimated from $\Delta g(r, \tau)$. In these samples, a larger fraction of nearest neighbors originate from different molecules, and we speculate that the correction terms used to fit these additional neighbors contribute to this bias. Lastly, we note that the NeNA estimates do not capture broadening of the LSF by factors that erode localization precision over time. This is expected since only nearest neighbors from adjacent frames are incorporated in the analyzed distributions.

5.5 Conclusions

Here we present a method to estimate the localization spread function of a SMLM measurement directly from acquired localizations, relying on a few reasonable assumptions. The basic method is validated through simulations and demonstrated using experimental data of three commonly used localization microscopy probes. The described method performs best when used alongside fluorophores that exhibit blinking dynamics that remain correlated in time out to time-scales relevant to sources of error present in the imaging experiment. The width of the LSF, here reported by fitting to a Gaussian shape, directly reports on how accurately the positions of molecules are recorded at the end of an experimental and analytical pipeline, and can be used to optimize imaging protocols or assist in the interpretation or further processing of imaged structures.

Directly measuring the LSF allows experimenters to validate and optimize imaging and processing methods which can be difficult to accomplish using existing metrics of image resolution. For example, while commonly used fitting algorithms return localization precisions, these estimates are only valid when raw image frames are fit, prohibiting the use of pre-processing steps such as filtering or background subtraction. In contrast, the FRC resolution metric can be used to compare processing steps accomplished on the same set of observations, but variation across samples can be hard to interpret since this measure depends on the LSF, the sampling of the image, and the types of structures being imaged. A distinct advantage of the LSF estimate described here is that it directly reports on how accurately one can measure distances between localized molecules in an image. We expect this method to be useful when interpreting experiments that involve the measurement of distances between localizations in images, for example in nanometer-precision distance measurement methods (309), where an accurate estimate of the localization error is essential. We also expect it to be useful in the interpretation of spatial

auto- or cross-correlation analysis which report on the statistical distribution or co-distribution of labeled components (83, 87, 310).

The values obtained using the described method depend on long time-scale correlations of fluorescent probes used for imaging. For the examples shown, we find that it is most important to characterize resolution lost on time-scales shorter than the time-scale of drift correction, which appears to be the most important source of time-dependent degradation of the LSF in these cases. Fortunately, numerous methods exist to correct for rigid drift on time-scales relevant to the temporal correlations of many SMLM probes (61–65, 282, 284, 285), suggesting that the method presented in this report is broadly applicable for a range of experimental conditions. Apart from drift, the estimated LSF is also sensitive to other time-dependent sources of error (or their absence). For example, a recent report observes residual motions of fluorescent labels in fixed samples, which would appear in the LSF as a degradation of LSF width over the timescale of the relevant motions (299).


The reported method estimates resolution by fitting the estimated LSF at each time-interval probed to a Gaussian shape, followed by a weighted average to extract the estimated error of the average localization in an image. This approach is convenient because the resolution is summarized as a single number. However, it is possible to extract more detailed information about the LSF by not averaging over angles, or by loosening the assumption of a Gaussian shape. Supplemental Figure 5.19 shows the autocorrelation of the LSF, $g_{LSF}(\vec{r}, \tau)$, for the NPC dataset of Figure 5.5, expanded in x , y , and τ . Supplemental Figure 5.20 shows the weighted time-average $\langle g_{LSF} \rangle := \langle g_{LSF}(\vec{r}, \tau) \rangle_{\tau}$ for the six datasets of Figs 3-8, with weights given by the observed number of localization pairs separated by each time delay τ . Beyond quantifying resolution, we anticipate that the full average effective PSF could prove useful for other purposes, such as deconvolution of reconstructed images or tabulated spatial correlation functions, or as an input to clustering algorithms or other analysis tools.

Table 5.1 Summary of values obtained for several resolution measures for the datasets shown in figures

All units are nm and errors, when evaluated, are included in parenthesis. LSF widths are estimated from $\Delta g(r, \tau)$ and by associating localizations with molecules (from segments, not possible for cellular images). $\sigma_{xy}(\tau_{frame})$ is the LSF width for τ close to the frame time (τ_{frame}) and $\langle \sigma_{xy} \rangle$ is the LSF widths averaged over all τ . FRC are values obtained using the Fourier Ring Correlation. Full FRC curves are included in Supplemental Figure 5.16. CRLB values describe features of the distribution of Cramer-Rao lower bounds returned by the fitting procedure for datasets processed without background subtraction. Full distributions and $\sigma_{xy}(\tau)$ estimated from these differently processed images are included in Supplemental Figure 5.17. NeNA are obtained by fitting nearest neighbor distributions from localizations in adjacent frames. Distributions and fits are included in Supplemental Figure 5.18.


Dataset	LSF (from $\Delta g(r, \tau)$)		LSF (from segments)		FRC	CRLB		NeNA
	$\sigma_{xy}(\tau_{frame})$	$\langle \sigma_{xy} \rangle$	$\sigma_{xy}(\tau_{frame})$	$\langle \sigma_{xy} \rangle$		peak	avg	
Figure 1	10.1 (0.2)	10.02 (0.08)	10.1	10.0	30 (1)	N/A		10.0
Figure 2	10.1 (0.1)	11.79 (0.07)	10.1	11.5	35 (1)	N/A		10.0
Figure 3	6.5 (0.1)	7.48 (0.07)	6.4	8.1	20 (1)	6.3	9.0	6.7
Figure 4	6.78 (0.02)	8.8 (0.2)	6.8	9.7	27 (1)	6.8	7.8	6.8
Figure 5	8.3 (0.3)	10.9 (0.8)	N/A	N/A	35 (1)	5.6	9.5	7.8
Figure 6	10.1 (0.1)	11.6 (0.3)	N/A	N/A	48 (2)	10.9	14.1	9.3
Figure 7	8.1 (0.1)	11.8 (1.5)	N/A	N/A	285 (6)	N/A	N/A	7.5
Figure 8	11.0 (0.1)	13.7 (0.2)	N/A	N/A	59 (6)	9.9	14.2	9.9

5.6 Supplementary Figures



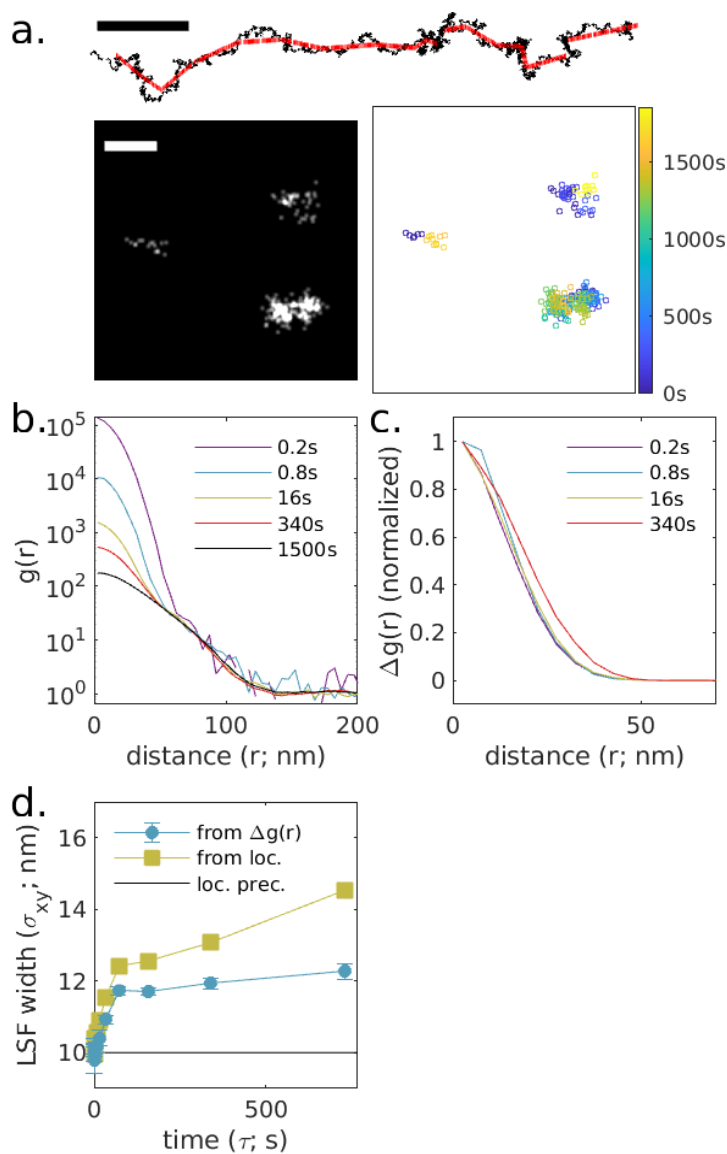
Supplemental Figure 5.9 Plots of $g(r < 25 \text{ nm}, \tau)$ for the experimental samples shown the main text.

These curves capture $g_e(\tau)$ up to a numerical offset that is dependent on the structure present in the image. Black lines are fit to a sum of exponentials and are present to highlight the monotonically decreasing trend. Dashed lines indicate the average value over the last $\frac{1}{4}$ of the dataset. The red vertical line indicates where $g(r < 25 \text{ nm}, \tau)$ falls below 1.5 times the dashed line, indicating the maximum τ expected to yield an unbiased estimate of the LSF from $\Delta g(r, \tau)$.




Supplemental Figure 5.10 Subtracting Gaussian shapes with different width leads to distortion in $\Delta g(r) = g(r, \tau) - g(r, \tau_{max})$ when $g(r, \tau)$ and $g(r, \tau_{max})$ have similar amplitudes but different widths.

(top) plots of $\Delta g(r) = (A + \Delta A) \exp(-r^2/4\sigma^2) - A \exp(-r^2/4\sigma_{\tau_{max}}^2)$ for $\sigma_{\tau_{max}} = 1.1\sigma$ (left) and $\sigma_{\tau_{max}} = 0.9\sigma$ (right) and $\Delta A = 0.25, 0.5, 1, 2$ from purple to red. Curves are normalized so they pass through 1 at $r = 0$. The legend shows the width extracted when fitting $\Delta g(r)$ to a single Gaussian shape $\Delta g(r) = A \exp(-r^2/4\sigma_{fit}^2)$. A broader $\sigma_{\tau_{max}}$ leads to systematic narrowing of σ_{fit} , while a narrow $\sigma_{\tau_{max}}$ leads to systematic broadening of σ_{fit} when the difference in amplitudes is order 1. (bottom) a summary of results over a broad range of ΔA and $\sigma_{\tau_{max}}$ indicates that distortion is not a major concern over broad range of values interrogated.

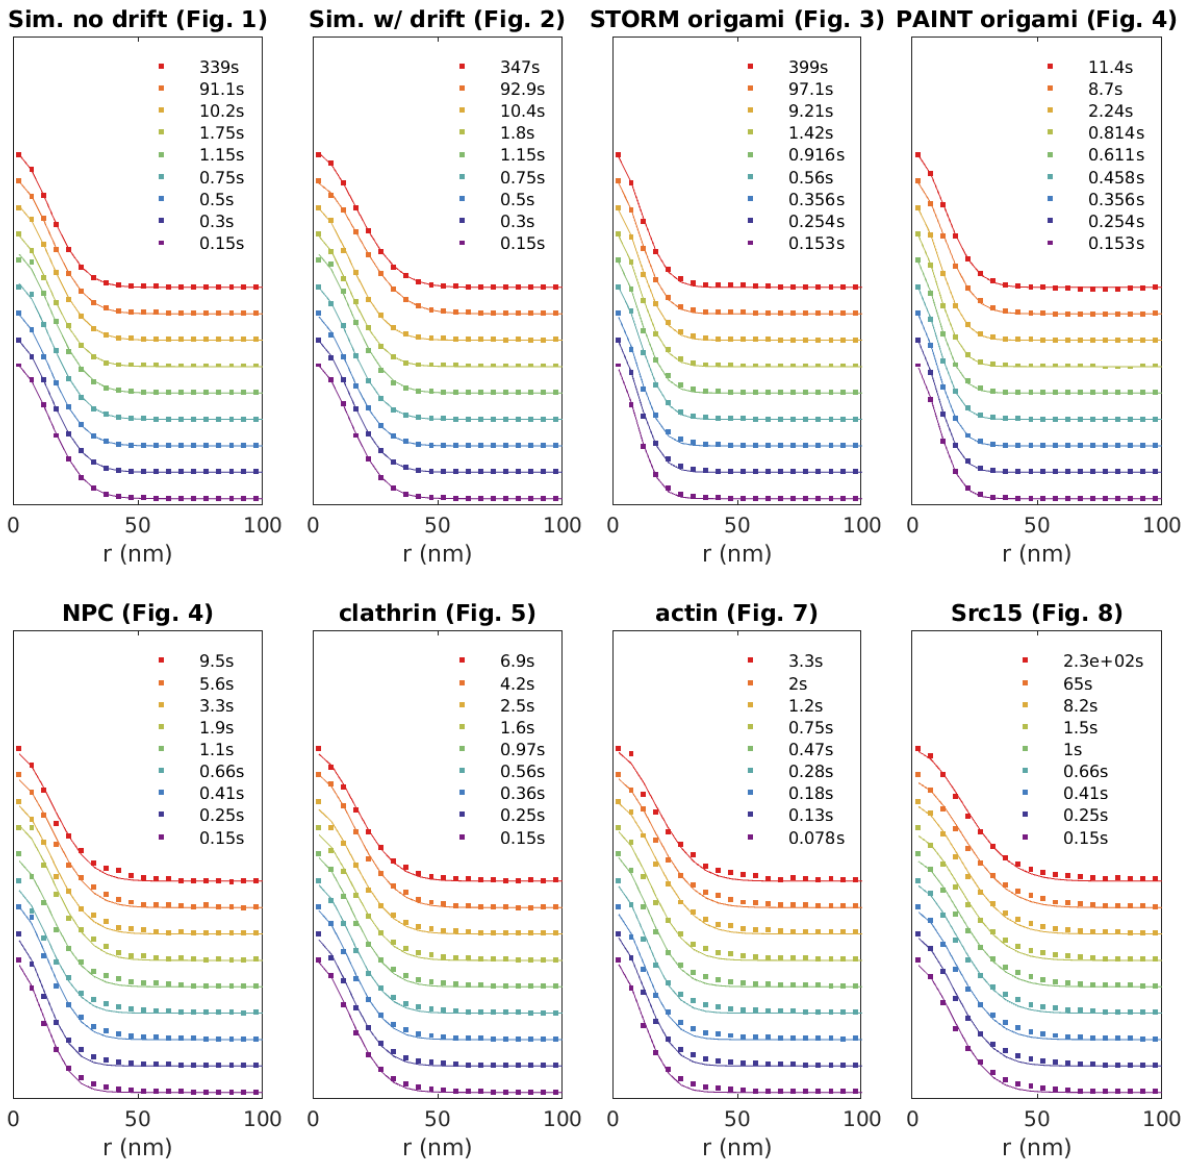


Supplemental Figure 5.11 Simulation with drift, drift correction and incoherent single molecule motions.

(a.) The simulation from Figure 5.1 with applied drift (black) and drift correction (red) as shown in the trajectory above as well as single molecule diffusion with $D = 1 \text{ nm}^2/\text{s}$. Reconstructed image (left; 1 nm pixels, 3 nm Gaussian blur) and scatterplot of localizations with color representing the observation time (right) for a small subset of the simulated plane. Scale-bar is 100 nm. (b.) Auto-correlations as a function of displacement, $g(r, \tau)$, tabulated from simulations for time-interval windows centered at the values shown. (c.) $\Delta g(r, \tau) = g(r, \tau) - g(r, \tau = 1500 \text{ s})$ for the examples shown in b. (d.) $\Delta g(r, \tau)$ are fit to $\Delta g(r, \tau) \propto \exp(-r^2/4\sigma_{xy}^2)$ to extract the LSF width in each lateral dimension (from $\Delta g(r)$). The LSF width from $\Delta g(r, \tau)$ varies with time-interval and is systematically narrower than the LSF measured by grouping localizations with molecules (from loc.). Error bars represent estimates of the standard error obtained through bootstrapping. This is due to the distortion effect demonstrated in Supplemental Figure 5.9 and is characterized by a σ_{xy} that increases with τ .




Supplemental Figure 5.12 10 μm by 10 μm region showing simulated localizations from Figs 1-2. The full simulated area was 40 μm by 40 μm . Scale bar is 1 μm .



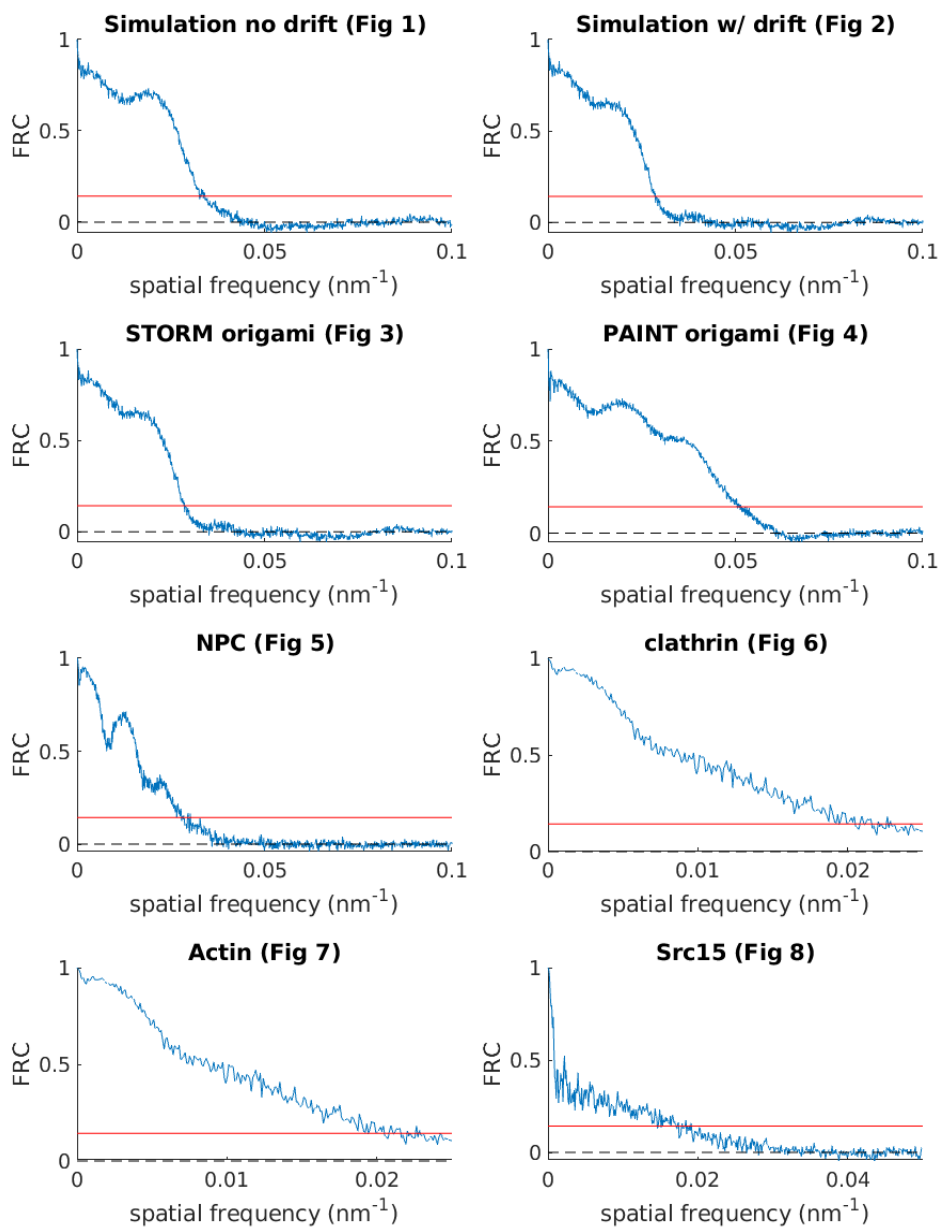
Supplemental Figure 5.13 Localization spread functions (LSFs) estimated as $\Delta g(r, \tau)$ (points) are well described by Gaussian functions (Eqn. 8 of main text; solid lines). Fits are used to extract $\sigma_{xy}(\tau)$ reported in the main text.



Supplemental Figure 5.14 10 μm by 10 μm region showing DNA origami rulers analyzed in Figure 5.3. The full imaged area was 40 μm by 40 μm . Scale bar is 1 μm .

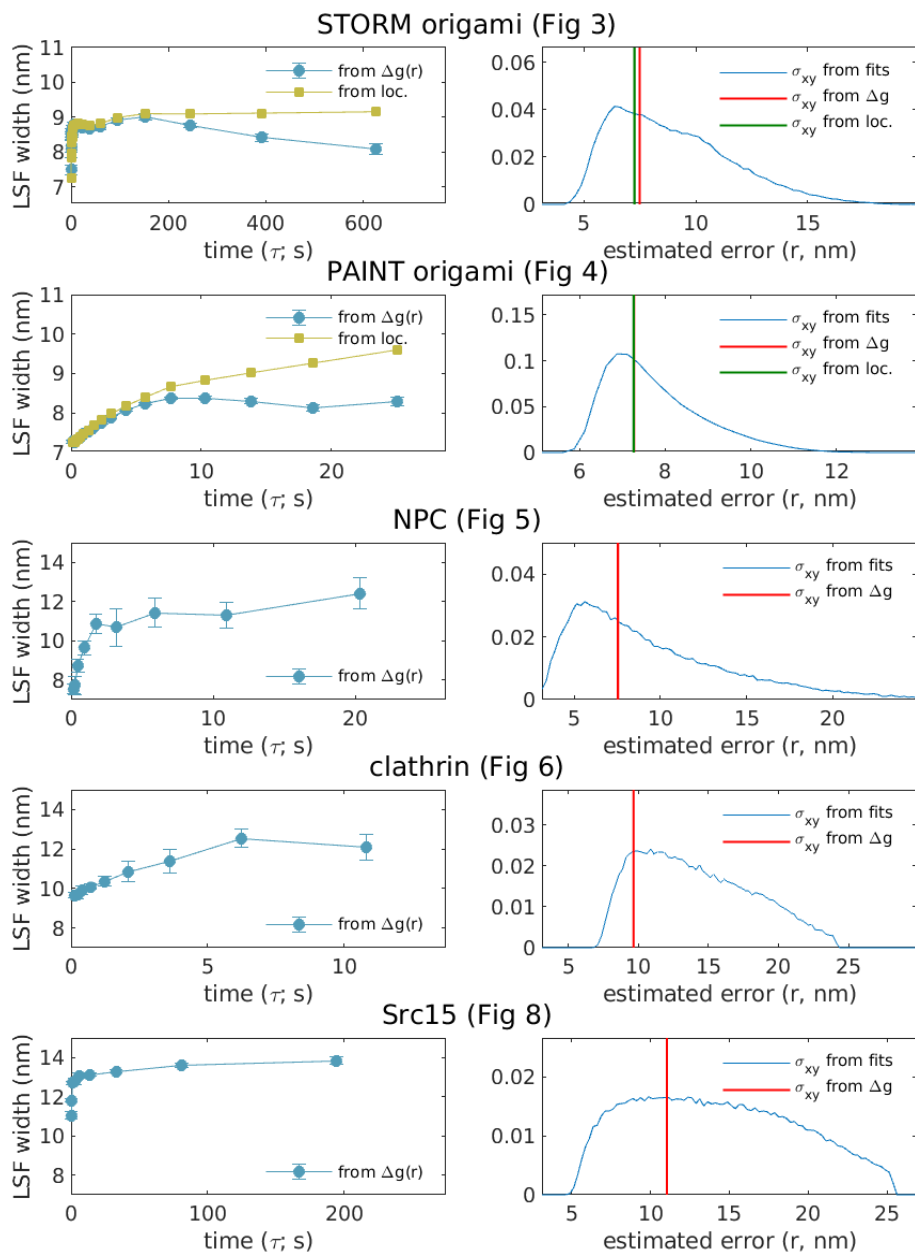


Supplemental Figure 5.15 10 μm by 10 μm region showing DNA origami rulers analyzed in Figure 5.4. The full imaged area was 40 μm by 40 μm. Scale bar is 1 μm.



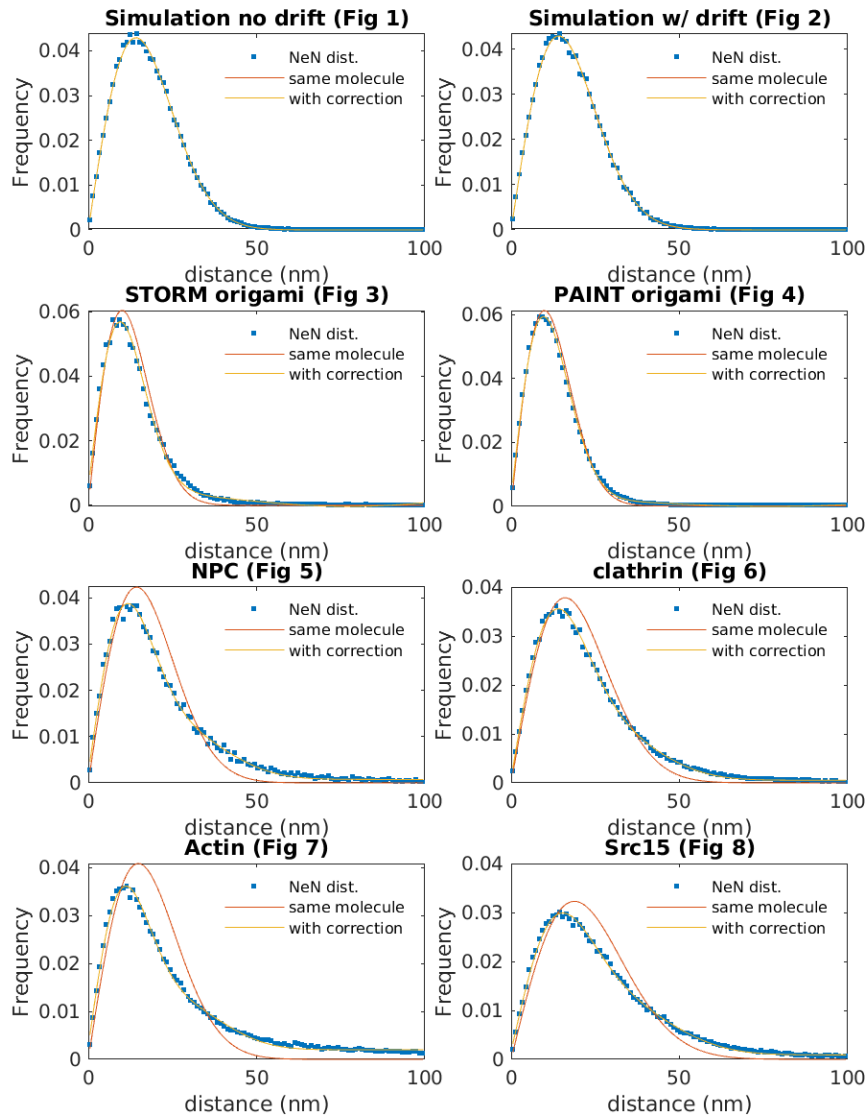
Supplemental Figure 5.16 Resolution calculation with Fourier Ring Correlation (FRC).

FRC curves (81) are presented for all the datasets in the main text. The FRC is computed as a function of spatial frequency and smoothed with the LOESS method. The red line indicates the fixed $\frac{1}{7}$ threshold. The first intersection of the FRC curve with this line yields the resolution $R = \frac{1}{q_{res}}$.




Supplemental Figure 5.17 Comparison of LSF widths (left) and distributions of CRLB localization errors returned from fits (right) for datasets fit without background subtraction.

For the origami samples, the LSF width was also estimated from the segmented images (from loc.) Error bars represent estimates of the standard error obtained through bootstrapping. The LSF widths (σ_{xy}) at the shortest time-interval interrogated are plotted on the right as lines for comparison purposes.




Supplemental Figure 5.18 Localization precision determined by the Nearest Neighbor analysis (NeNA).

Distributions were generated and fit as described in Methods. The “same molecule” fit excludes the correction term that accounts for localizations originating from different molecules.



Supplemental Figure 5.19 Estimates of the autocorrelation of the 2-dimensional LSF, $g_{LSF}(\vec{r}, \tau)$, for the Nuclear Pore Complex data of Figure 5.5.

The first panel shows the weighted average $\langle g_{LSF} \rangle$ over all time-delays τ , with weights determined by the number pairs of localizations observed at each time-delay. Subsequent panels show estimates of $g_{LSF}(\vec{r}, \tau)$ at the indicated time-delay τ . The scale bars are 40 nm.



Supplemental Figure 5.20 Estimates of the weighted average of the autocorrelation of the 2-dimensional LSF $\langle g_{LSF} \rangle$, for the indicated datasets.

Weights are determined by the number of pairs of localizations observed at each time-delay τ . The scale bars are 40 nm.

5.7 Supplementary Note: Derivation of spacetime pair correlation function estimator, and related computations.

N localizations $\mathbf{u}_i = (\vec{r}_i, t_i) = (x_i, y_i, t_i), i = 1, \dots, N$ are observed on a spatial window (region of interest/ROI) W during a temporal window T . This set of points is considered as a realization of a space-time point process X , so that we may define a (first-order) density $\rho(\mathbf{u}) = \rho(\vec{r}, t)$ notionally as

$$\frac{\text{Expected \# of points in area } d\vec{r} \text{ and time-interval } dt \text{ around } (\vec{r}, t)}{d\vec{r} \cdot dt}$$

Or more formally by the following:

$$E \sum_{\mathbf{u} \in X \cap W \times T} 1[\mathbf{u} \in A] = \int_A \rho(\mathbf{u}) d\mathbf{u} \quad (0.9)$$

For any set $A \subset W \times T$, where $1[\cdot]$ is an indicator function, taking the value 1 when its argument is true, and 0 otherwise. For the purposes of this paper, we assume that $\rho = \rho(t)$ is constant in space but may vary in time, e.g. due to bleaching of the fluorophores of the sample.

Further, define the second-order density $\rho^{(2)}(\mathbf{u}_1, \mathbf{u}_2)$ notionally by

$$\frac{\text{Expected \# of pairs of points in the } d\vec{r} \cdot dt \text{ neighborhoods of } \mathbf{u}_1 \text{ and } \mathbf{u}_2 \text{ respectively}}{(d\vec{r} \cdot dt)^2}$$

Or more formally

$$E \sum_{\mathbf{u}_1, \mathbf{u}_2 \in X \cap W \times T} 1[\mathbf{u}_1 \in A \text{ and } \mathbf{u}_2 \in B] = \int_A \int_B \rho^{(2)}(\mathbf{u}_1, \mathbf{u}_2) d\mathbf{u}_2 d\mathbf{u}_1 \quad (0.10)$$

Now $\rho^{(2)}$ describes the second-order properties of X , for example attraction or repulsion between points. It is convenient to normalize $\rho^{(2)}$ so that it is dimensionless and easier to interpret. To that end, define the pair autocorrelation function $g(\mathbf{u}_1, \mathbf{u}_2)$:

$$g(\mathbf{u}_1, \mathbf{u}_2) = \frac{\rho^{(2)}(\mathbf{u}_1, \mathbf{u}_2)}{\rho(\mathbf{u}_1, \mathbf{u}_2)} \quad (0.11)$$

Loosely, the pair autocorrelation function is the ratio of the actual probability of finding points at both \mathbf{u}_1 and \mathbf{u}_2 to the hypothetical probability under the assumption that \mathbf{u}_1 and \mathbf{u}_2 are independent. We typically assume that g is translation invariant in both space and time, and often further assume that it is rotationally invariant in space, so that it only depends on the separation of \mathbf{u}_1 and \mathbf{u}_2 in space and time, and we may write $g(\mathbf{u}_1, \mathbf{u}_2) = g(\|\vec{r}_2 - \vec{r}_1\|, t_2 - t_1)$.

We estimate g using the standard kernel-based framework as laid out in e.g. (86, 311). Specifically, we use a box kernel with bandwidth δ_r in space and δ_t in time, and an isotropic edge-

correction in space, and a density correction for the temporal edge correction, following the approach of (312). Briefly, consider the family of estimators for $g(r, \tau)$ given by:

$$\hat{g}(r, \tau) := \frac{1}{\gamma_{\text{sp.}}(r)\gamma_t(\tau)} \sum_{\mathbf{u}_i, \mathbf{u}_j \in X \cap W \times T}^{\neq} \mathbb{1}[\|\vec{r}_j - \vec{r}_i\| - r < \delta_r/2, |t_j - t_i - \tau| < \delta_t/2] \quad (0.12)$$

We wish to derive functions $\gamma_{\text{sp.}}$ and γ_t such that the resulting estimator is unbiased. The expectation value of the sum in the above expression can be determined from an appropriate Campbell's theorem:

$$\begin{aligned} \mathbb{E}\hat{g}(r, \tau) &= \frac{1}{\gamma_{\text{sp.}}\gamma_t} \int_{W \times T} \int_{W \times T} \rho^{(2)}(\mathbf{u}_1, \mathbf{u}_2) \mathbb{1}[\|\vec{r}_2 - \vec{r}_1\| - r < \delta_r/2, |t_2 - t_1 - \tau| < \delta_t/2] d\mathbf{u}_1 d\mathbf{u}_2 \\ &= \frac{1}{\gamma_{\text{sp.}}\gamma_t} \int_{W \times T} \int_{W \times T} g(\|\vec{r}_2 - \vec{r}_1\|, t_2 - t_1) \rho(t_1) \rho(t_2) \mathbb{1}[\|\vec{r}_2 - \vec{r}_1\| - r < \delta_r/2] \mathbb{1}[|t_2 - t_1 - \tau| < \delta_t/2] d\mathbf{u}_1 d\mathbf{u}_2 \\ &\approx \frac{g(r, \tau)}{\gamma_{\text{sp.}}\gamma_t} \int_W \int_W \mathbb{1}[\|\vec{r}_2 - \vec{r}_1\| - r < \delta_r/2] d\vec{r}_1 d\vec{r}_2 \int_T \int_T \rho(t_2) \rho(t_1) \mathbb{1}[|t_2 - t_1 - \tau| < \delta_t/2] dt_1 dt_2 \end{aligned}$$

Where the approximation in line 3 is due to the assumption that $g(r, \tau)$ is almost constant within $\delta_r/2$ in space and $\delta_t/2$ in time.

From the above derivation, it follows that $\hat{g}(r, \tau)$ is unbiased for the choices

$$\begin{aligned} \gamma_{\text{sp.}}(r) &= \int_W \int_W \mathbb{1}[\|\vec{r}_2 - \vec{r}_1\| - r < \delta_r/2] d\vec{r}_2 d\vec{r}_1 \\ \gamma_t(\tau) &= \int_T \int_T \rho(t_1) \rho(t_2) \mathbb{1}[|t_2 - t_1 - \tau| < \delta_t/2] dt_2 dt_1 \end{aligned}$$

For computational considerations, we make further approximations on $\gamma_{\text{sp.}}$:

$$\begin{aligned} \gamma_{\text{sp.}}(r) &= \int_W \int_W \mathbb{1}[\|\vec{h}\| - r < \delta_r/2] \mathbb{1}[\vec{r}_1 + \vec{h} \in W] d\vec{h} d\vec{r}_1 \\ &= \int_W \int_0^{2\pi} \int \mathbb{1}[|h - r| < \delta_r/2] \mathbb{1}[\vec{r}_1 + (h \cos \theta, h \sin \theta) \in W] h dh d\theta d\vec{r}_1 \\ &= \int_W \int_{r-\delta_r/2}^{r+\delta_r/2} h \int_0^{2\pi} \mathbb{1}[\vec{r}_1 + (h \cos \theta, h \sin \theta) \in W] d\theta dh d\vec{r}_1 \\ &\approx r \delta_r \int_W \int_0^{2\pi} \mathbb{1}[\vec{r}_1 + (r \cos \theta, r \sin \theta) \in W] d\theta d\vec{r}_1 \\ &= r \delta_r \int_0^{2\pi} |\mathcal{W} \cap \mathcal{W}_{-(r \cos \theta, r \sin \theta)}| d\theta \end{aligned}$$

where $|A|$ indicates the area of the set A , and $A_{\vec{h}}$ indicates the translation of the set A by the vector \vec{h} . The first line is a change of variables to $\vec{h} = \vec{r}_2 - \vec{r}_1$, with the extra indicator functions reflecting the integration bounds on \vec{r}_2 , followed by a change to polar coordinates for \vec{h} . The approximation in the fourth line is justified by the fact that angular integral varies slowly with h , so the radial part of the integral can be approximately separated.

For the purposes of our MATLAB code, we represent the spatial window/ROI W as a polygon with vertices $[\text{mask.x}(i), \text{mask.y}(i)]$. We translate the ROI by a vector $[\text{hx}, \text{hy}]$ by simply adding hx and hy to mask.x and mask.y , respectively. Matlab provides functions `polybool` to compute the intersection $W \cap W_{-\vec{h}}$, and `polyarea` to compute the area of the resulting polygon. It remains to complete the angular integral, which we compute by discretizing theta into 32 equally spaced points

$$\gamma_{\text{sp.}}(r) = \frac{2\pi r \delta_r}{32} \sum_{i=1}^{32} \left| W \cap W_{-(r \cos \theta_i, r \sin \theta_i)} \right|, \quad \theta_i = \frac{2\pi i}{32}$$

Chapter 6 Globally Intensity-reweighted Estimators for K - and Pair Correlation Functions

This chapter presents a method for estimating spatial interactions between points of a point pattern such as is produced by a single molecule localization microscopy experiment, under conditions where the probability of detecting a point vary over space, focusing on formal proofs of the validity of the method, and formal analysis of the kinds of bias and variance that might occur in practical experiments. It was written in collaboration with Jesper Møller and Rasmus Waagepetersen, both professors of statistics at Aalborg University in Denmark. I proposed the basic form of the proofs, to which Jesper Møller and Rasmus Waagepetersen added important formal details and improved the presentation. I also wrote the analysis software and performed the simulation study, with guidance from Rasmus Waagepetersen. It has been published in the Australia and New Zealand Journal of Statistics (312).

6.1 Abstract

We introduce new estimators of the inhomogeneous K -function and the pair correlation function of a spatial point process as well as the cross K -function and the cross pair correlation function of a bivariate spatial point process under the assumption of second-order intensity-reweighted stationarity. These estimators rely on a “global” normalization factor which depends on an aggregation of the intensity function, whilst the existing estimators depend “locally” on the intensity function at the individual observed points. The advantages of our new global estimators over the existing local estimators are demonstrated by theoretical considerations and a simulation study.

6.2 Introduction

Functional summary statistics like the nearest-neighbor-, the empty space-, and Ripley's K -function have a long history in statistics for spatial point processes (85, 86, 313). For many years the theory of these functional summary statistics was confined to the case of stationary point processes with consequently constant intensity functions. The paper Baddeley et al. 2000 (88) was

therefore a big step forward since it relaxed substantially the assumption of stationarity in case of the K -function and the closely related pair correlation function.

Baddeley et al (88) introduced the notion of second-order intensity-reweighted stationarity (*soirs*) for a spatial point process. When the pair correlation function g exists for the

point process, *soirs* is equivalent to that g is translation invariant. However, the intensity function does not need to be constant which is a great improvement compared to assuming stationarity, see e.g. (314). When the point process is *soirs*, Baddeley et al further introduced a generalization of Ripley's K -function, the so-called inhomogeneous K -function which is based on the idea of intensity-reweighting the points of the spatial point process, and they discussed its estimation (88). The inhomogeneous K -function has found applications in a very large number of applied papers and has also been generalized e.g. to the case of space-time point processes (89, 90) and to point processes on spheres (315, 316). Moreover, van Lieshout used the idea of intensity-reweighting to generalize the so-called J -function to the case of inhomogeneous point processes (317).

A generic problem in spatial statistics, when just one realization of a spatial process is available, is to separate variation due to random interactions from variation due to a non-constant intensity or mean function. In general, if an informed choice of a parsimonious intensity function model is available for a point process, the intensity function can be estimated consistently. Consistent estimation of the inhomogeneous K -function is then also possible when the consistent intensity function estimate is used to reweight the point process, see e.g. the case of regression models for the intensity function(318). When a parsimonious model is not available, one may resort to non-parametric kernel estimation of the intensity function as considered initially in (88). However, kernel estimators are not consistent for the intensity function and they are strongly upwards biased when evaluated at the observed points. This implies strong bias of the resulting inhomogeneous K -function estimators when the kernel estimators are plugged in for the true intensity.

In this paper, we introduce a new approach to non-parametric estimation of the (inhomogeneous) K - and g -functions for a spatial point process, or of the cross K -function and the cross pair correlation function for a bivariate spatial point process, assuming *soirs* in both cases. This formalizes an approach that has been used to estimate space-time cross pair correlation functions in live-cell single molecule localization microscopy experiments with spatially varying

localization probabilities (53, 83). In the univariate case, our new as well as the existing estimators are given by a sum over all distinct points x and y from an observed point pattern. For the new estimators, each term in the sum depends on an aggregation of the intensity function through a “global” normalization factor $\gamma(y - x)$ instead of depending “locally” on the intensity function at x and at y as for the existing estimators (a similar remark applies in the bivariate case). Intuitively one may expect this to mitigate the problem of using biased kernel estimators of the intensity function in connection to non-parametric estimation of the K -function or pair correlation function. Moreover, to reduce bias when using a non-parametric kernel estimator of γ , we propose a “leave-out” modification of our γ estimator. Our simulation study shows that our new globally intensity reweighted estimators are superior to the existing local estimators in terms of bias and estimation variance regardless of whether the intensity function is estimated parametrically or non-parametrically.

The remainder of the paper is organized as follows. Some background on spatial point processes and notational details are provided in Section 6.3. Section 6.4 introduces our global estimator for the K -function or the cross K -function, discusses modifications to account for isotropy, and compares with the existing local estimators. Section 6.5 is similar but for our new global estimator of the g -function or cross pair correlation function. Section 6.6 describes sources of bias in the local and global estimators when kernel estimators are used, and modifications to reduce bias. In Section 6.8, the global and local estimators of K and g are compared in a simulation study. Possible extensions are discussed in Section 0. Finally, Section 6.10 contains some concluding remarks.

6.3 Preliminaries

We consider the usual setting for a spatial point process X defined on the d -dimensional Euclidean space \mathbb{R}^d , that is, X is a random locally finite subset of \mathbb{R}^d . This means that the number of points from X falling in A , denoted $N(A)$, is almost surely finite for any bounded subset A of \mathbb{R}^d . For further details we refer to (85). In our examples, $d = 2$.

For any integer $n \geq 1$, we say that X has n -th order intensity function $\rho^n: (\mathbb{R}^d)^n \mapsto [0, \infty)$ if for any disjoint bounded Borel sets $A_1, \dots, A_n \subset \mathbb{R}^d$,

$$E\{N(A_1) \cdots N(A_n)\} = \int_{A_1} \cdots \int_{A_n} \rho^{(n)}(x_1, \dots, x_n) dx_1 \cdots dx_n < \infty.$$

By the so-called standard proof we obtain the n -th order Campbell's formula(85): for any Borel function $k: (\mathbb{R}^d)^n \mapsto [0, \infty)$,

$$E \sum_{\substack{\neq \\ x_1, \dots, x_n \in X}} k(x_1, \dots, x_n) = \int \cdots \int k(x_1, \dots, x_n) \rho^{(n)}(x_1, \dots, x_n) dx_1 \cdots dx_n, 6.1$$

which is finite if the left or right hand side is so. Here, \neq over the summation sign means that x_1, \dots, x_n are pairwise distinct.

Throughout this paper, we assume that X has an intensity function ρ and a translation invariant pair correlation function g . This means that for all $x, y \in \mathbb{R}^d$, $\rho^{(1)}(x) = \rho(x)$ and $\rho^{(2)}(x, y) = \rho(x)\rho(y)g(x, y)$, where $g(x, y) = g_0(x - y)$ with $g_0: \mathbb{R}^d \mapsto [0, \infty)$ a symmetric Borel function. If ρ is constant we say that X is (first-order) homogeneous. In particular, if X is stationary, that is, the distribution of X is invariant under translations in \mathbb{R}^d , then ρ is constant and g is translation invariant.

Following (88), the translation invariance of g implies that X is second-order intensity reweighted stationary (soirs) and the *inhomogeneous K-function* (or just *K-function*) is then given by

$$K(t) := \int_{|h| \leq t} g_0(h) dh, \quad t \geq 0.$$

This is Ripley's *K-function* when X is stationary.

Suppose X_1 and X_2 are locally finite point processes on \mathbb{R}^d such that X_i has intensity function ρ_i , $i = 1, 2$, and (X_1, X_2) has a translation invariant *cross pair correlation function* $g_{12}(x_1, x_2) = c(x_1 - x_2)$ for all $x_1, x_2 \in \mathbb{R}^d$. That is, for bounded Borel sets $A_1, A_2 \subset \mathbb{R}^d$ and $N_i(A_i)$ denoting the cardinality of $X_i \cap A_i$, $i = 1, 2$, we have

$$E\{N_1(A_1)N_2(A_2)\} = \int_{A_1} \int_{A_2} \rho_1(x_1)\rho_2(x_2)c(x_1 - x_2) dx_1 dx_2.$$

Then the *cross K-function* is defined by

$$K_{12}(t) := \int_h c(h) dh, \quad t \geq 0.$$

In practice X, X_1, X_2 are observed within a bounded window $W \subset \mathbb{R}^d$, and we use the following notation. The translate of W by $x \in \mathbb{R}^d$ is denoted $W_x := \{w + x \mid w \in W\}$. For a Borel set $A \subseteq \mathbb{R}^d$, $1[x \in A]$ denotes the indicator function which is 1 if $x \in A$ and 0 otherwise. The

Lebesgue measure of A (or area of A when $d = 2$) is denoted $|A|$, and $|x|$ is the usual Euclidean length of $x \in \mathbb{R}^d$.

6.4 Global and local intensity-reweighted estimators for K -functions

6.4.1 The case of one spatial point process

Considering the setting in Section 6.3 for the spatial point process X , we define

$$\gamma(h) := \int_{W \cap W_{-h}} \rho(u) \rho(u+h) du, \quad h \in \mathbb{R}^d. \quad (1)$$

Clearly, γ is symmetric, that is, $\gamma(h) = \gamma(-h)$ for all $h \in \mathbb{R}^d$. We assume that with probability 1, $\gamma(y-x) > 0$ for all distinct $x, y \in X \cap W$. Then, for $t \geq 0$, we can define

$$\widehat{K}_{\text{global}}(t) := \sum_{x, y \in X \cap W}^{\neq} \frac{1[|y-x| \leq t]}{\gamma(y-x)}. \quad (2)$$

If $\gamma(h) > 0$ whenever $|h| \leq t$, then $\widehat{K}_{\text{global}}(t)$ is an unbiased estimator of $K(t)$. This follows from the second-order Campbell's formula:

$$\begin{aligned} E\widehat{K}_{\text{global}}(t) &= \int \int \frac{1[x \in W, y \in W, |y-x| \leq t]}{\gamma(y-x)} \rho(x) \rho(y) g_0(y-x) dx dy \\ &= \int \int \frac{1[x \in W \cap W_{-h}, |h| \leq t]}{\gamma(h)} \rho(x) \rho(x+h) g_0(h) dx dh \\ &= \int_{|h| \leq t} \frac{\gamma(h)}{\gamma(h)} g_0(h) dh. \end{aligned}$$

We call $\widehat{K}_{\text{global}}$ the *global estimator* since it contrasts with one of the estimators suggested in (88): assuming that almost surely $|W \cap W_{y-x}| > 0$ for distinct $x, y \in X \cap W$,

$$\widehat{K}_{\text{local}} := \sum_{x, y \in X \cap W}^{\neq} \frac{1[|y-x| \leq t]}{\rho(x) \rho(y) |W \cap W_{y-x}|}, \quad (3)$$

which we refer to as the *local estimator*. Note that $\widehat{K}_{\text{local}}(t)$ is also an unbiased estimator of $K(t)$ provided $|W \cap W_h| > 0$ for $|h| \leq t$.

In the homogeneous case,

$$\gamma(h) = \rho^2 |W \cap W_{-h}|,$$

whereby $\widehat{K}_{\text{local}} = \widehat{K}_{\text{global}}$, and in the stationary case, these estimators coincide with the translation estimator of Ohser and Stoyan (319).

In practice ρ and hence γ must be replaced by estimates. Estimators of ρ and γ and the bias of these estimators are discussed in Section 6.6.

6.4.2 Modifications to account for isotropy

In addition to soirs, it is frequently assumed that the pair correlation function is isotropic meaning that $g_0(h) = g_1(|h|)$ for some Borel function $g_1: [0, \infty) \mapsto [0, \infty)$. We benefit from this by integrating over the sphere: for $r > 0$, define

$$\gamma^{\text{iso}}(r) := \int_{\mathbb{S}^{d-1}} \gamma(rs) \, d\nu_{d-1}(s) / \zeta_d, \quad (4)$$

where $\mathbb{S}^{d-1} = \{s \in \mathbb{R}^d \mid |s| = 1\}$ denotes the $(d - 1)$ -dimensional unit-sphere, ν_{d-1} is the $(d - 1)$ -dimensional surface measure on \mathbb{S}^{d-1} , and $\zeta_d = 2\pi^{d/2}/\Gamma(d/2)$ is the surface area of the \mathbb{S}^{d-1} . Thus $\gamma^{\text{iso}}(r)$ is the mean value of $\gamma(H)$ when H is a uniformly distributed point on the $(d - 1)$ -dimensional sphere of radius r and center at the origin.

Assuming that almost surely $\gamma^{\text{iso}}(|y - x|) > 0$ for distinct $x, y \in X \cap W$, this naturally leads to another global estimator for K when the pair correlation function is isotropic, namely

$$\widehat{K}_{\text{global}}^{\text{iso}}(t) := \sum_{x, y \in X \cap W}^{\neq} \frac{1[|y - x| \leq t]}{\gamma^{\text{iso}}(|y - x|)}. \quad (5)$$

That $\widehat{K}_{\text{global}}^{\text{iso}}$ is unbiased follows from a similar derivation as for $\widehat{K}_{\text{global}}$: for any $t \geq 0$ such that $\gamma^{\text{iso}}(r) > 0$ whenever $r \leq t$,

$$\begin{aligned} \mathbb{E} \widehat{K}_{\text{global}}^{\text{iso}}(t) &= \int_{|h| \leq t} \frac{\gamma(h)}{\gamma^{\text{iso}}(|h|)} g_0(h) \, dh \\ &= \int_0^t g_1(r) r^{d-1} \int_{\mathbb{S}^{d-1}} \frac{\gamma(rs)}{\gamma^{\text{iso}}(r)} \, d\nu_{d-1}(s) \, dr \end{aligned} \quad (6)$$

$$\begin{aligned} &= \int_0^t g_1(r) \zeta_d r^{d-1} \, dr \\ &= \int_{|h| \leq t} g_1(|h|) \, dh = K(t) \end{aligned} \quad (7)$$

where equations 6 and 7 employ changes of variables to and from polar coordinates, respectively.

When X is homogeneous, equation 5 coincides with the isotropic estimator of Ohser and Stoyan (319). A local estimator of this form can also be defined:

$$\widehat{K}_{\text{local}}^{\text{iso}}(t) := \sum_{x,y \in X \cap W} \frac{1[|y-x| \leq t]}{\rho(x)\rho(y)a_W(|y-x|)}, \quad (8)$$

where

$$a_W(r) = \int_{\mathbb{S}^{d-1}} |W \cap W_{-rs}| \, d\nu_{d-1}(s) / \zeta_d \quad (9)$$

is an isotropized edge correction factor, and where it is assumed that almost surely $a_W(|y-x|) > 0$ for distinct $x, y \in X \cap W$. The local estimator is unbiased when $a_W(r) > 0$ for $r \leq t$.

6.4.3 Comparison of local and global estimators

The global and local estimators of equations 2 and 3 differ in the relative weighting of distinct points $x, y \in X \cap W$. Namely, $\widehat{K}_{\text{local}}$ weights pairs x, y from low-density areas more strongly than those from high-density areas, whilst for $\widehat{K}_{\text{global}}$, the weight only depends on the difference $y-x$. Theoretical expressions for the variances of the global and local K -function estimators are very complicated, not least when the intensity function is replaced by an estimate. This makes it difficult to make a general theoretical comparison of the estimators in terms of their variances. However, under some simplifying assumptions insight can be gained as explained in the following.

Consider a quadratic observation window W of sidelength nm . Then W is a disjoint union of n^2 quadrats W_1, \dots, W_{n^2} each of sidelength m . Assume that the intensity function is constant and equal to ρ_i within each W_i , with ρ naturally estimated by $\rho(u) = \widehat{\rho}_i = N(W_i) / m^2$ for $u \in W_i$. For fixed t and large m , when ρ is replaced by its estimator $\widehat{\rho}$, we can now approximate the local estimator:

$$\begin{aligned} \widehat{K}_{\text{local}}(t) &= \sum_{u,v \in X \cap W}^{\neq} \frac{1[|u-v| \leq t]}{\widehat{\rho}(u)\widehat{\rho}(v)|W \cap W_{u-v}|} \simeq \sum_{i=1}^{n^2} \sum_{u,v \in X \cap W_i}^{\neq} \frac{1[|u-v| \leq t]}{\widehat{\rho}_i^2 |W \cap W_{u-v}|} \\ &\simeq \sum_{i=1}^{n^2} \sum_{u,v \in X \cap W_i}^{\neq} \frac{1[|u-v| \leq t]}{\widehat{\rho}_i^2 |W_i \cap (W_i)_{u-v}| n^2} = \frac{1}{n^2} \sum_{i=1}^{n^2} \widehat{K}_{i,\text{local}}(t), \end{aligned}$$

where $\widehat{K}_{i,\text{local}}$ is the local estimator based on $X \cap W_i$. We use here \simeq in a rather loose sense, meaning that asymptotically, as m tends to infinity, the difference between the two quantities on each side of \simeq tends to zero in a suitable sense (e.g. in mean square) under appropriate regularity conditions. The first approximation above follows because contributions from $u \in X_i$ and $v \in X_j$,

$i \neq j$, are negligible for fixed t and m large, and the second approximation is justified since for $|h| \leq t$, $|W| / |W \cap W_h|$ and $|W_i| / |W_i \cap (W_i)_h|$ will tend to 1 as m increases. Following similar steps, we obtain for the global estimator,

$$\widehat{K}_{\text{global}}(t) \simeq \sum_{i=1}^{n^2} \widehat{K}_{i,\text{local}}(t) \frac{\widehat{\rho}_i^2}{\sum_{l=1}^{n^2} \widehat{\rho}_l^2}.$$

Suppose X is a Poisson process. Note that $\widehat{K}_{\text{local}}(t)$ is an equally weighted average of the $\widehat{K}_{i,\text{local}}(t)$, but since the $\widehat{K}_{i,\text{local}}(t)$ are independent, the optimal weighted average is obtained with weights inversely proportional to the variances of the $\widehat{K}_{i,\text{local}}(t)$. For large m , the variance of $\widehat{K}_{i,\text{local}}(t)$ is well approximated by $2\pi t^2 / (\rho_i^2 m^2)$ (320, 321) and the optimal weights w_i are thus proportional to ρ_i^2 . Our global estimator is obtained from the optimal weighted average by replacing the optimal weights by natural consistent estimates. Hence one may anticipate that the global estimator has smaller variance than the local estimator. In a small-scale simulation study this was indeed the case, and the global estimator with (random) weights proportional to $\widehat{\rho}_i^2$ even had slightly smaller variance than when the optimal fixed weights $w_i \propto \rho_i^2$ were used.

6.4.4 The case of two spatial point processes

For two spatial point processes X_1 and X_2 observed on the same observation window W (cf. Section 6.3), we define the following global estimator for the cross K -function: for $t \geq 0$,

$$\widehat{K}_{12,\text{global}}(t) := \sum_{x \in X_1 \cap W, y \in X_2 \cap W} \frac{1[|y-x| \leq t]}{\gamma_{12}(y-x)}, \quad (10)$$

where

$$\gamma_{12}(h) := \int_{W \cap W_{-h}} \rho_1(u) \rho_2(u+h) du$$

and it is assumed that almost surely $\gamma_{12}(y-x) > 0$ for $x \in X_1 \cap W$ and $y \in X_2 \cap W$. It is straightforwardly verified that $\widehat{K}_{12,\text{global}}(t)$ is unbiased for any $t \geq 0$ such that $\gamma_{12}(h) > 0$ whenever $|h| \leq t$.

The corresponding local estimator is

$$\widehat{K}_{12,\text{local}}(t) := \sum_{x \in X_1 \cap W, y \in X_2 \cap W} \frac{1[|y-x| \leq t]}{\rho_1(x) \rho_2(y) |W \cap W_{y-x}|}, \quad (11)$$

assuming that almost surely $|W \cap W_{y-x}| > 0$ for $x \in X_1 \cap W$ and $y \in X_2 \cap W$. The local estimator is unbiased when $|W \cap W_h| > 0$ for $|h| \leq t$.

Interchanging X_1 and X_2 does not affect equation 10: $\widehat{K}_{12,\text{global}}(t) = \widehat{K}_{21,\text{global}}(t)$ when $\widehat{K}_{21,\text{global}}(t)$ is defined as in equation 10 with γ_{12} replaced by

$$\gamma_{21}(h) := \int_{W \cap W_{-h}} \rho_1(u+h)\rho_2(u) du.$$

This follows since by a change of variable, γ_{12} is symmetric, $\gamma_{21}(h) = \gamma_{12}(-h) = \gamma_{12}(h)$.

When the cross pair correlation function $c(h)$ is also isotropic, additional unbiased estimators of K_{12} are readily obtained in the same way as for the one point process case. Thus, defining

$$\gamma_{12}^{\text{iso}}(r) := \int_{\mathbb{S}^{d-1}} \gamma_{12}(rs) dv_{d-1}(s)/\zeta_d, \quad r \geq 0, \quad (12)$$

and assuming that almost surely $\gamma_{12}^{\text{iso}}(|y-x|) > 0$ for $x \in X_1 \cap W$ and $y \in X_2$, we define an isotropic global estimator by

$$\widehat{K}_{12,\text{global}}^{\text{iso}}(t) := \sum_{x \in X_1 \cap W, y \in X_2 \cap W} \frac{1[|y-x| \leq t]}{\gamma_{12}^{\text{iso}}(|y-x|)}. \quad (13)$$

This is easily seen to be unbiased when $\gamma_{12}^{\text{iso}}(r) > 0$ for $r \leq t$. Finally, the isotropic local estimator is

$$\widehat{K}_{12,\text{local}}^{\text{iso}}(t) := \sum_{x \in X_1 \cap W, y \in X_2 \cap W} \frac{1[|y-x| \leq t]}{\rho_1(x)\rho_2(y)a_W(|y-x|)}, \quad (14)$$

with $a_W(r)$ as defined in Section , and it becomes unbiased if $a_W(r) > 0$ for $r \leq t$.

6.5 Global and local intensity-reweighted estimators for pair correlation functions

6.5.1 The case of one spatial point process

Considering again the setting in Section 6.3 for the spatial point process X , this section introduces global and local estimators for the translation invariant pair correlation function given by g_0 . Note that it may be easier to interpret g_0 than K , but non-parametric kernel estimation of g_0 involves the choice of a bandwidth.

Let $\kappa_b: \mathbb{R}^d \mapsto [0, \infty)$ be a (normalized) kernel with bandwidth $b > 0$, that is, $\kappa_b(h) = \kappa_1(h/b) / b^d$ for $h \in \mathbb{R}^d$, where κ_1 is a probability density function. We assume that κ_1 has support centered in the origin and contained in $[-k, k]^d$ for some $k > 0$; e.g. κ_1 could be a standard d -dimensional normal density truncated to $[-k, k]^d$ (this choice is convenient when W is rectangular with sides parallel to the usual axes in \mathbb{R}^d). Note that the bounded support of κ_b shrinks to $\{0\}$ when b tends to zero. Then, for $h \in \mathbb{R}^d$,

$$\begin{aligned} \mathbb{E} \sum_{x,y \in X \cap W}^{\neq} \kappa_b(h - (y - x)) &= \int_W \int_W \kappa_b(h - (y - x)) \rho(x) \rho(y) g_0(y - x) \, dx \, dy \end{aligned} \quad (15)$$

$$\begin{aligned} &= \int_W \left\{ \int_{W-h-x} \kappa_b(-z) \rho(x) \rho(x + h + z) g_0(h + z) \, dz \right\} \, dx \\ &\simeq g_0(h) \int_W \rho(x) \left\{ \int_{W-h-x} \kappa_b(-z) \rho(x + h + z) \, dz \right\} \, dx \end{aligned} \quad (16)$$

$$\simeq g_0(h) \gamma(h) \quad (17)$$

where $\gamma(h)$ is defined in equation 1. Here, equation 15 follows from the second-order Campbell's formula and \simeq in equations 16 and 17 means that the difference between the quantities on each side of \simeq converges to zero as the bandwidth b tends to zero, under appropriate continuity conditions on ρ and g_0 . The expression in equation 16 is expected to be more accurate but 17 is simpler to compute.

From equation 17 we conclude that $g_0(h)$ can be estimated by the following *global estimator*,

$$\hat{g}_{\text{global}}(h) := \sum_{x,y \in X \cap W}^{\neq} \kappa_b(h - (y - x)) / \gamma(h),$$

provided $\gamma(h) > 0$. This contrasts with the *local estimator*

$$\hat{g}_{\text{local}}(h) := \sum_{x,y \in X \cap W}^{\neq} \frac{\kappa_b(h - (y - x))}{\rho(x) \rho(y) |W \cap W_{x-y}|},$$

which is analogous to the estimator suggested in (88) for an isotropic pair correlation function, see also Section 6.5.2.

6.5.2 Modifications to account for isotropy

For isotropic point processes as defined in Section 6.4.2, the global pair correlation function estimator may be modified to estimate the isotropic pair correlation function given by g_1 : for $r > 0$ such that $\gamma^{\text{iso}}(r) > 0$, define

$$\hat{g}_{\text{global}}^{\text{iso}}(r) := \frac{1}{\zeta_d r^{d-1}} \sum_{x,y \in X \cap W}^{\neq} \frac{\tilde{\kappa}_b(r - |x - y|)}{\gamma^{\text{iso}}(r)}, \quad (18)$$

where for $b > 0$, $\tilde{\kappa}_b(t) = \tilde{\kappa}_1(t/b)/b$, $t \in \mathbb{R}$, for a probability density $\tilde{\kappa}_1: \mathbb{R} \mapsto [0, \infty)$ with support centered at 0 and contained in the interval $[-k, k]$ for some constant $k > 0$, and where $\gamma^{\text{iso}}(r)$ is defined in equation 4. This definition is motivated by the following derivation:

$$\begin{aligned} \mathbb{E} \sum_{x,y \in X \cap W}^{\neq} \tilde{\kappa}_b(r - |y - x|) \\ = \int_W \int_W \tilde{\kappa}_b(r - |y - x|) \rho(x) \rho(y) g_1(|y - x|) dy dx \end{aligned} \quad (19)$$

$$= \int_W \left\{ \int_0^\infty \tilde{\kappa}_b(r - \xi) g_1(\xi) \xi^{d-1} \int_{\mathbb{S}^{d-1}} \rho(x) \rho(x + \xi s) 1[x + \xi s \in W] dv_{d-1}(s) d\xi \right\} dx \quad (20)$$

$$\simeq g_1(r) \zeta_d \gamma^{\text{iso}}(r) r^{d-1} \int_0^\infty \tilde{\kappa}_b(r - \xi) d\xi \quad (21)$$

$$\simeq g_1(r) \zeta_d \gamma^{\text{iso}}(r) r^{d-1}, \quad (22)$$

using the second-order Campbell formula in equation 19, a “shift to polar coordinates” in equation 20, the assumption that b is small in equation 21, and that the kernel is a probability density function in equation 22. Note regarding equation 22 that

$$\int_0^\infty \tilde{\kappa}_b(r - \xi) d\xi = \int_{-\infty}^r \tilde{\kappa}_b(\xi) d\xi,$$

which is not 1 in general. Since $\tilde{\kappa}_b(\xi) = 0$ for $\xi \notin [-bk, bk]$, the integral is 1 if $bk < r$. From equation 22 we obtain the estimator in equation 18.

In the isotropic case the most commonly used local estimators (88) are

$$\hat{g}_{\text{local}}^{\text{iso}}(r) := \frac{1}{\zeta_d r^{d-1}} \sum_{x,y \in X \cap W}^{\neq} \frac{\tilde{\kappa}_b(r - |y - x|)}{\rho(x) \rho(y) |W \cap W_{x-y}|}$$

and

$$\hat{g}_{\text{local}}^{\text{iso}}(r) := \frac{1}{\zeta_d} \sum_{x,y \in X \cap W}^{\neq} \frac{\tilde{\kappa}_b(r - |y - x|)}{\rho(x) \rho(y) |W \cap W_{x-y}| (|y - x|)^{d-1}}$$

assuming that almost surely $|W \cap W_{x-y}| > 0$ for distinct $x, y \in X \cap W$. These estimators suffer from strong positive respectively negative bias for values of r close to 0.

6.5.3 Two point processes

A similar derivation is possible for the cross pair correlation function of a bivariate point process (X_1, X_2) , yielding similar global and local estimators of $c(h)$: for $\gamma_{12}(h) > 0$,

$$\hat{c}_{\text{global}}(h) := \sum_{x \in X_1 \cap W, y \in X_2 \cap W} \kappa_b(h - (y - x)) / \gamma_{12}(h);$$

for $\gamma_{12}^{\text{iso}}(r) > 0$,

$$\hat{c}_{\text{global}}^{\text{iso}}(r) := \frac{1}{\zeta_d r^{d-1}} \sum_{x \in X_1 \cap W, y \in X_2 \cap W} \tilde{\kappa}_b(r - |y - x|) / \gamma_{12}^{\text{iso}}(r);$$

and for $|W \cap W_{x-y}| > 0$ almost surely when $x \in X_1 \cap W$ and $y \in X_2 \cap W$,

$$\hat{c}_{\text{local}}(h) = \sum_{x \in X_1 \cap W, y \in X_2 \cap W} \frac{\kappa_b(h - (y - x))}{\rho_1(x) \rho_2(y) |W \cap W_{x-y}|}$$

and

$$\hat{c}_{\text{local}}^{\text{iso}}(r) = \frac{1}{\zeta_d r^{d-1}} \sum_{x \in X_1 \cap W, y \in X_2 \cap W} \frac{\tilde{\kappa}_b(r - |y - x|)}{\rho_1(x) \rho_2(y) |W \cap W_{x-y}|}.$$

Also an intermediate estimator is possible, with the intensity weighting for one of the processes applied locally, and the other applied globally: with X_1, X_2 , and κ_b as above, we have

$$\begin{aligned} \mathbb{E} \sum_{x \in X_1 \cap W, y \in X_2 \cap W} \frac{\kappa_b(h - (y - x))}{\rho_2(y)} &= \int_W \int_W \kappa_b(h - (y - x)) c(y - x) \rho_1(x) \, dy \, dx \\ &= \int_W \int_{W_{-x-h}} \kappa_b(-z) c(h + z) \, dz \, dx \\ &\simeq c(h) \int_{W \cap W_{-h}} \rho_1(x) \, dx \end{aligned}$$

for a small bandwidth $b > 0$, which suggests the partially-reweighted estimator

$$\hat{c}_{\text{partial}}(h) := \sum_{x \in X_1 \cap W, y \in X_2 \cap W} \frac{\kappa_b(h - (y - x))}{\rho_2(y) \int_{W \cap W_{-h}} \rho_1(x) \, dx'}$$

provided $\int_{W \cap W_{-h}} \rho_1(x) dx > 0$. This estimator may be useful when ρ_2 is much easier to estimate than ρ_1 , e.g. when X_2 is homogeneous.

6.6 Sources of bias when ρ is estimated

All of the estimators of $K(t)$, $K_{12}(t)$, $g_0(h)$, and $g_1(r)$ discussed above are unbiased (at least when $t, |h|, r$ are sufficiently small) when the true intensity function ρ is used to compute the weight functions $\rho(x)\rho(y)$ in the local estimators or $\gamma, \gamma^{\text{iso}}, \gamma_{12}$, or γ_{12}^{iso} in the global estimators. However, in most applications ρ is not known, and must be replaced by an estimate. When the source of inhomogeneity is well understood, it is recommended to fit a model with an appropriate parametric intensity function and use it as the estimate (88, 318).

In the absence of such a model, the most common alternative is a kernel estimator

$$\hat{\rho}(x) := \sum_{y \in X \cap W} \frac{\kappa_\sigma(y - x)}{w_W(x; y)}, \quad (23)$$

where κ_σ is a symmetric kernel on \mathbb{R}^d with bandwidth $\sigma > 0$, and where $w_W(x; y)$ is an appropriate edge correction weight. We take the standard choice from (322),

$$w_W(x; y) = \int_W \kappa_\sigma(u - x) du,$$

see also (323) (other types of edge corrections may depend on both x and y which is why we write $w_W(x; y)$ although the weight here only depends on x .)

In the following we discuss estimators for ρ and γ with particular focus on the implications of estimation bias when kernel estimators are used to replace the true γ or ρ in the global and local estimators.

6.6.1 Bias of local estimators with estimated ρ

We start by considering a single spatial point process X . For each point pair $x, y \in X$ ($x \neq y$), the corresponding term in the local K - and pair correlation function estimators is normalized by the product $\rho(x)\rho(y)$. While an exact expression for the bias of the estimators with estimated ρ is not analytically tractable, we can understand major sources of bias by considering the expression $1 / (\hat{\rho}(x)\hat{\rho}(y))$, which appears in each of the local estimators.

First, following (88), we note that $\hat{\rho}$ as defined in equation 23 is subject to bias when evaluated at the points of X , and that a “leave-one-out” kernel estimator given by

$$\bar{\rho}(x) := \sum_{y \in (X \cap W) \setminus \{x\}} \frac{\kappa_\sigma(y - x)}{w_W(x; y)}, \quad x \in W, \quad (24)$$

is a better choice, with reduced bias in most cases.

Second, we note that

$$E(1/\bar{\rho}(x)) > 1 / E(\bar{\rho}(x))$$

(if $E(1/\bar{\rho}(x))$ exists; in some cases it may be infinite). This follows from Jensen's inequality, since $x \mapsto 1/x$ is strictly convex for $x > 0$. In addition, note that the leading contribution to $E(1/\bar{\rho}(x)) - 1 / E(\bar{\rho}(x))$ is proportional to $\text{Var}\bar{\rho}(x)$ (324). This discrepancy leads to a strong positive bias of the local K - and pair correlation function estimators, especially at large $|y - x|$, where $1 / \bar{\rho}(x)$ and $1 / \bar{\rho}(y)$ are almost independent. This effect becomes more pronounced for smaller σ , since $\text{Var}\bar{\rho}(x)$ typically increases as σ decreases.

Third, we note that for distinct points $x, y \in W$ that are close compared to the bandwidth σ , the covariance of $\bar{\rho}(x)$ and $\bar{\rho}(y)$ leads to bias. For the local (and global) estimators, we consider sums over distinct $x, y \in X \cap W$, which leads us to condition on $x, y \in X$ as follows (325). By X conditioned on distinct points $x, y \in X$ with $\rho^{(2)}(x, y) > 0$, we mean that X is equal to $X_{xy} \cup \{x, y\}$ in distribution, where X_{xy} follows the second-order reduced Palm distribution of X at x, y :

$$P(X \in F \mid x, y \in X) = P(X_{xy} \cup \{x, y\} \in F).$$

Assuming X has n -th order joint intensity functions $\rho^{(n)}$ for $n \leq 4$, X_{xy} has intensity function $\rho_{xy}(u) = \rho^{(3)}(x, y, u) / \rho^{(2)}(x, y)$ and second order joint intensity function $\rho_{xy}^{(2)}(u, v) = \rho^{(4)}(x, y, u, v) / \rho^{(2)}(x, y)$. Now, for distinct $x, y \in W$ with $\rho^{(2)}(x, y) > 0$, neglecting the edge correction in equation 24 for simplicity, we obtain the following by the first and second-order Campbell's formulas for X_{xy} and using that κ_σ is symmetric:

$$\begin{aligned} E[\bar{\rho}(x)\bar{\rho}(y) \mid x, y \in X \cap W] &= E \left\{ \sum_{u \in (X_{xy} \cap W) \cup \{y\}} \kappa_\sigma(x - u) \sum_{v \in (X_{xy} \cap W) \cup \{x\}} \kappa_\sigma(y - v) \right\} \\ &= E \sum_{u, v \in X_{xy} \cap W}^{\neq} \kappa_\sigma(x - u) \kappa_\sigma(y - v) + E \sum_{u \in X_{xy} \cap W} \kappa_\sigma(x - u) \kappa_\sigma(y - u) \\ &\quad + \kappa_\sigma(x - y) \kappa_\sigma(y - x) \end{aligned}$$

$$+\kappa_\sigma(x-y)\mathbb{E} \sum_{v \in X_{xy} \cap W} \kappa_\sigma(y-v) + \kappa_\sigma(y-x)\mathbb{E} \sum_{u \in X_{xy} \cap W} \kappa_\sigma(x-u) \quad (25)$$

$$= \int_W \int_W \kappa_\sigma(x-u)\kappa_\sigma(y-v) \frac{\rho^{(4)}(x,y,u,v)}{\rho^{(2)}(x,y)} du dv \quad (26)$$

$$+ \int_W \kappa_\sigma(x-u)\kappa_\sigma(y-u) \frac{\rho^{(3)}(x,y,u)}{\rho^{(2)}(x,y)} du \quad (27)$$

$$+\kappa_\sigma(x-y)^2 + \kappa_\sigma(x-y) \int_W \{\kappa_\sigma(x-u) + \kappa_\sigma(y-u)\} \frac{\rho^{(3)}(x,y,u)}{\rho^{(2)}(x,y)} du. \quad (28)$$

If X is a Poisson process, then X and X_{xy} are identically distributed, and so the term in equation 26 simplifies to $\mathbb{E}\bar{\rho}(x)\mathbb{E}\bar{\rho}(y)$, which differs from $\rho(x)\rho(y)$ only by the inherent bias of the kernel estimators. In general, the joint intensity $\rho^{(4)}(x,y,u,v)$ in the integrand of that term represents the additional covariance of $\bar{\rho}(x)$ and $\bar{\rho}(y)$ due to interactions between the points of the process, and induces further bias. For example, this bias will tend to overestimate $\rho(x)\rho(y)$ for clustered processes, and lead to an underestimate of K , g_0 , and g_1 . The terms in equations 27 and 28 are non-negative, and in particular the term in equation 27 can be large when x and y are close together compared to σ . This positive bias leads to substantial negative bias at short distances of the local estimators of K , g_0 , and g_1 .

In comparison, the conditional expectation $\mathbb{E}\{\hat{\rho}(x)\hat{\rho}(y) \mid x, y \in X\}$ would have additional positive terms depending on $\kappa(0)$. In the two point process case, the relevant conditional expectation $\mathbb{E}\{\bar{\rho}_1(x)\bar{\rho}_2(y) \mid x \in X_1, y \in X_2\}$ has an expression (of which we omit the details) analogous to equation 27. However, since X_1 and X_2 are assumed to have a cross pair correlation function, almost surely $u = v$ does not occur for $u \in X_1$ and $v \in X_2$, so no term analogous to the second term in equation 27 occurs in $\mathbb{E}\{\bar{\rho}_1(x)\bar{\rho}_2(y) \mid x \in X_1, y \in X_2\}$. This reduces the bias problem in the two point process case compared to the single point process case.

For distinct $x, y \in W$ with $\rho^{(2)}(x,y) > 0$, a superior estimator for $\rho(x)\rho(y)$ might be given by

$$\overline{\rho(x)\rho(y)} := \sum_{u,v \in X \cap W \setminus \{x,y\}}^{\neq} \frac{\kappa_\sigma(x-u)\kappa_\sigma(y-v)}{w_W(x;u)w_W(y;v)}. \quad (29)$$

Then the terms in equations 27 and 28 are avoided, since

$$E\left\{\overline{\rho(x)\rho(y)}\middle|x, y \in X \cap W\right\} = \int_W \int_W \frac{\kappa_\sigma(x-u)\kappa_\sigma(y-v)}{w_W(x;u)w_W(y;v)} \frac{\rho^{(4)}(x,y,u,v)}{\rho^{(2)}(x,y)} du dv.$$

We do not investigate this idea further in the current work.

6.6.2 Bias of global estimators with estimated γ

Given the kernel estimate in equation 23 an immediate estimator of $\gamma(h)$, $h \in \mathbb{R}^d$, is

$$\hat{\gamma}(h) := \int_{W \cap W_{-h}} \hat{\rho}(z)\hat{\rho}(z+h) dz. \quad (30)$$

To understand properties of this estimator we evaluate its expected value. We start with the simplest case where h is a fixed vector in \mathbb{R}^d . This case is relevant for the global estimator of the pair correlation function. We return in the end of this section to the case where h is an observed difference $h = y - x$ for distinct $x, y \in X$, which occurs for the global estimator of the K -function.

Neglecting edge corrections for simplicity, we get

$$E\hat{\gamma}(h) = \int_{W \cap W_{-h}} \int_W \kappa_\sigma(z-u)\rho(u) \int_W \kappa_\sigma(z+h-v)\rho(v)g_0(u-v) dv du dz \quad (31)$$

$$+ \int_{W \cap W_{-h}} \int_W \kappa_\sigma(z-u)\kappa_\sigma(z+h-u)\rho(u) du dz. \quad (32)$$

The two resulting terms are analogous to the terms in equations 26 and 27.

When $g_0 = 1$ as for a Poisson process, the term in the right hand side of equation 31 simplifies to

$$\int_{W \cap W_{-h}} E\hat{\rho}(x) E\hat{\rho}(x+h) dx.$$

This differs from $\gamma(h)$ due to the inherent bias of the kernel estimators which depends on the spatial structure of the intensity function: $E\hat{\rho}(x) - \rho(x)$ becomes large when σ is large compared to the length scale of spatial variation of $\rho(x)$. On the other hand, when $g_0 \neq 1$, the term in the right hand side of equation 31 includes an additional bias due to the interaction between points. For example, this bias will tend to overestimate γ for clustered processes, and therefore lead to an underestimate of K or the pair correlation function. This interaction bias is most pronounced when σ is small. In particular, as $\sigma \rightarrow 0$, this term approaches $g_0(y-x)\gamma(y-x)$, so that e.g. $E\hat{g}_{\text{global}}(h) \rightarrow 1$ for all $h \in \mathbb{R}^d$. However, in the typical case where the strength of pairwise interactions decreases with distance, increasing σ reduces bias due to interactions. Therefore, it is important to choose σ to be larger than the length-scale of interesting correlations.

The term in equation 32, though, is always positive when $h / 2$ is in the support of κ_σ . We can avoid this term by using the following “leave-out” estimator

$$\bar{\gamma}(h) = \int_{W \cap W_{-h}} \sum_{\substack{\neq \\ u, v \in X \cap W}} \frac{\kappa_\sigma(z - u) \kappa_\sigma(z + h - v)}{w_W(z; u) w_W(z + h; v)} dz, \quad (33)$$

where leave-out refers to omitting “diagonal terms” $u = v$ in $\hat{\rho}(z) \hat{\rho}(z + h)$ (with $u, v \in X \cap W$).

Similarly, when X is isotropic, an estimator of γ^{iso} can be defined in terms of $\bar{\gamma}$, as

$$\bar{\gamma}^{\text{iso}}(r) := r^{d-1} \int_{\mathbb{S}^{d-1}} \bar{\gamma}(rs) dv_{d-1}(s). \quad (34)$$

For the global K -function estimators, γ is evaluated at $y - x$ for distinct $x, y \in X \cap W$. In this case the relevant expectation is $E\{\bar{\gamma}(y - x) \mid x, y \in X\}$. As in Section 6.6.1 we obtain this by considering the second-order reduced Palm distribution at distinct $x, y \in W$ with $\rho^{(2)}(x, y) > 0$, by assuming that X has n -th order intensity functions $\rho^{(n)}$ for $n \leq 4$, and by neglecting the edge corrections for simplicity:

$$\begin{aligned} E\{\bar{\gamma}(y - x) \mid x, y \in X\} = & \int_{W \cap W_{-(y-x)}} \left(\int_W \int_W \frac{\rho^{(4)}(x, y, u, v)}{\rho^{(2)}(x, y)} \kappa_\sigma(z - u) \kappa_\sigma(z + (y - x) - v) du dv + \kappa_\sigma(z - x)^2 \right. \\ & + \kappa_\sigma(z - y) \kappa_\sigma(z + y - 2x) \\ & + \int_W \frac{\rho^{(3)}(x, y, u)}{\rho^{(2)}(x, y)} [\{\kappa_\sigma(z - x) + \kappa_\sigma(z - y)\} \kappa_\sigma(z + (y - x) - u) \\ & \left. + \kappa_\sigma(z - u) \{\kappa_\sigma(z - x) + \kappa_\sigma(z + y - 2x)\}] du \right) dz. \end{aligned}$$

Again, in case of a Poisson process, $\rho^{(4)}(x, y, u, v) / \rho^{(2)}(x, y) = \rho(u) \rho(v)$ and the first term is approximately $\gamma(y - x)$, subject to the subtleties discussed above. The other three terms are related to the terms with $u, v \in \{x, y\}$ of the double sum of equation 33, and yield a positive bias. We expect this bias to be small when σ is reasonably small, since the excess terms become negligible far from x and y , and the integral is over all of $W \cap W_{-h}$. The three terms could be avoided by considering the further modified “leave-one pair-out” estimator

$$\tilde{\gamma}(h; x, y) = \int_{W \cap W_{-h}} \sum_{\substack{\neq \\ u, v \in (X \cap W) \setminus \{x, y\}}} \frac{\kappa_\sigma(z - u) \kappa_\sigma(z + h - v)}{w_W(z; u) w_W(z + h; v)} dz, \quad \text{with } h = y - x,$$

but this depends on (x, y) not only through $h = y - x$ which precludes the use of interpolation schemes as discussed in Section 6.7.

In case of two point processes we just use

$$\hat{\gamma}_{12}(h) = \int_{W \cap W_{-h}} \hat{\rho}_1(z) \hat{\rho}_2(z + h) dz,$$

for kernel estimators $\hat{\rho}_1$ and $\hat{\rho}_2$, since in this case almost surely there are no diagonal terms $u = v$ in $\hat{\rho}_1(z) \hat{\rho}_2(z + h)$ (with $u \in X_1$ and $v \in X_2$).

6.7 Computation of γ and γ^{iso}

We compute $\gamma(h)$ for a given intensity function ρ using a simple Monte Carlo integration algorithm: we generate uniform random samples $U_i, i = 1, \dots, n$, on $W \cap W_{-h}$ and approximate $\gamma(h)$ by the unbiased Monte Carlo estimate

$$\gamma_{\text{MC}}(h) = \frac{|W \cap W_{-h}|}{n} \sum_{i=1}^n \rho(U_i) \rho(U_i + h). \quad (35)$$

To achieve a desired precision, we consider the standard error $\sigma_{\text{MC}} / \sqrt{n}$ of $\gamma_{\text{MC}}(h)$ and choose n so that the coefficient of variation becomes less than a selected threshold α : $\sigma_{\text{MC}} / (\sqrt{n} \mu_{\text{MC}}) < \alpha$. For the simulation studies in Section 6.8, we used $\alpha = .001$ or $\alpha = .005$. In practice, we wish to evaluate γ at many values of h . Thus it is convenient to generate a single sequence of random samples $V_j, j = 1, \dots, n'$ on W , and for each h use a subsequence $\{U_i^{(h)}\} = \{V_j \mid V_j \in W \cap W_{-h}\}$. We choose n' sufficiently large to produce the requisite length of sub-sequence for each h .

For $\gamma^{\text{iso}}(r)$, we follow a similar approach, generating also random independent s_i uniformly on $\{s \mid s \in \mathbb{S}^{d-1}, U_i + rs \in W\}$, and computing

$$\gamma_{\text{MC}}^{\text{iso}} = \frac{\int_{\mathbb{S}^{d-1}} |W \cap W_{-rs}| dv_{d-1}(s)}{\zeta_d n} \sum_{i=1}^n \rho(U_i) \rho(U_i + rS_i). \quad (36)$$

The integral $\int_{\mathbb{S}^{d-1}} |W \cap W_{-rs}| dv_{d-1}(s)$ is easy to compute when W is a rectangular window. As above, U_i and s_i are typically generated for each r as appropriate subsequences of shared larger sequences V_j and t_j , respectively, sampled uniformly on W and \mathbb{S}^{d-1} , respectively.

In practice ρ is replaced by an estimate. Then for the kernel-based leave-out estimator given by equation 33, $\rho(U_i) \rho(U_i + h)$ in equation 35 is replaced by

$$\sum_{u,v \in X \cap W}^{\neq} \frac{\kappa_{\sigma}(U_i - u)\kappa_{\sigma}(U_i + h - v)}{w_W(z; u)w_W(z + h; v)},$$

which is evaluated using a fast routine written in C. In a similar way, when X is isotropic and equation 34 is used, $\rho(U_i)\rho(U_i + rs_i)$ in equation 36 is replaced by a double sum.

Since γ and γ^{iso} are quite smooth, it is possible to interpolate them very accurately based on a moderate number of points h_j or r_j . This is especially helpful for γ^{iso} because it is one-dimensional. For the kernel-estimated $\bar{\gamma}^{\text{iso}}$ or $\hat{\gamma}^{\text{iso}}$, we find that linear interpolation based on sample spacing of $|r_{j+1} - r_j| < \sigma / 10$ gives estimates within .01% of the true values. The interpolation scheme is especially helpful for the K -functions as the number of points grows large, in which case we must evaluate γ (or γ^{iso} in the isotropic case) at a very large number of pairs of points.

The proposed Monte Carlo computation becomes very slow when especially precise coefficient of variation α is desired, or when using kernel-based estimates with very small kernel bandwidth σ or large number of points N . For these cases, it may be beneficial to apply a variance reduction technique such as antithetic variables, or to consider an approximate convolution based on discrete Fourier transforms, with a kernel-based estimate of ρ , when desired, based on quadrat counts. When the side length of the quadrats is much less than σ , we expect this method to produce accurate estimates of γ (or γ^{iso} in the isotropic case).

6.8 Simulation study

To compare global and local estimators for K and g , we simulated 100 point patterns on the unit square $W = [0,1]^2$ for each of nine point process models obtained by combining three different types of point process interactions with four types of intensity functions. For plots of estimated K or g we simulated a further 1000 point patterns of the considered point process model.

More specifically we simulated stationary point processes of the types Poisson (no interaction), log-Gaussian Cox (LGCP – these are clustered/aggregated, see (326)), and determinantal (DPP – these are regular/repulsive, see (327)), and subsequently subjected them to independent thinning to obtain various types of intensity functions. Note that independent thinnings of stationary point processes are soirs (88). The intensities of the stationary point

processes were adjusted to obtain on average 200 or 400 points in the simulated point patterns (that is, after independent thinning).

For the Gaussian random field underlying the LGCP we used an exponential covariance function with unit variance and correlation scale 0.05 resulting in the isotropic pair correlation function

$$g_{\text{LGCP}}(r) = \exp\{\exp(-r / .05)\}.$$

For the DPP we used a Gaussian kernel with scaling parameter $\alpha = 0.02$ leading to

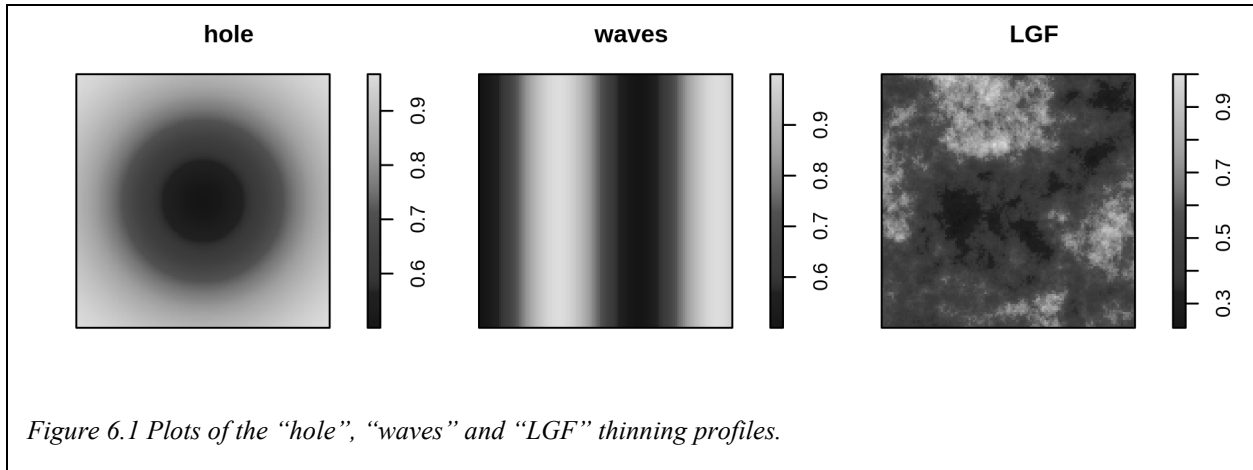
$$g_{\text{DPP}}(r) = 1 - \exp\{-2(r/.02)^2\}.$$

The intensity functions were of type constant (no thinning), “hole”, “waves”, or log-Gaussian random field (“LGF”). Intensity functions of the “hole” and “waves” types were obtained by independent thinning using spatially varying retention probabilities

$$\begin{aligned} p_{\text{hole}}(x, y) &= 1 - .5 \exp[-\{(x - .5)^2 + (y - .5)^2\} / .18], \\ p_{\text{waves}}(x, y) &= 1 - .5 \cos^2(5x), \\ p_{\text{LGF}}(x, y) &= \lambda(x, y) / \sup_{(u,v) \in W} \lambda(u, v), \end{aligned}$$

for $(x, y) \in [0,1]^2$. In case of “LGF”, $\log \lambda$ was generated as a realization of a Gaussian random field with exponential covariance function, with variance .1 and correlation scale .3. The resulting “LGF” retention probability surface is much less smooth than for “hole” and “waves” but similar to “hole” and “waves” in terms of intensity contrast and spatial separation of high-intensity and low-intensity regions. The surfaces of retention probabilities are shown in Figure 6.1.

Simulations were carried out and analyzed using the R package `spatstat`, and a new package `globalKinhom` that implements the global K - and pair correlation function estimators using Monte-Carlo estimates of γ as described in Section 6.7 (328–330). In most cases we set the precision of the Monte-Carlo estimates to $\alpha = .005$. When probability intervals and root integrated mean square error (RIMSE) values are shown, we use $\alpha = .001$ instead, where the more precise calculation produced slightly smaller RIMSE values. We also tested smaller values of α in a few particular cases, and did not observe any reduction in RIMSE values below $\alpha = .001$. We do not show simulation results for all scenarios since in many cases the different scenarios led to qualitatively similar conclusions.



To investigate our cross K and cross pair correlation function estimators we generated simulations from a bivariate LGCP detailed in Section 6.8.2.

6.8.1 Estimation of K - and pair correlation functions

We initially compare the bias of global and local estimators of the K -function using in both cases kernel estimators of the intensity function obtained with a Gaussian kernel with bandwidth σ chosen by the method of Cronie & van Lieshout(331), as implemented in the spatstat procedure `bw.CvL` (CvL for convenience in the following). The selected bandwidths vary around .05 (see third column in Table 6.1), with slightly larger bandwidths for LGCP than for Poisson and DPP. For the global estimator we consider the isotropic estimator from equation 5, since the pair correlation functions of the point processes tested here are all isotropic, as in the setting of Section 6.4.2, and the estimation of γ^{iso} is less computationally intensive than that of γ . We consider both the estimator of equation 30 and the leave-out estimator of equation 33 of the function γ . Similarly we also consider the local estimator using either the original kernel estimator of equation 23 or the leave-out estimator of equation 24 as suggested in (88).

For better visualization of the simulation results we transform the K -function estimators into estimators of the so-called $\{L(r) - r\}$ -function via the one-to-one transformation

$$L(r) - r = \sqrt{K(r) / \pi} - r.$$

We only show results in case of the waves intensity function with on average 400 simulated points, since the results for the other intensity functions and with on average 200 simulated points give the same qualitative picture.

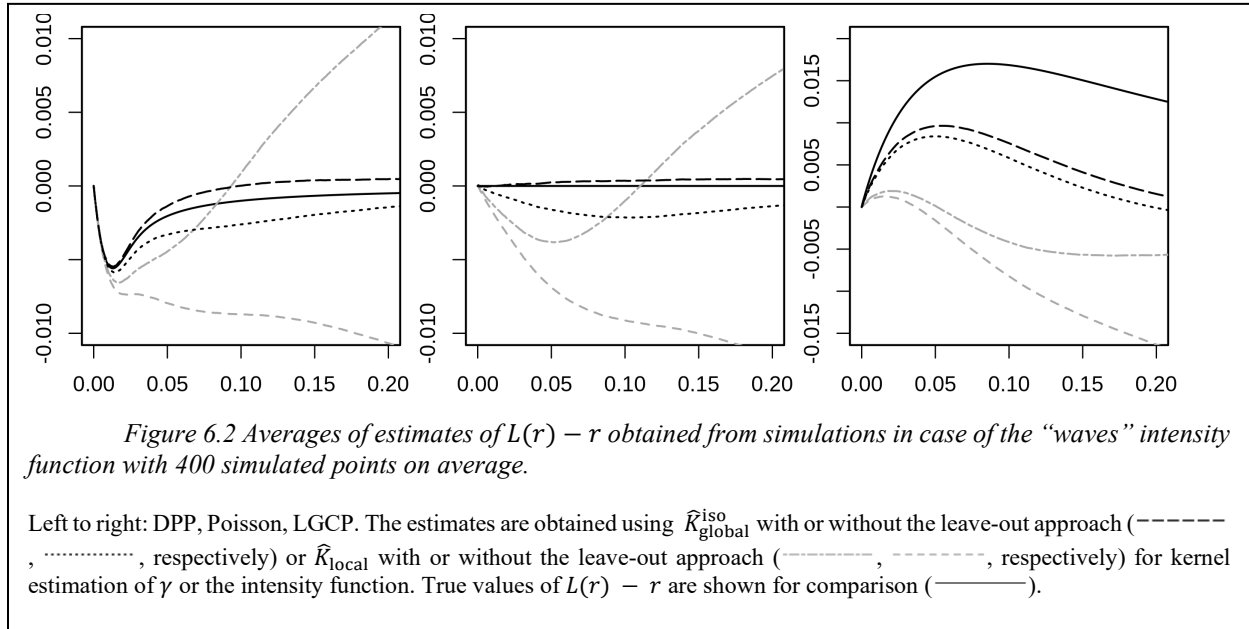


Figure 6.2 shows averages of the simulated estimates and it is obvious that the global estimators are much less biased than the local estimators. It is clearly advantageous to use the leave-out versions for the global estimator. The leave-out approach is also advantageous for the local estimator, at least for small distances r . The biases of the leave-out local estimator are as discussed in Section 6.6.1: strong negative bias at short distances due to the covariance of $\bar{\rho}(x)$ and $\bar{\rho}(y)$, and strong positive bias at large distances due to Jensen's inequality $E(1/\bar{\rho}) > 1/E(\bar{\rho}(x))$. The leave-out global estimator appears to be close to unbiased in case of DPP and Poisson but is too small on average in case of LGCP.

There exist a number of alternatives to the CVL approach to choosing the bandwidth for the kernel estimation. We therefore also investigate bias in the case where the bandwidth is selected using the likelihood cross validation (LCV) method implemented in the spatstat procedure `bw.ppl`. Results regarding the LCV selected bandwidths are summarized in the fourth column of Table 6.1. Comparison of the CVL and LCV results in Table 6.1 shows that the LCV approach tends to select considerably larger bandwidths σ than the CVL method for the DPP and Poisson process, and somewhat smaller σ for the LGCP.

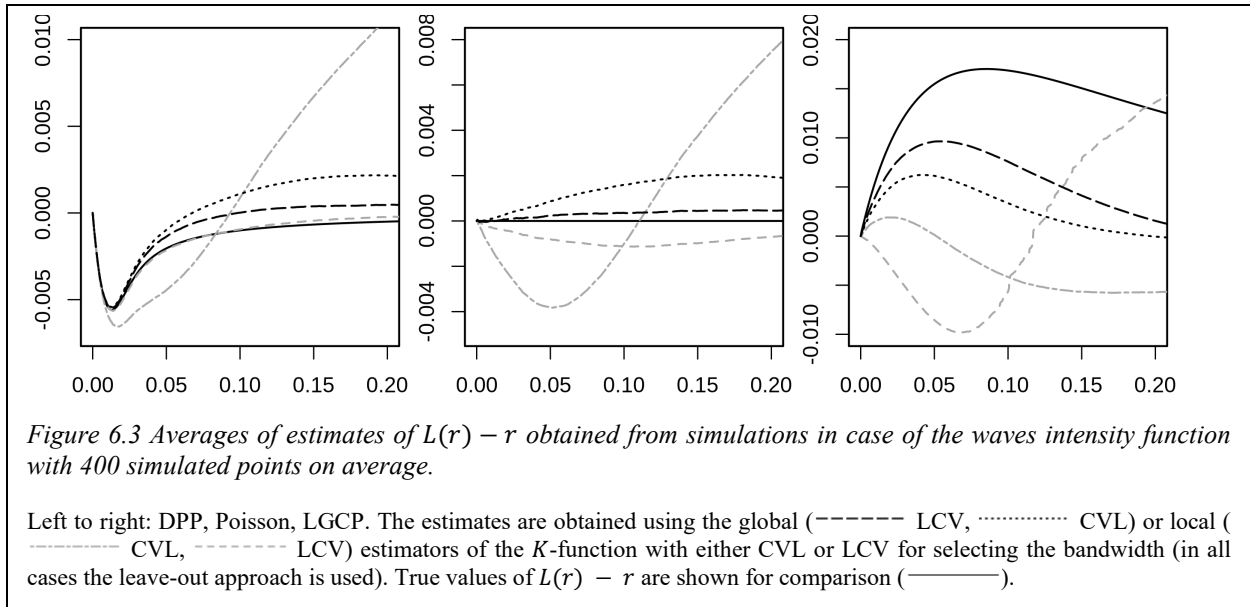
Figure 6.3 compares averages of the global and local estimators using either of the two approaches to bandwidth selection and with leave-out in all cases. Again we show only results for the waves intensity function and expected number of points equal to 400.

The bias of the estimators is quite sensitive to the choice of bandwidth selection method. In case of DPP and Poisson, the global estimator using CVL and the local estimator using LCV perform similarly with the global estimator a bit more biased than the local for DPP and vice versa for Poisson. The global estimator performs slightly worse when combined with LCV than with CVL, likely due to the inherent biases of the kernel estimator $\bar{\rho}$, which become more pronounced as σ increases. The local estimator with CVL is strongly biased for almost all r considered. The improved performance with LCV is likely due to the reduced variances and covariances for $\bar{\rho}$ when a larger bandwidth is used. This also explains the strong bias of the local estimator with LCV for the LGCP, since σ_{LCV} is typically smaller than σ_{CVL} in that case. The global estimator for the LGCP has the smallest bias with the CVL method and has much less bias than the local estimator regardless of whether CVL or LCV is used. It is not surprising that the LGCP is the most challenging case for both the global and local estimators, since the random aggregation of the LGCP tends to be entangled with the variation in the intensity function.

We finally compare the sampling variability of the leave-out global estimator using CVL and the leave-out local estimator using LCV. Figure 6.4 shows 95% pointwise probability intervals and averages for the two estimators, again with 400 simulated points on average and the “waves” intensity function, and Table 6.2 gives root integrated mean square error (RIMSE) values for the K -function estimators applied to each process, for each combination of CVL or LCV with the local

Table 6.1. Mean (\pm st. dev.) of CVL and LCV bandwidths, for each type of spatial point process we considered. The expected number of points for each listed process is 400.

Interaction type	Intensity function	σ_{CVL}	σ_{LCV}
DPP	constant	0.046 (0.005)	0.63 (0.15)
	hole	0.045 (0.004)	0.33 (0.22)
	waves	0.048 (0.004)	0.28 (0.25)
	LGF	0.047 (0.005)	0.22 (0.16)
Poisson	constant	0.047 (0.006)	0.59 (0.21)
	hole	0.048 (0.007)	0.29 (0.23)
	waves	0.050 (0.006)	0.14 (0.11)
	LGF	0.050 (0.006)	0.17 (0.13)
LGCP	constant	0.066 (0.009)	0.040 (0.007)
	hole	0.064 (0.012)	0.044 (0.008)
	waves	0.071 (0.011)	0.042 (0.008)
	LGF	0.066 (0.011)	0.042 (0.007)



or global leave-out estimator. Figure 6.4 indicates that the global estimator has smaller variance than the local estimator. This should also result in smaller mean square error for Poisson and LGCP where the bias is also smallest for the global estimator. For DPP the picture is not completely clear regarding mean square error since in this case the global estimator has larger bias than the local estimator. Table 6.2 gives more insight where a first observation is that the leave-out local estimator is very sensitive to the choice of bandwidth selection method with LCV performing much better than CVL for DPP and Poisson and vice versa for LGCP. The leave-out global estimator is much less sensitive to choice of bandwidth selection method. Best results in terms of RIMSE are obtained with the leave-out global estimator combined with CVL.

Figure 6.5 shows averages of leave-out global and local estimators of the isotropic pair correlation function using either CVL or LCV in case of the wave intensity with 400 points on average. Once again, local estimators are most strongly biased with the bandwidth selection method that produces the smaller bandwidth: CVL for the DPP and Poisson processes, and LCV for the LGCP. The bias is small to moderate for the global estimators with largest bias in case of LGCP. For the DPP and Poisson case positive bias of the local and global estimator occurs for very small distances.

6.8.2 Estimation of cross K - and cross pair correlation functions

To investigate the cross K and cross pair correlation function estimators, we simulated 100 bivariate point patterns for each model of a bivariate point process (X_1, X_2) , where either X_1 and




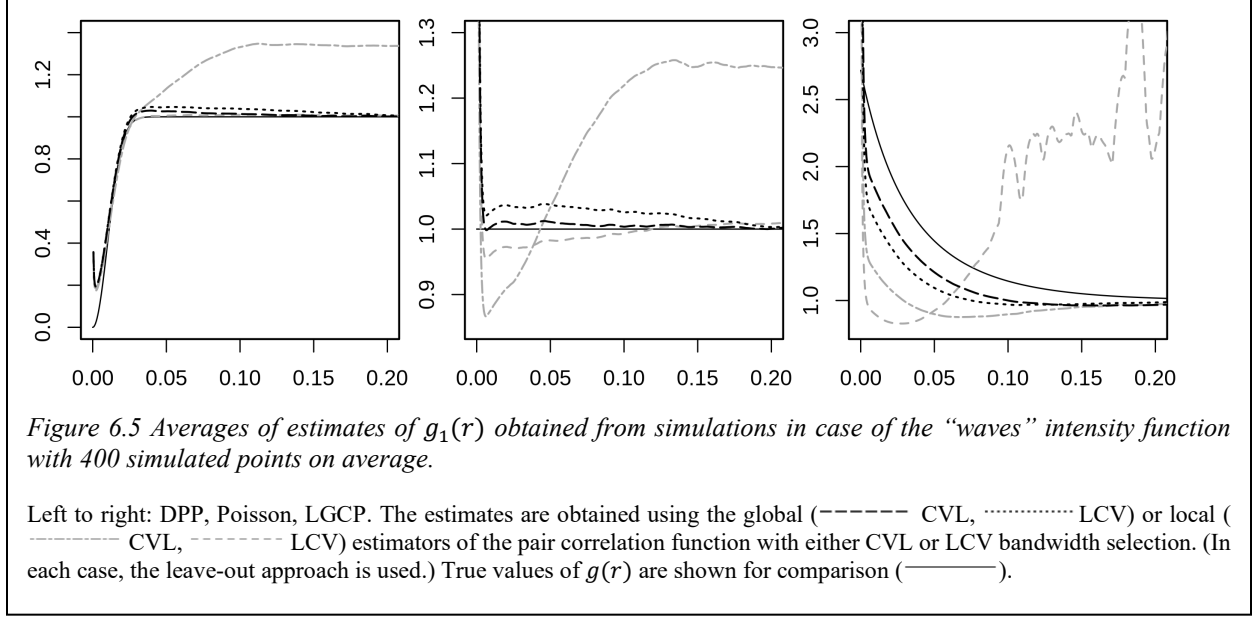
Figure 6.4 Averages and 95 % pointwise probability intervals for estimates of $L(r) - r$ in case of the waves intensity function with 400 simulated points on average.

Left to right: DPP, Poisson, LGCP. The estimators used are the leave-out global estimator using CVL (-----) and the leave-out local estimator using LCV (-·-·-·-) with pointwise probability intervals shown in like shade. True values of $L(r) - r$ are also shown (———).

X_2 are independent or display segregation or co-clustering. Processes that were chosen for plotting were simulated an additional 1000 times. Inhomogeneous intensity functions were subsequently obtained using independent thinning of stationary bivariate point processes, where the two point processes have the same intensity, and the constant, “hole”, and “waves” retention probabilities p

Table 6.2 $RIMSE \times 10^2$ of local and global K -function estimators with CVL and LCV bandwidths.

Interaction	Intensity	\hat{K}_{local}		\hat{K}_{global}	
		CVL	LCV	CVL	LCV
DPP	flat	0.59	0.069	0.029	0.060
	hole	0.64	0.107	0.031	0.128
	waves	0.60	0.052	0.049	0.121
	LGF	0.59	0.060	0.050	0.110
Poisson	flat	0.45	0.083	0.028	0.069
	hole	0.45	0.120	0.034	0.103
	waves	0.40	0.061	0.037	0.093
	LGF	0.37	0.087	0.050	0.089
LGCP	flat	0.89	0.999	0.573	0.628
	hole	0.87	1.554	0.576	0.636
	waves	0.89	1.146	0.528	0.613
	LGF	0.90	1.506	0.542	0.625



as described in connection to Figure 6.6 were used. This implies $\rho_1(x) = \rho_2(x)$ for $x \in [0,1]^2$ (we did not investigate any scenarios where $\rho_1 \neq \rho_2$).

In the case of independence, X_1 and X_2 are independent Poisson processes. For the dependent cases, we considered a bivariate LGCP. Specifically, for $i = 1,2$, X_i has random intensity function

$$\Lambda_i(u) = p(u) \exp\{\mu_i + \alpha_i Y(u) + \beta U_i(u)\}, \quad i = 1,2,$$

where Y , U_1 , and U_2 are independent zero-mean unit-variance Gaussian random fields with isotropic exponential correlation functions given by $\exp(-r/\phi)$ and $\exp(-r/\psi_i)$ ($r \geq 0$), $i = 1,2$, respectively, and where $\mu_i \in \mathbb{R}$, $\alpha_i \in \mathbb{R}$, and $\beta > 0$ are parameters. This means that X_1 and X_2 conditioned on (Λ_1, Λ_2) are independent Poisson processes with intensity functions Λ_1 and Λ_2 , respectively. The (cross) pair correlation functions for this class of bivariate LGCP are isotropic, where the pair correlation function of X_i is given by

$$g_i^{\text{iso}}(r) = \exp\{\alpha_i^2 \exp(-r/\phi) + \beta \exp(-r/\psi_i)\}, \quad i = 1,2,$$

and the cross pair correlation function of (X_1, X_2) is given by

$$c^{\text{iso}}(r) = \exp\{\alpha_1 \alpha_2 \exp(-r/\phi)\}.$$

Note that $c^{\text{iso}} < 1$ if $\alpha_1 \alpha_2 < 0$ (the case of segregation between X_1 and X_2), and $c^{\text{iso}} > 1$ if $\alpha_1 \alpha_2 > 0$ (the case of co-clustering between X_1 and X_2). For the segregated processes, we chose $\alpha_1 = -\alpha_2 = 1$, $\phi = .03$, $\beta = .25$, $\psi_1 = .02$, and $\psi_2 = .01$. For the co-clustered case, we used

$\alpha_1 = \alpha_2 = 1$ and the other parameters as for the segregated case. With these choices, the cross pair correlation functions become

$$c_{\text{segr}}^{\text{iso}}(r) = \exp\{-\exp(-r/.03)\}$$

for the segregation case and

$$c_{\text{cluster}}^{\text{iso}}(r) = \exp\{\exp(-r/.03)\}$$

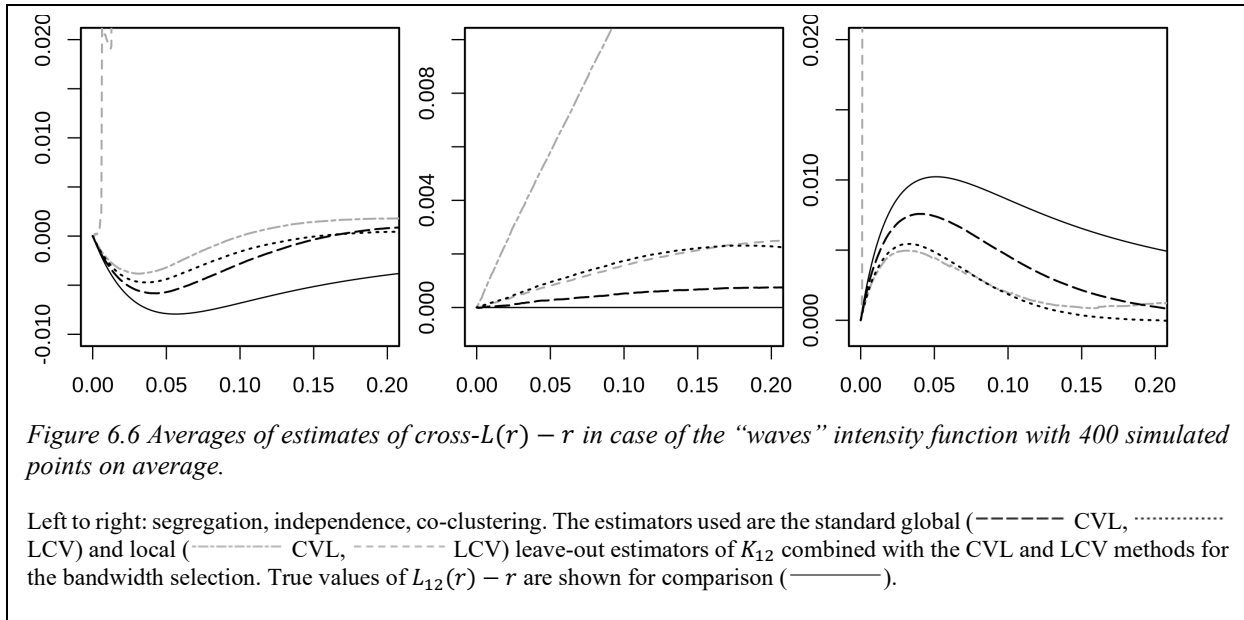
for the co-clustered case. Finally, we adjusted μ_1 and μ_2 so that the expected number of points after independent thinning is 200 or 400.

For the global estimator of K_{12} , we consider again the isotropic estimator (equation 13), since in each case the cross pair correlation function is isotropic, and estimation of $\gamma_{12}^{\text{iso}}(r)$ is less computationally intensive than that of $\gamma_{12}(h)$. For the local estimator we consider the estimator from equation 11, with ρ_i estimated by the leave-out kernel estimator $\bar{\rho}$ from equation 24. Similar to the $\{L(r) - r\}$ -function used above, we transform the K_{12} -function estimators into estimators of the $\{L_{12}(r) - r\}$ -function, by the one-to-one transformation

$$L_{12}(r) - r = \sqrt{K_{12}(r) / \pi} - r.$$

Figure 6.6 shows averages of estimators of $L_{12}(r) - r$ in case of the waves intensity and expected number of points equal to 400. The bandwidth is selected using the CVL or LCV procedure applied to X_1 . Table 6.3 gives selected bandwidth values for the pairs of spatial point processes we considered. The results are similar to the one point process case. Both the segregated and co-clustered LGCP typically yield $\sigma_{\text{LCV}} < \sigma_{\text{CVL}}$ while the opposite is true for the Poisson case. Further, the local estimators are strongly biased, and the bias increases as the bandwidth σ decreases: in the case of segregation and co-clustering, the local estimators are better with CVL, while LCV is better in the case of independence. Note also that the negative bias that is observed at small distances r for \hat{K}_{local} is absent here as predicted in the discussion in Section 6.6.1. The bias for the global estimator with CVL is smaller than for the best local estimators in each case.

To compare sampling variability for the estimators of the cross K -function, we show pointwise 95% probability intervals for estimated $L_{12}(r) - r$ in Figure 6.7. The bandwidth selection method that produces the least bias in each case is shown. Table 6.4 shows root integrated mean square error of the estimators of K_{12} . In every case, the best global estimator has smaller integrated mean square error than the best local estimator, as expected from the considerations of Section 6.4.3.




For the estimation of the cross pair correlation functions, the conclusions are similar to those for the cross K -functions, see Figure 6.8. The average of the global estimator is quite close to the true cross pair correlation function, while the local estimator is strongly biased. Note that $\hat{c}_{\text{local}}^{\text{LCV}}$ is missing for the segregated and co-clustered processes, because the average values of that estimator were extremely large.

6.8.3 Estimation of K -function using a parametric estimate for ρ

Returning to the setting of a single point process X as in Section 6.8.1, we also consider

Table 6.3 Mean (\pm st. dev.) of CVL and LCV selected bandwidths for the simulated two point process cases. Expected number of points is 400 for each listed process.

Interaction type	Intensity function	σ_{CVL}	σ_{LCV}
Segregated	constant	0.063 (0.008)	0.038 (0.006)
	hole	0.062 (0.009)	0.039 (0.008)
	waves	0.064 (0.010)	0.040 (0.008)
Poisson	constant	0.048 (0.006)	0.60 (0.19)
	hole	0.048 (0.006)	0.28 (0.22)
	waves	0.051 (0.006)	0.19 (0.20)
Co-clustered	constant	0.062 (0.008)	0.040 (0.008)
	hole	0.060 (0.009)	0.040 (0.007)
	waves	0.064 (0.011)	0.040 (0.009)




the case of a parametric model where the intensity $\alpha > 0$ of the underlying stationary point process (that is, before thinning) is unknown but the retention probability function p that was used to thin the point process is known. Then a simple parametric estimator for ρ is given by

$$\hat{\rho}_p(x) = Np(x) / \int_W p(x) dx$$

where N is the number of points in $X \cap W$. We apply this intensity estimator to $\hat{K}_{12,local}$ and $\hat{K}_{12,global}$ for 1000 realizations of each interaction type, with the “waves” intensity function and expected

Table 6.4 Root integrated mean squared errors $\times 10^2$ of local and global K_{12} -function estimators with CVL and LCV bandwidths.

Interaction	Intensity	$\hat{K}_{12,local}$		$\hat{K}_{12,global}$	
		CVL	LCV	CVL	LCV
Segregated	flat	0.65	390.125	0.161	0.181
	hole	0.69	4.574	0.171	0.185
	waves	0.64	270.633	0.208	0.201
Independent	flat	1.03	0.066	0.024	0.049
	hole	1.09	0.112	0.026	0.109
	waves	0.95	0.191	0.037	0.104
Co-clustered	flat	0.92	18.783	0.234	0.262
	hole	0.97	3.510	0.239	0.265
	waves	0.92	5.238	0.195	0.244

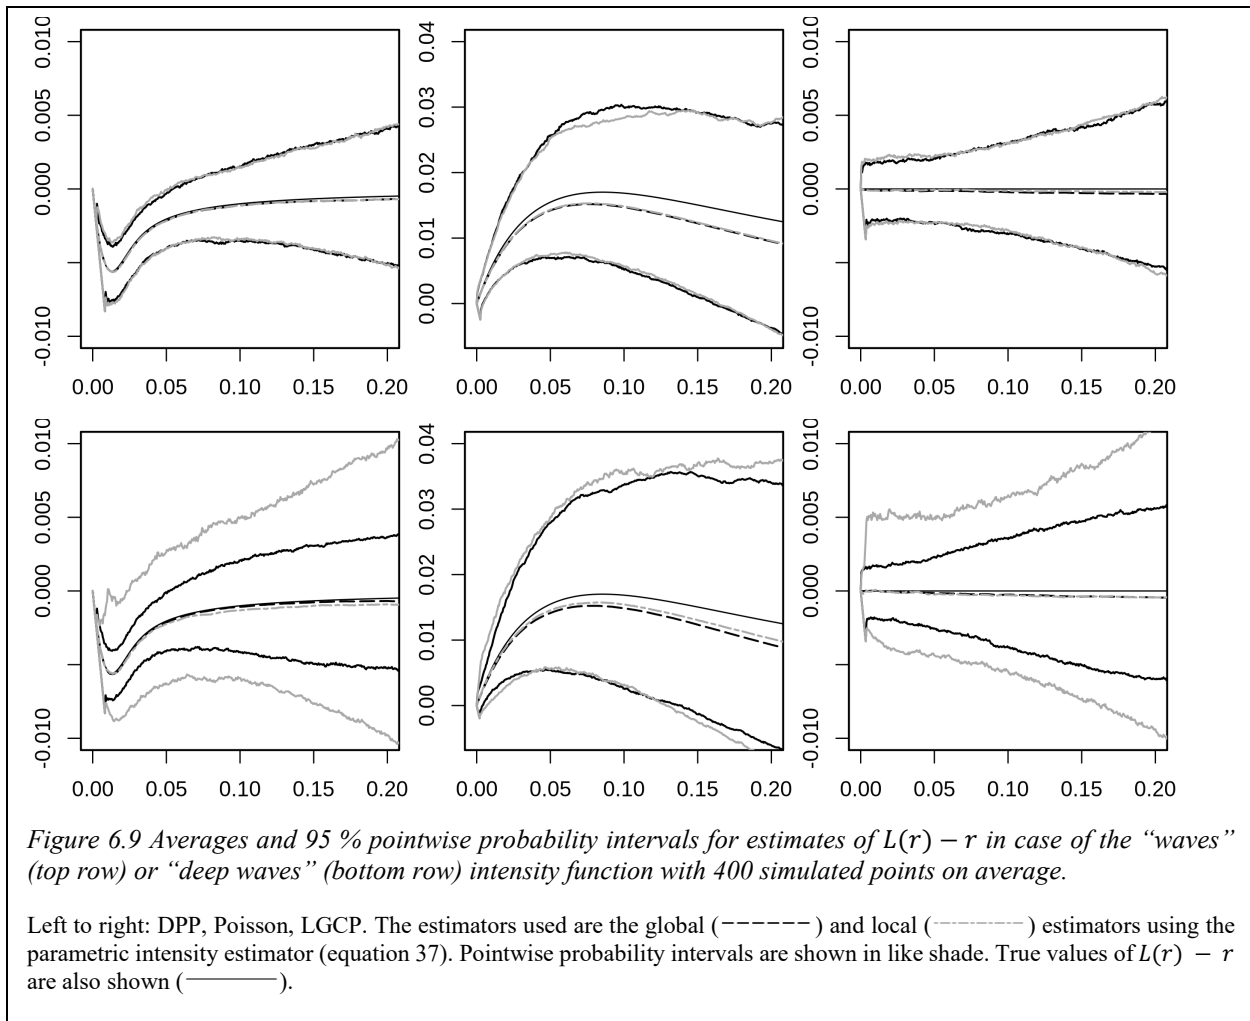


number of points equal to 400. In addition, we generate 1000 simulations for each interaction type with a new thinning profile, “deep waves”, given by

$$p_{\text{deep}}(x, y) = 1 - .9 \cos^2(5x), \quad (x, y) \in [0, 1]^2.$$

The deep waves profile is similar to the waves profile, but with much more extreme intensity variations.

Pointwise probability intervals for estimates of $L(r) - r$ are shown in Figure 6.9, and root integrated mean square error for estimates of K are given in Table 6.5. We observe that in all cases the error of the global estimator is comparable to or better than the corresponding local estimator. For the “waves” intensity function, the difference is small. Both estimators have larger error when applied to the patterns with the “deep waves” intensity function. However, the performance of the local estimator degrades much more strongly, reflecting the fact that regions of low intensity are weighted more heavily in \hat{K}_{local} than in \hat{K}_{global} , as discussed in Section 6.4.3. The LGCP yielded the largest errors with the parametric intensity estimates, similar to our observations with the kernel-based intensity estimates. We also note that for the DPP and the Poisson process, using the parametric estimates for the “waves” intensity function results in higher integrated mean square error than for the kernel-based estimates (Table 6.2). We believe this is because the kernel-based estimates of ρ are adapted to the random local fluctuations of the point processes, similar to how homogeneous K -function estimates have lower variance when using estimated intensity than true intensity. However, for the LGCP, best results are obtained with the parametric estimates, which



presumably are less prone to confounding of random clustering with variations in the intensity function.

Table 6.5 Root integrated mean squared errors $\times 10^2$ of local and global K -function estimators with parametric intensity estimator, applied to point processes with intensity function “waves” or “deep waves”.

Interaction type	Intensity function	\hat{K}_{local}	\hat{K}_{global}
DPP	waves	0.111	0.102
	deep waves	0.227	0.103
Poisson	waves	0.132	0.122
	deep waves	0.239	0.133
LGCP	waves	0.416	0.417
	deep waves	0.601	0.516

6.9 Extensions

The same sort of analysis as in Sections 6.4 and 6.5 could be applied to point processes defined on a non-empty manifold on which a group acts transitively (a so-called homogeneous space), where the space is equipped with a reference measure which is invariant under the group action. In this paper, the space was \mathbb{R}^d , the group action was given by translations, and the reference measure was Lebesgue measure. For example, instead we could consider the space to be a d -dimensional sphere, with the group action given by rotations and where the reference measure is the corresponding d -dimensional surface measure. Then the global and local estimators considered in this paper are simply modified to the case of the sphere by replacing Lebesgue with surface measure and using appropriate edge correction factors as defined in (315). Similarly, our global estimators could also be extended to the case of spatio-temporal point processes, as in (89, 90).

6.10 Conclusion

According to our simulation studies, our new global estimators outperform the existing local estimators in terms of bias and mean integrated squared error when kernel or parametric estimators are used for the intensity function. The kernel intensity function estimators depend strongly on the choice of bandwidth and we considered two different data-driven approaches, CVL and LCV, to bandwidth selection. In our simulation studies the two approaches gave similar selected bandwidths in the LGCP case but very different results in case of Poisson and DPP. This has a considerable impact on the K - and pair correlation function estimators but the global estimators appear to be much less sensitive to the choice of bandwidth selection method than the local estimators. The simulation studies with parametric estimates of the intensity function, along with the theory of Section 6.4.3, indicate that the global estimators are also much less sensitive to regions of especially low intensity. The improved statistical efficiency comes at a considerable extra computational cost. Therefore, we especially recommend the global estimators for situations where intensity variations are large and where computational speed is not a primary concern.

Chapter 7 Spatial Pair Correlation Analysis of Localization Microscopy Data under Spatially Nonuniform Sampling

Non-uniform sampling can arise in localization microscopy in the process of labeling the molecules of interest with fluorophores, or in the detection process, for instance due to non-uniform illumination over a large field of view. This chapter extends the methods proposed in Chapter 6 for estimating pair correlation functions in conditions where sampling of points is not uniform and applies them to the particularities of localization microscopy. In addition, I propose and validate a novel edge-correction method for the kernel-based estimator of γ , concluding that it is superior to the previously studied estimators. Finally, I present an example application of this method to cross-correlations in a B cell SMLM dataset. All of the work included in this chapter is my own, except the dataset that is used for demonstration of the method, which was collected by Sarah Veatch. This chapter has not been published at this time.

7.1 Abstract

Pair correlation functions, especially pair cross correlation functions have proved to be valuable tools in quantifying subtle statistical interactions in single molecule localization microscopy. Traditional methods for estimating pair correlation functions rely on the assumption that the localization density is uniform on the region of interest. However, even if the molecules of interest are distributed with uniform density, experimental challenges can cause local variations in the probability of observing the molecule of interest. This can either be due to non-uniform labeling, where some molecules are not observed at all, or due to non-uniform detection, where molecules in some regions are localized more times than others, on average. In turn, correlation functions computed using the standard methods include artifactual correlations. In this paper, we present methods for correcting these artifacts, and show their applicability to single molecule localization microscopy, using both simulated and experimental data. We also discuss a diagnostic that can be useful for determining whether the correction is necessary, and what parameters should be used in applying the correction.

7.2 Introduction

In single molecule localization microscopies (SMLM) such as PALM (37, 39), (d)STORM (38, 40) and PAINT (47), single probe molecules are made to fluoresce stochastically, so that the emission patterns of the emitters that are “on” are separated in space, and can be fitted to determine the location of the probe to high precision, often around 20 nm. The resulting dataset is a list of coordinates, accumulated from many thousands to millions of localizations. The localizations can then be assembled into an image that resembles the structure that was labeled. Alternatively, statistical analysis can be applied to the pattern of localizations to extract subtle structure that is not apparent when examining the images by eye. The pair (cross) correlation function is a useful way of characterizing statistical attraction or repulsion between points, which has been broadly applied, including to gain insight into localization microscopy data (53, 83, 87, 310, 332).

Correlation functions quantify interactions by comparing the observed distribution of separations of pairs of points to a benchmark distribution in which points are independent, i.e. have no interactions. For a given random process that produces a set of points X , the pair correlation function $g(\mathbf{x}, \mathbf{x}')$ at locations \mathbf{x} and \mathbf{x}' can be thought of as a ratio of probabilities:

$$g(\mathbf{x}, \mathbf{x}') = \frac{P(\mathbf{x} \in X \text{ and } \mathbf{x}' \in X)}{P(\mathbf{x} \in X)P(\mathbf{x}' \in X)}.$$

Note that when the points of X are independent, $g(\mathbf{x}, \mathbf{x}') = 1$. Values greater than one indicate that finding a point at \mathbf{x}' is more likely when there is also a point at \mathbf{x} , for example when the points tend to cluster. By contrast, values less than one indicate that a point at \mathbf{x} makes a point at \mathbf{x}' less likely. Similarly, when two types of points are observed, as sets X_1 and X_2 , an analogous cross-correlation function may be defined:

$$c(\mathbf{x}, \mathbf{x}') = \frac{P(\mathbf{x} \in X_1 \text{ and } \mathbf{x}' \in X_2)}{P(\mathbf{x} \in X_1)P(\mathbf{x}' \in X_2)}$$

with a similar interpretation. In SMLM experiments, X_1 and X_2 may represent localizations of two different proteins, labeled with fluorophores of distinct colors. The cross correlation function is often more useful than the autocorrelation function, since the autocorrelation function is typically dominated by multiple localizations of the same fluorophore at short distances, which often obscures more interesting interactions (87).

Although $g(\mathbf{x}, \mathbf{x}')$ or $c(\mathbf{x}, \mathbf{x}')$ may depend on both locations \mathbf{x} and \mathbf{x}' in principle, the pair correlation function is most useful when it is *translation invariant*, in which case it only depends on the displacement $\Delta\mathbf{x} = (\mathbf{x}' - \mathbf{x})$. $g(\Delta\mathbf{x}) = g(\mathbf{x}, \mathbf{x} + \Delta\mathbf{x})$.

$$g(\Delta\mathbf{x}) = g(\mathbf{x}, \mathbf{x} + \Delta\mathbf{x}).$$

Translation invariance holds in general for fluids. Translation invariance is typically required in order to estimate the pair correlation function, since we assemble the estimate by counting pairs near $(\mathbf{x}, \mathbf{x} + \Delta\mathbf{x})$ for all observed points. In many cases the pair correlation function is also *isotropic*, meaning that it depends only on the distance $r = \|\mathbf{x}' - \mathbf{x}\|$. Again, this is frequently true in fluids. However, note that the considerations presented here can be applied to data where the assumption of isotropy does not hold by omitting azimuthal averages.

A standard further simplifying assumption in correlation analysis is that the density of points is uniform everywhere in the region of interest. This is often a reasonable assumption about the distribution of the molecules of interest, for example in a fluid plasma membrane, where we expect membrane proteins to be equally likely to occupy any location on the membrane. However, the experimental realities of localization microscopy experiments can lead to non-uniform detection of localizations. In some cases, the fluorescent labeling of the molecule of interest is non-uniform in space, for example in situations with poor antibody access. In others, detection of the fluorescent label is non-uniform, e.g. when photoswitching or excitation illumination varies across the field of view, or when detection is impeded by non-uniform background fluorescence. Since we observe localizations, rather than the molecules of interest themselves, our correlation function estimate must account for these non-uniform detection scenarios. Thus, we will focus here on estimating the pair cross-correlation function $c(r)$ when the density of localizations is not uniform. Importantly, even in these non-uniform sampling conditions, the auto- and cross-correlation functions of the localizations are the same as those of the underlying proteins, as long as localizations are sampled independently from the locations of the proteins (88) (up to a convolution with the localization precision (87)).

To our knowledge, two approaches have been developed to correct correlation function estimators for point datasets with non-uniform density. The first weights each pair of points by the inverse of the local density at those points, which we call a local density correction (88). The second counts each pair with equal weight, and instead normalizes the entire count with a global density correction factor γ that represents the total number of pairs that would be expected

if the points were independently distributed according to the density profile (312). This second approach was originally proposed in our lab for live-cell space-time cross correlations, which suffer from bleaching probabilities that depend on the distance from the edge of the cell, due to diffusion (83). Recent work formalizing the global correction has demonstrated that it performs better than the local correction in cases with a few hundred points, and argued that it is generally expected to yield lower bias and variance (312). In this paper, we describe how these density corrected correlation functions may be computed for SMLM data with large numbers of points, demonstrate their usefulness in simulated and real SMLM data, emphasizing the superiority of the global correction. While the method is straightforward, best results are sensitive to a choice of parameter for local density estimation. We provide a diagnostic to help choose good values for this parameter. Finally, we discuss some practicalities that should be considered when applying the global correction, especially in SMLM data.

7.3 Results and Discussion

7.3.1 Naive correlation function estimate yields artifactual correlations

In the following, we consider an SMLM dataset with two types of localizations, as one would obtain from two-color fluorescence localization microscopy. As above, X_i denotes the set of localizations of type i , and $\rho_i(\mathbf{x})$ the density of localizations of type i at location \mathbf{x} . X_i is called a spatial point process, and the localizations take values in \mathbb{R}^2 for purposes of this paper. The density $\rho_i(\mathbf{x})$ should be thought of as an expected number of points per unit area in the vicinity of \mathbf{x} , rather than a particular realized number of points per unit area that are found near \mathbf{x} .

In practice, we observe localizations on a finite region of interest or observation window $W \subset \mathbb{R}^2$, for example the ventral membrane of a cell. We write $|A|$ for the area of A , and define the translation of A by a displacement \mathbf{h} with the notation $A_{\mathbf{h}}$, so that $A_{\mathbf{h}} = \{\mathbf{x} + \mathbf{h} | \mathbf{x} \in A\}$.

With these definitions in place, we may write the standard, uniform density cross-correlation function estimator:

$$\hat{c}_{\text{unif}}(r) = \frac{1}{\rho_1 \rho_2 a(r)} \sum_{\mathbf{x} \in X_1, \mathbf{y} \in X_2} \frac{\mathbf{1}_{[|\|\mathbf{y}-\mathbf{x}\|-r| < \frac{b}{2}]} }{2\pi r b}, \quad (1)$$

where $b > 0$ is a finite bandwidth as described below, $\rho_i(\mathbf{x}) = \rho_i$ is assumed to be uniform and $a(r)$ is an azimuthally averaged edge correction factor

$$a(r) = \langle |W \cap W_{-h}| \rangle_{\|h\|=r}.$$

As a whole, equation 1 can be thought of as a ratio where the numerator is the sum, which counts an observed density of pairs at the distance r , and the denominator is the normalization factor $\rho_1\rho_2a(r)$, which represents the expected number of pairs in an uncorrelated point process. The sum effectively counts pairs of points with separation near r . The finite bandwidth b sets the range of distances that are considered. Adjusting b leads to a tradeoff between variance and spatial resolution: large values produce low-variance estimates because they average over more data, but blur the underlying correlations. The factor of $2\pi rb$ is a geometrical factor to account for the size of the set of separations that are considered.

In the normalization factor $\rho_1\rho_2a(r)$, the edge correction $a(r)$ accounts for the finite observation window: it effectively counts how much space there is in W that has a pair in W at separation r . The factor of $\rho_1\rho_2$ accounts for the effect of density: the simple count of pairs at separation r is proportional to the density of each type of point.

When the assumption of uniform density does not hold, one may replace the densities ρ_i with the average density $N_i/|W|$, where N_i is the number of points of X_i . However, this practice generally induces bias in the correlation function estimates. To make this concrete, consider the simulated dataset of Figure 7.1, representing an idealized two-color SMLM experiment. Points of two types are distributed uniformly and independently across a circular observation window (Figure 7.1a). Since they are distributed independently, $c(r) = 1$. The points are poorly labeled, with points near the edge of the observation window more likely to be labeled than those in the interior (Figure 7.1b). This scenario could represent an experiment using antibodies to label a membrane protein of interest with a fluorophore, in a dense sample where antibody access is poor. Figure 7.1c shows the two parts of equation 1: the sum that counts pairs of points with separation near r , and the normalization factor $\rho_1\rho_2a(r)$, where densities ρ_i correspond to average density over the entire region of interest. The cross correlation function estimate $\hat{c}_{\text{unif}}(r)$ is simply the ratio of these lines, and is shown in Figure 7.1d. Clearly, $\hat{c}_{\text{unif}}(r)$ overestimates $c(r) = 1$ in this case.

7.3.2 Local and global density corrections for cross correlation functions

Baddeley et al (88) introduced a density-corrected cross-correlation function that we call local because it weights each counted pair in the numerator with the local density:

$$\hat{c}_{\text{local}}(r) = \frac{1}{a(r)} \sum_{\mathbf{x} \in X_1, \mathbf{y} \in X_2} \frac{\mathbf{1}[\|\mathbf{y}-\mathbf{x}\|-r| < b/2]}{2\pi r b \rho_1(\mathbf{x}) \rho_2(\mathbf{y})}.$$

As a result, points from high density regions are effectively down-weighted, to compensate for the excess pairs in that region that are simply due to the high density. The sum can be directly compared to the sum in \hat{c}_{unif} by multiplying by the average densities:

$$\langle \rho_1 \rangle \langle \rho_2 \rangle \sum_{\mathbf{x} \in X_1, \mathbf{y} \in X_2} \frac{\mathbf{1}[\|\mathbf{y}-\mathbf{x}\|-r| < b/2]}{2\pi r b \rho_1(\mathbf{x}) \rho_2(\mathbf{y})},$$

which is the quantity shown as “weighted pairs” in Figure 7.2a. The resulting estimate $\hat{c}_{\text{local}}(r)$ is the ratio of these weighted counts (dashed black line) and the edge-corrected density normalization from before (red line), and is shown in yellow in Figure 7.2b. $\hat{c}_{\text{local}}(r)$ is an unbiased estimator for $c(r)$, as demonstrated here by the fact that it closely straddles $c(r) = 1$.

By contrast, the global estimator accounts for density variations by moving the density functions into the edge correction convolution:

$$\hat{c}_{\text{global}}(r) = \frac{1}{\gamma(r)} \sum_{\mathbf{x} \in X_1, \mathbf{y} \in X_2} \frac{\mathbf{1}[\|\mathbf{y}-\mathbf{x}\|-r| < b/2]}{2\pi r b},$$

where $\gamma(r)$ is an angularly averaged global normalization factor

$$\gamma(r) = \langle \gamma(\mathbf{h}) \rangle_{\|\mathbf{h}\|=r}$$

and

$$\gamma(\mathbf{h}) = \int_{W \cap W_{-\mathbf{h}}} \rho_1(\mathbf{x}) \rho_2(\mathbf{x} + \mathbf{h}) \, d\mathbf{x}.$$

In this case, the numerator is the same as that of \hat{c}_{unif} : an unweighted count of the observed number of pairs separated by a distance r , as shown by the black line in Figure 7.2a. However, the denominator $\gamma(r)$ gives the average number of such pairs that would be present if the points were distributed independently according to the densities $\rho_i(\mathbf{x})$. The quantity $\gamma(r)$ is also shown in Figure 7.2a. The resulting $\hat{c}_{\text{global}}(r)$ estimate is the ratio of these lines, and is shown in Figure 7.2b. $\hat{c}_{\text{global}}(r)$ is also an unbiased estimator of $c(r)$.

Note that when the densities are constant, $\gamma(r) = \rho_1 \rho_2 a(r)$, and all three definitions are exactly equivalent.

7.3.3 Variance of global estimator is lower than that of local estimator

As we note above, both $\hat{c}_{\text{local}}(r)$ and $\hat{c}_{\text{global}}(r)$ are unbiased: they accurately reproduce the true value $c(r)$ on average. However, the variances of the estimates are different in general. In practice, we find that $\hat{c}_{\text{global}}(r)$ has a smaller variance than $\hat{c}_{\text{local}}(r)$, and is therefore preferred.

Figure 7.3 demonstrates the improved variance of $\hat{c}_{\text{global}}(r)$ compared to $\hat{c}_{\text{local}}(r)$ for simulated SMLM data that is somewhat more realistic than that of Figure 7.1. Figure 7.3a shows an example dataset. As in Figure 7.1, proteins are distributed uniformly on the cell surface with density $20 \mu\text{m}^{-2}$, and proteins near the edge are more likely to be labeled by a fluorescent antibody. However, in this case each fluorescent label produces a Poisson-distributed number of localizations, with Gaussian localization error. The localization precision is 25 nm, and mean number of localizations per fluorophore is 10. Figure 7.3b shows mean and 95% pointwise probability intervals for the three cross-correlation estimates, calculated from simulated point processes. The probability intervals for $\hat{c}_{\text{global}}(r)$ are substantially tighter than those for $\hat{c}_{\text{local}}(r)$, as expected.

We can develop an intuitive explanation for this trend by recalling that the inverse density weights in the local estimator weight rare events (from low-density regions) strongly, and common events (from high-density regions) weakly. If the range of densities is extreme, the sum in (1.3) is dominated by the few pairs that are observed in low-density regions, effectively reducing the sample size. By contrast, the pairs in the sum in (1.4) are all weighted equally, making the effective sample size as large as possible.

We note that the same reasoning does not apply when the variation of localization density is entirely induced by non-uniform overcounting. In that case the weighting in the local estimator effectively weights each molecule equally, by down-weighting localizations from molecules with more overcounting. By contrast, the global estimator effectively weights molecules with more localizations more strongly. In simulations with uniform labeling but non-uniform overcounting, we find that the error of the global and local estimators is roughly equal. For an example, see Figure 7.4.

7.3.4 Estimating density corrections when density is unknown

Both of the density corrections described above depend on accurate knowledge of the local density as a function of space. Except in a simulation setting as in Figure 7.1 and Figure 7.2, the local density is usually unknown, and must be estimated from the data. In this paper, we will use

a kernel density estimate, which is a common non-parametric density estimate. Kernel density estimates are a useful baseline, since they do not depend on any assumptions about the density variations. Often, they are the only option available. However, if more is known about the source of density variations, such that it can be accurately modeled with a parametric density function, a parametric estimate of the density will generally lead to somewhat more precise and accurate cross-correlation function estimates, using either the local or global density correction.

A kernel density estimator $\hat{\rho}(\mathbf{x})$ is based on a kernel $\kappa_\sigma(\mathbf{h})$, which determines the contribution of a point at $\mathbf{x} + \mathbf{h}$ to the density estimated at \mathbf{x} . The parameter σ is a scale parameter for κ_σ , so

$$\kappa_\sigma(\mathbf{h}) = (1/\sigma^2)\kappa_1(\mathbf{h}/\sigma).$$

In other words, σ sets the length scale over which observed points contribute to the density estimate.

Intuitively, the density estimate is obtained by blurring the data itself, with the extent of blurring determined by κ_σ . For \mathbf{x} near the edge of the observation window, κ_σ places significant weight on the area outside W where points are not observed. Therefore an “edge correction” is required. Overall, $\hat{\rho}(\mathbf{x})$ can be written

$$\hat{\rho}(\mathbf{x}) = \sum_{\mathbf{x}' \in X} \frac{\kappa_\sigma(\mathbf{x}' - \mathbf{x})}{w(\mathbf{x})}$$

with the edge correction

$$w(\mathbf{x}) = \int_W \kappa_\sigma(\mathbf{x} - \mathbf{x}') \, d\mathbf{x}$$

due to Diggle (322). To make a kernel estimator concrete, we must choose the shape of the kernel κ_σ and the bandwidth σ . It has been noted that the shape of the kernel is less important than the choice of bandwidth (86). In the following, we use a Gaussian kernel for convenience, since it can be separated into a product of functions that each depend on only one coordinate:

$$\kappa_\sigma(\mathbf{h}) = \frac{1}{2\pi\sigma^2} e^{-\mathbf{h}^2/(2\sigma^2)}.$$

Shaw et al (312) note that the density corrections for auto-correlations depend on products of these estimates $\hat{\rho}(\mathbf{x})\hat{\rho}(\mathbf{y})$, and that these products are prone to bias due to correlations between nearby density estimates. They propose instead a direct estimate of the product

$$\overline{\rho(\mathbf{x})\rho(\mathbf{y})} = \sum_{\mathbf{x}', \mathbf{y}' \in X}^{\neq} \frac{\kappa_\sigma(\mathbf{x} - \mathbf{x}')\kappa_\sigma(\mathbf{y} - \mathbf{y}')}{w(\mathbf{x}')w(\mathbf{y}')}$$

i.e. excluding the “diagonal” term $\mathbf{x}' = \mathbf{y}'$. From this, an estimator for γ can be obtained by applying the definition directly:

$$\tilde{\gamma}_\sigma(\mathbf{h}) = \int_{W \cap W_{-\mathbf{h}}} \overline{\rho_1(\mathbf{x})\rho_2(\mathbf{x} + \mathbf{h})} \, d\mathbf{x}$$

In addition, we define one more γ estimator inspired by $\bar{\gamma}(\mathbf{h})$:

$$\tilde{\gamma}_\sigma(\mathbf{h}) = |W \cap W_{-\mathbf{h}}| \sum_{\mathbf{x} \in X_1, \mathbf{y} \in X_2} \frac{\kappa_{\sigma\sqrt{2}}(\mathbf{y}-\mathbf{x}-\mathbf{h})}{|W \cap W_{-(\mathbf{y}-\mathbf{x})}|}$$

Intuitively, this estimator improves on $\bar{\gamma}(\mathbf{h})$ by weighting each pair \mathbf{x}, \mathbf{y} separated by \mathbf{u} equally to the estimate of $\gamma(\mathbf{h})$, rather than weighting points near the edge more strongly. Here, the edge correction factors are important when σ is large, in which case $\mathbf{y} - \mathbf{x}$ may be significantly different from \mathbf{h} .

Both kernel estimators are constructed to be unbiased estimators for γ under the rather strict condition that $\rho(\mathbf{x}) = \rho$ is constant and no pairwise correlations are present. In the general case, where density variation and correlations are present, bias may arise due to several roughly separable considerations, which we attempt to disentangle here.

7.3.5 Pairwise edge-correction for the normalization factor gamma reduces systematic errors

The choice of edge-correction is most important when the localization density at the edges is substantially different from the localization density in the interior of the region of interest. We find that the pairwise edge correction yields a more accurate estimate of γ than the pointwise one. This effect is demonstrated by the example of Figure 7.5. The density shown in Figure 7.5a is half again as dense at the edge as in the interior of the square observation window, with the transition between sparse and dense occurring over about 2 μm . The systematic errors of the two kernel-based γ estimates are shown in Figure 7.5b, for three choices of σ . Clearly, the pairwise edge corrected estimator $\tilde{\gamma}_\sigma$ yields smaller systematic errors than the pointwise one $\bar{\gamma}_\sigma$. Furthermore, the errors in $\tilde{\gamma}_\sigma$ decay over a shorter range of r . The pairwise edge correction is clearly superior with this density pattern. Moreover, we have not found any circumstance where the pointwise edge correction is superior to the pairwise one.

In light of this comparison, we recommend using the pairwise edge correction, especially for situations where the localization density at the edge is very different from the localization density in the interior. This is an important scenario, especially in live-cell experiments using TIRF

microscopy to illuminate the ventral membrane of a cell. Molecules that reach the center of the cell are more likely to have bleached, since they have been exposed to the TIRF illumination for a longer time than those at the edge, on average. Edge-heavy labeling could also occur due to poor antibody access in cell lines that adhere strongly to the cover slip.

7.3.6 Short range density variation induces bias in kernel based gamma estimates, especially for large sigma

Density variations on short length scales compared to the kernel bandwidth σ tend to be smoothed away, and can therefore lead to bias in $\tilde{\gamma}_\sigma$. For example, the simulated local density shown in Figure 7.6a is a sum of perpendicular sine waves, of wavelength 1 μm . Smoothing by length-scales of more than 1 μm are thus expected to substantially erode the peaks of the density profile. Figure 7.6b shows the systematic error of the estimate $\tilde{\gamma}_\sigma$ that corresponds to smoothed densities using $\sigma = .3, 1, \text{ and } 3 \mu\text{m}$. The error grows with σ , with negligible error for $\sigma = .3 \mu\text{m}$, which is much less than the wavelength of the density profile, but a noticeable error of about 2.5% when $\sigma = 3 \mu\text{m}$ is used. This error is due to smoothing away of details on the 1 μm length scale.

7.3.7 Density corrected estimators underestimate the strength of correlations, especially for small sigma

Finally, when the points of the sample are co-clustered or mutually segregated, the estimates of γ are biased by the resulting excess or paucity of point pairs at short distances. In particular, in the $\sigma \rightarrow 0$ limit, both γ estimators are biased by a factor of the correlation function (up to edge corrections):

$$E[\tilde{\gamma}(r)] = \gamma(r)g(r), \quad (0.13)$$

so that the globally-corrected correlation estimate that uses $\hat{\gamma}$ will always produce an estimate of $g(r) = 1$. For finite σ , the bias is reduced by blurring the peak of $g(r)$, yielding a bias of

$$E[\hat{\gamma}_\sigma(r)] = \gamma(r) \left\langle \int_{\mathbb{R}^2} g(|\mathbf{x}|) \kappa_\sigma(\mathbf{y} - \mathbf{x}) \, d\mathbf{x} \right\rangle_{|\mathbf{y}|=r} \quad (0.14)$$

In some cases $g(0)$ can be quite large. This is especially true for auto-correlations of fixed-cell SMLM data, where molecules tend to produce several or many localizations; in the absence of correlations between the position of the molecules, $g(r)$ exhibits a strong peak at $r = 0$, with amplitude inversely proportional to the density of molecules. In such a case, it is important to choose $\sigma \gg r_{cor}$. Figure 7.7 shows how large correlations induce bias in estimates of gamma

and the resulting auto-correlation estimates for an overcounted, uncorrelated simulation of SMLM data. Figure 7.7a shows γ , as well as estimates of gamma using a range of σ . Corresponding autocorrelation functions $g(r)$ are shown in Figure 7.7b. With $\sigma = 100$ nm, the γ estimate overestimates γ by more than a factor of two at short r , so that $g(r)$ is underestimated by more than a factor of two. As σ is increased, this bias quickly decreases. Figure 7.7c shows how the overestimation of γ decreases as a function of σ , for three different correlation strengths.

This combination of observations leads to the conclusion that a good choice of σ is constrained both on the low end, by bias due to correlations, and on the high end, by bias due to smoothing of density variations. The length-scales where these features become important is experiment-dependent, and there may be experiments where the length-scales overlap, so that there is no choice of σ that results in an accurate estimate of the (cross-)correlation function. If the length-scale of the correlations of interest is approximately known, for example from the size of the structure that is being labeled, 10-20 times this value can be used as a starting point for σ .

One useful fact about the correlation-induced errors is that these errors are bounded between 1 and the (true) correlation function. As a result, choosing a too-small σ will lead to an underestimate of the strength of the correlations, but will not cause clustering ($c(r) > 1$) to be misidentified as segregation ($c(r) < 1$), or vice versa. Thus, when correlations and density variation are poorly understood, a useful strategy can be to estimate the correlation function using a range of sigma starting very small, and observing how the estimated short-range correlations develop with increasing σ .

Note that the same considerations can be applied in aid of choosing an observation window that excludes some portion of the observed localizations. In some cases, more accurate correlation functions may be obtained by excluding regions with density variation that is abrupt compared to the length scale of σ that will be used. Even though this reduces the effective number of observations, it also reduces a systematic error that is difficult to evaluate in a specific case.

7.3.8 Application to an SMLM dataset

Figure 7.7 demonstrates the application of density-corrected cross-correlation analysis to an SMLM dataset from a CH27 B cell with F-actin labeled by phalloidin-AlexaFluor647 shown in magenta and the ganglioside GM1 labeled by cholera toxin B subunit (CTxB) conjugated to AlexaFluor532, shown in green. We begin by estimating density profiles by the kernel estimator

$\hat{\rho}(\mathbf{x})$, with $\sigma = 1 \mu\text{m}$, shown in Figure 7.7b and c for CTxB and phalloidin, respectively. Both density profiles show substantial variation, with high-density regions up to twice as dense as low-density regions. These density variations are a strong indication that a density correction will be required for proper cross-correlation analysis. In addition to the density profiles, it is helpful to inspect the uncorrected radial and 2d correlation functions out to long distances, as shown in Figure 7.7d and e (radial and 2d, respectively). The 2d correlation function shows apparently anisotropic correlations at distances of several μm , with amplitudes of up to 25% more or less than would be expected from a random distribution. The radial correlation function is effectively a circular average over the 2d correlation function, so it also exhibits apparent correlations out to long distances, although the amplitudes are smaller, because the positive and negative correlations in different directions largely average out. Note that there is no kernel smoothing in Figure 7.7d and e, so the slowly-varying nature of these correlation functions at long distances indicates that the relevant features of the density profile are in fact slowly-varying. As a result, we expect that $\sigma \approx 1 \mu\text{m}$ will be small enough to adequately capture the spatial variation. Finally, Figure 7.7f shows $\hat{c}_{\text{global}}(r)$, calculating using a range of values for σ as indicated. When σ is in the range of 0.5-1.5 μm , $\hat{c}_{\text{global}}(r)$ shows a positive correlation of about 1.05, which represents 5% enrichment of CTxB around actin compared to the average density. Note that smaller σ leads to less pronounced positive correlations, as expected from the considerations of the previous section.

7.4 Conclusions

In this chapter I have adapted the density-correction of correlation functions, as introduced in chapter 3, to the context of SMLM. In addition, I have developed a superior kernel-based estimator of the normalization function γ that is less susceptible to bias, especially when the density is increasing or decreasing towards the edge of the dataset.

This method is most appropriate when there is a strong separation of scales between the length scales of biologically-relevant spatial correlations and those of experimentally-induced sampling variation. In that case one can choose σ to be simultaneously much smaller than the length scale of density variation and much larger than the length scale of correlations. Although that condition cannot always be guaranteed in practice, it is often feasible to distinguish between segregation and co-clustering by examining how \hat{c}_{global} varies as σ is increased. However, details

of correlation functions obtained by this method, such as their amplitude and spatial scale, should be interpreted cautiously and corroborated by alternative experimental strategies.

Some of the limitations of this method could be overcome by incorporating orthogonal estimates of the density profile. For example, if the nonuniform sampling is due to spatial variation in the intensity of the excitation or photoactivation light source, an experimentalist could directly measure that intensity profile and use it to model a detection probability for fluorophores at different locations in the sample.

7.5 Materials and Methods

7.5.1 Simulated SMLM datasets

For simulation figures, simulated localizations were generated using custom MATLAB (The MathWorks, inc., Natick, MA) software by first simulating an independent uniform distribution of true locations of two types of molecules on a region of interest, either circular or square as specified in the text. To simulate non-uniform labeling, the molecule locations were independently thinned by a spatially-varying thinning probability $p(\mathbf{x})$, so that a molecule at location \mathbf{x} is labeled with probability $p(\mathbf{x})$. Then each labeled molecule generates a Poisson-distributed number of localizations, each displaced from the labeled molecule by an independent Gaussian-distributed localization error with standard deviation as specified in the text.

To simulate uniform labeling with non-uniform optical detection (Figure 7.4), each molecule is labeled, but non-uniform independent thinning is applied to the localizations produced by the labeled molecules, so that the average number of localizations produced by a molecule at location \mathbf{x} is reduced by a factor of $p(\mathbf{x})$.

7.5.2 Computations for correlation functions

The correlation function estimators described above can be described in terms of various convolutions of various functions of space:

$$(f * g)(\mathbf{h}) = \int_{W \cap W_{-\mathbf{h}}} f(\mathbf{x})g(\mathbf{x} + \mathbf{h}) \, d\mathbf{x},$$

so that for example $\gamma = \rho_1 * \rho_2$.

Practical computational considerations lead us to discretize these functions on a lattice with spacing smaller than the correlations of interest, so that the function $f(\mathbf{x})$ is approximated by a

matrix $f(i, j)$ giving the average value of f over the (i, j) square pixel: $i < x/l < (i + 1), j < y/l < j + 1$, where l is the lattice spacing of the discretization and (x, y) are the x and y coordinates of \mathbf{x} . Unless otherwise specified, we use a lattice spacing of $l = 16$ nm. The sums that appear in the numerators of \hat{c}_{unif} and \hat{c}_{global} can also be treated as approximate convolutions:

$$\sum_{\mathbf{x} \in X_1, \mathbf{y} \in X_2} \frac{1[||\mathbf{y} - \mathbf{x}| - r| < (b/2)]}{2\pi r b} \approx \left\langle \sum_{\mathbf{x} \in X_1, \mathbf{y} \in X_2} \frac{1[|\Delta x - h_x| < l/2]1[\Delta y - h_y < l/2]}{l^2} \right\rangle_{||\mathbf{h}||=r}$$

$$= \langle I_1 * I_2 \rangle_{||\mathbf{h}||=r},$$

where u_x and u_y denote the x and y coordinates, respectively, of a vector \mathbf{u} , $\Delta x = y_x - x_x$, and $\Delta y = y_y - x_y$, and I_k is a discretization of the point set X_k , so that $I_k(i, j)$ is the number of points of X_i that fall in the (i, j) square pixel.

Then the convolution of two functions is computed using discrete Fourier transforms by

$$f * g = \mathcal{F}^{-1}(\mathcal{F}(g)^* \cdot \mathcal{F}(f)) / l^2$$

where \mathcal{F} is a 2D discrete fourier transform, all transforms are padded with zeros to twice their original size, and x^* is the complex conjugate of x .

Several of the estimators of γ also involve 2d kernel smoothing by a Gaussian kernel. These are computed by real-space discrete convolutions using MATLAB's built-in conv2 function to apply two 1d convolutions, each with a 1d Gaussian kernel that is extended to at least 3σ .

Finally, circular averages $\langle f(\mathbf{h}) \rangle_{||\mathbf{h}||=r}$ are performed on the discretized functions by calculating the distance from the center of the (0,0) pixel to each other pixel, finding the pixel with distances within $l / 2$ of r , and averaging the values of f at those pixels.

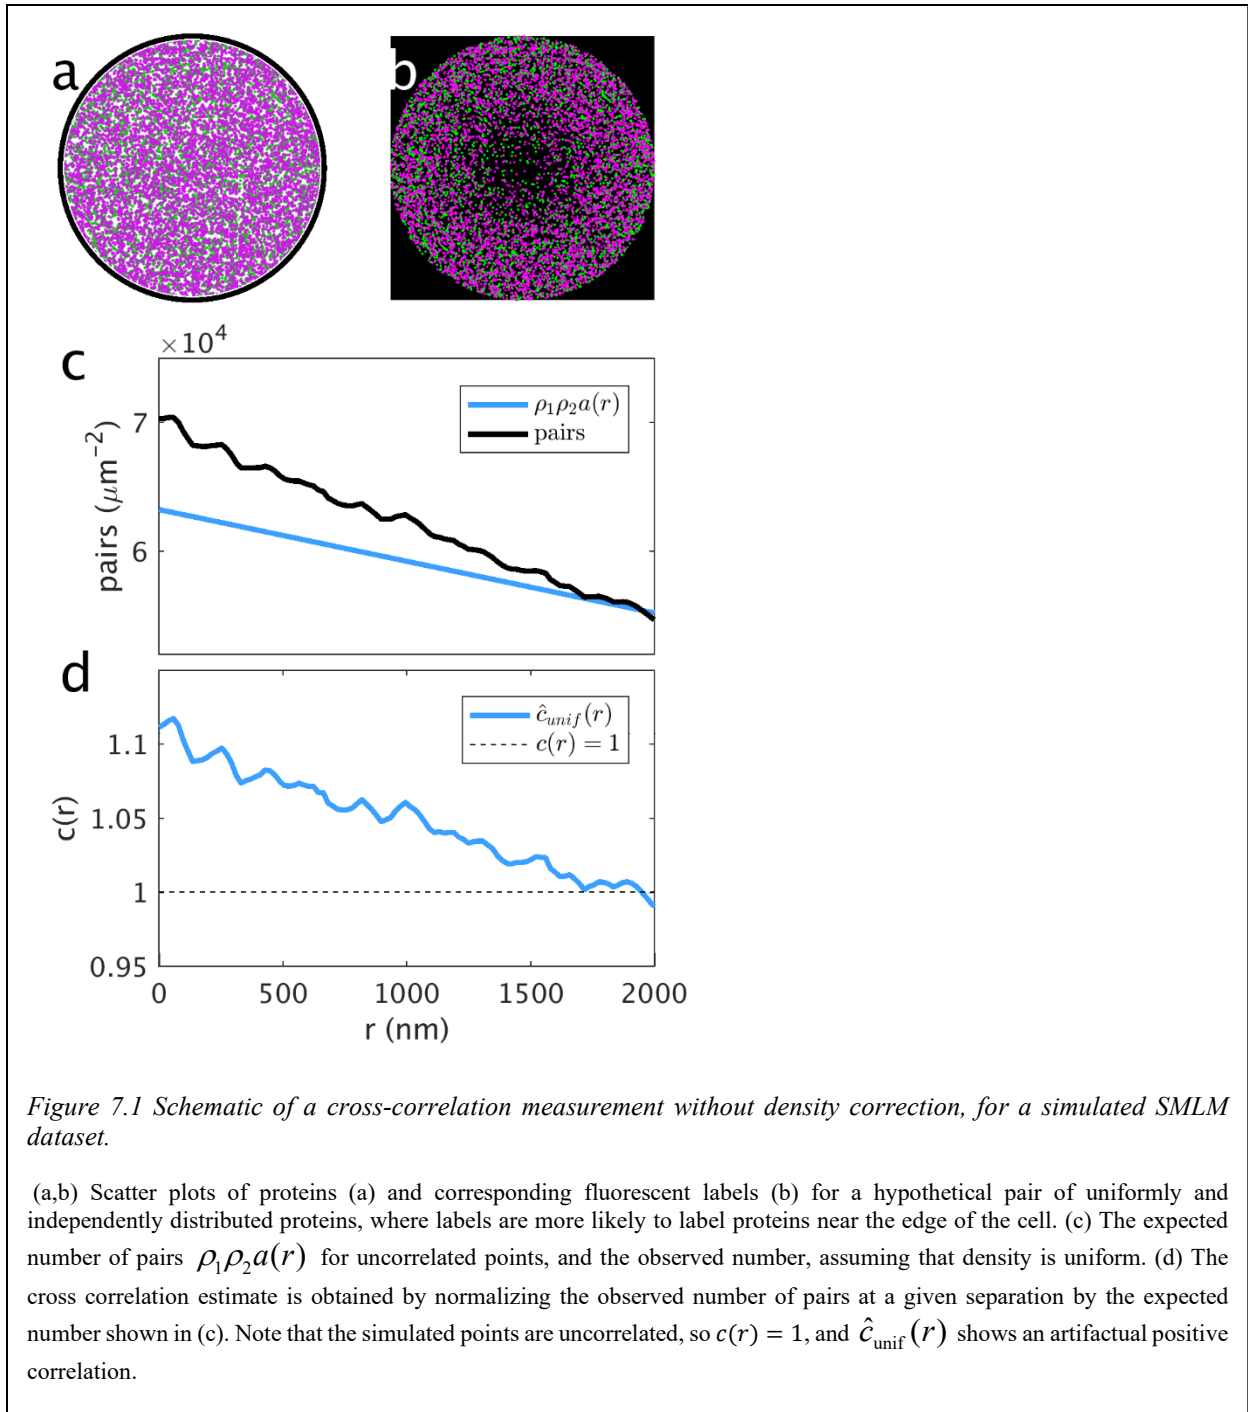
7.5.3 Sample preparation

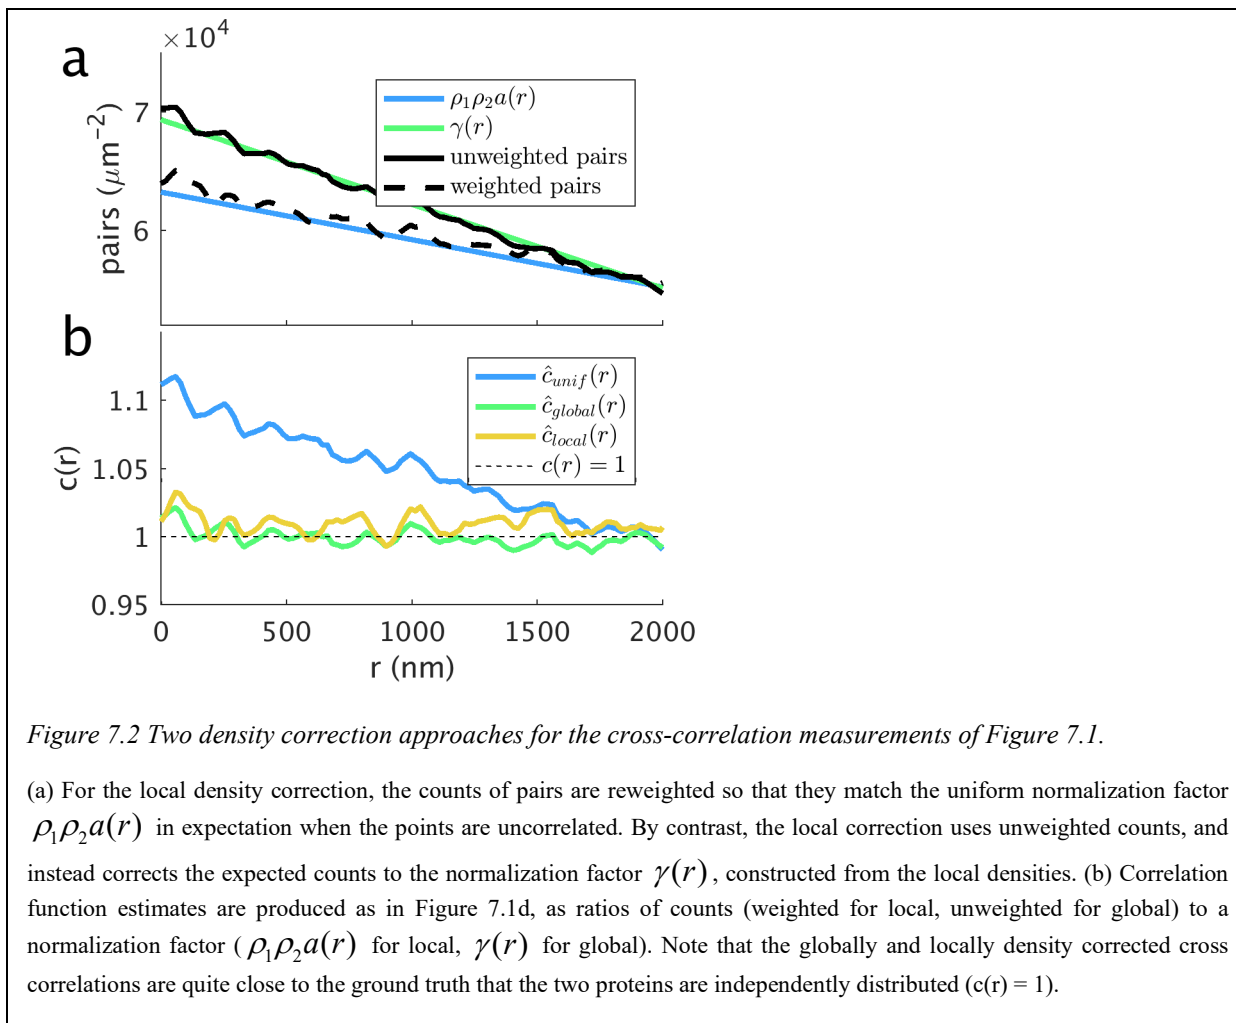
CH27 B-cells (mouse, Millipore Cat# SCC115, RRID:CVCL_7178), a lymphoma-derived cell line (292) were acquired from Neetu Gupta (Cleveland Clinic). CH27 Cells were maintained in culture as previously described (53). Cells were plated in MatTek glass-bottomed dishes and allowed to adhere overnight, followed by chemical fixation in 2% PFA and 0.2% glutaraldehyde (Electron Microscopy Sciences) for 10 min. F-Actin was stained by permeabilizing cells with 0.1% Triton-X-100 prior to incubation with 3.3 μM phalloidin-AlexaFluor647 (Invitrogen) for at least 15 min and imaged immediately after removing label.

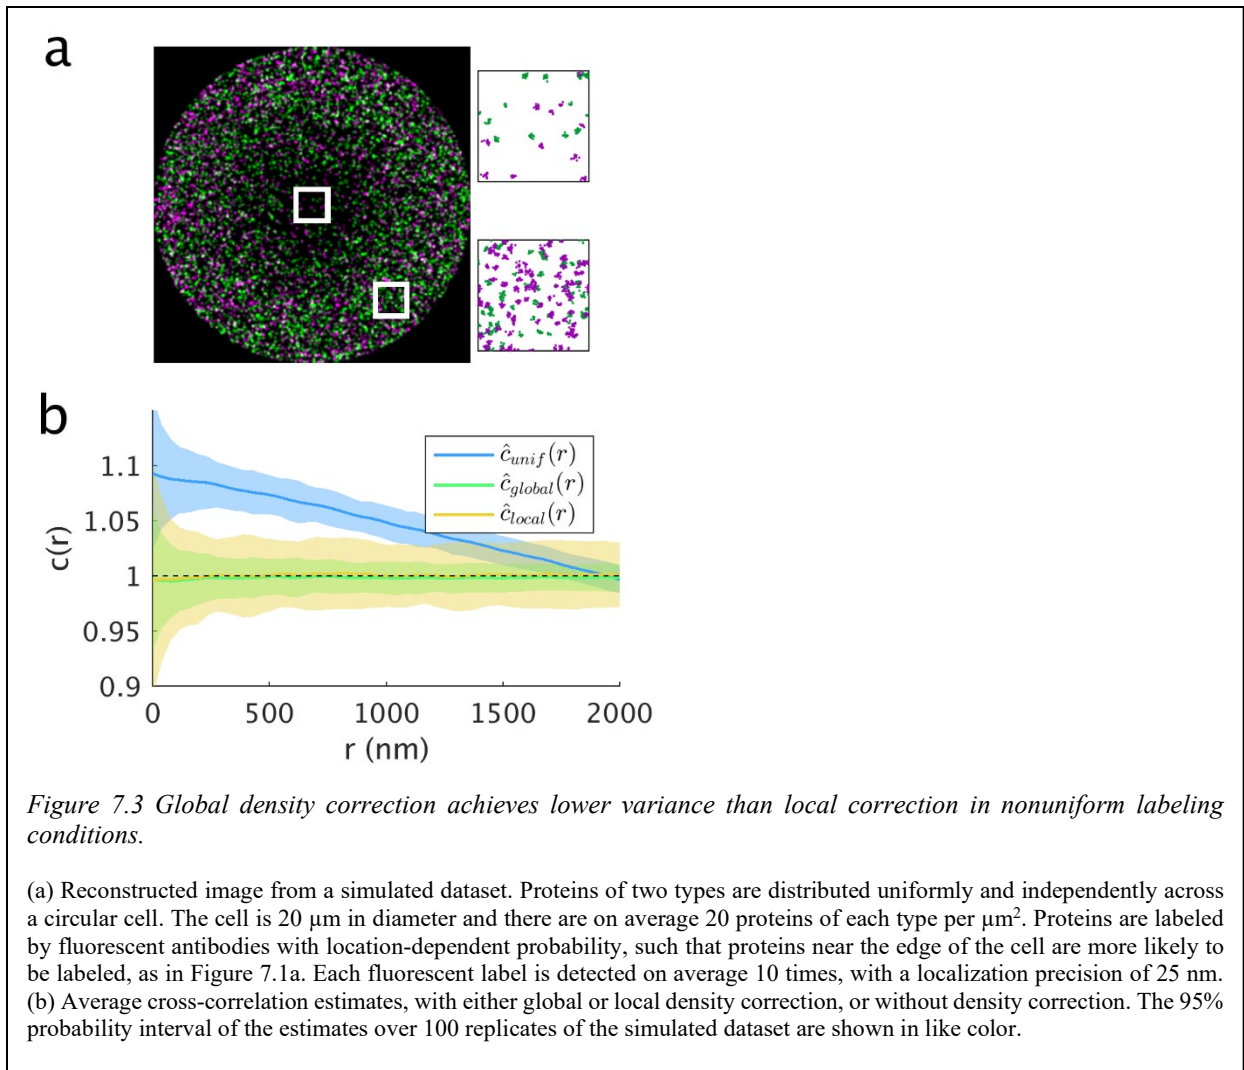
7.5.4 Imaging and localization

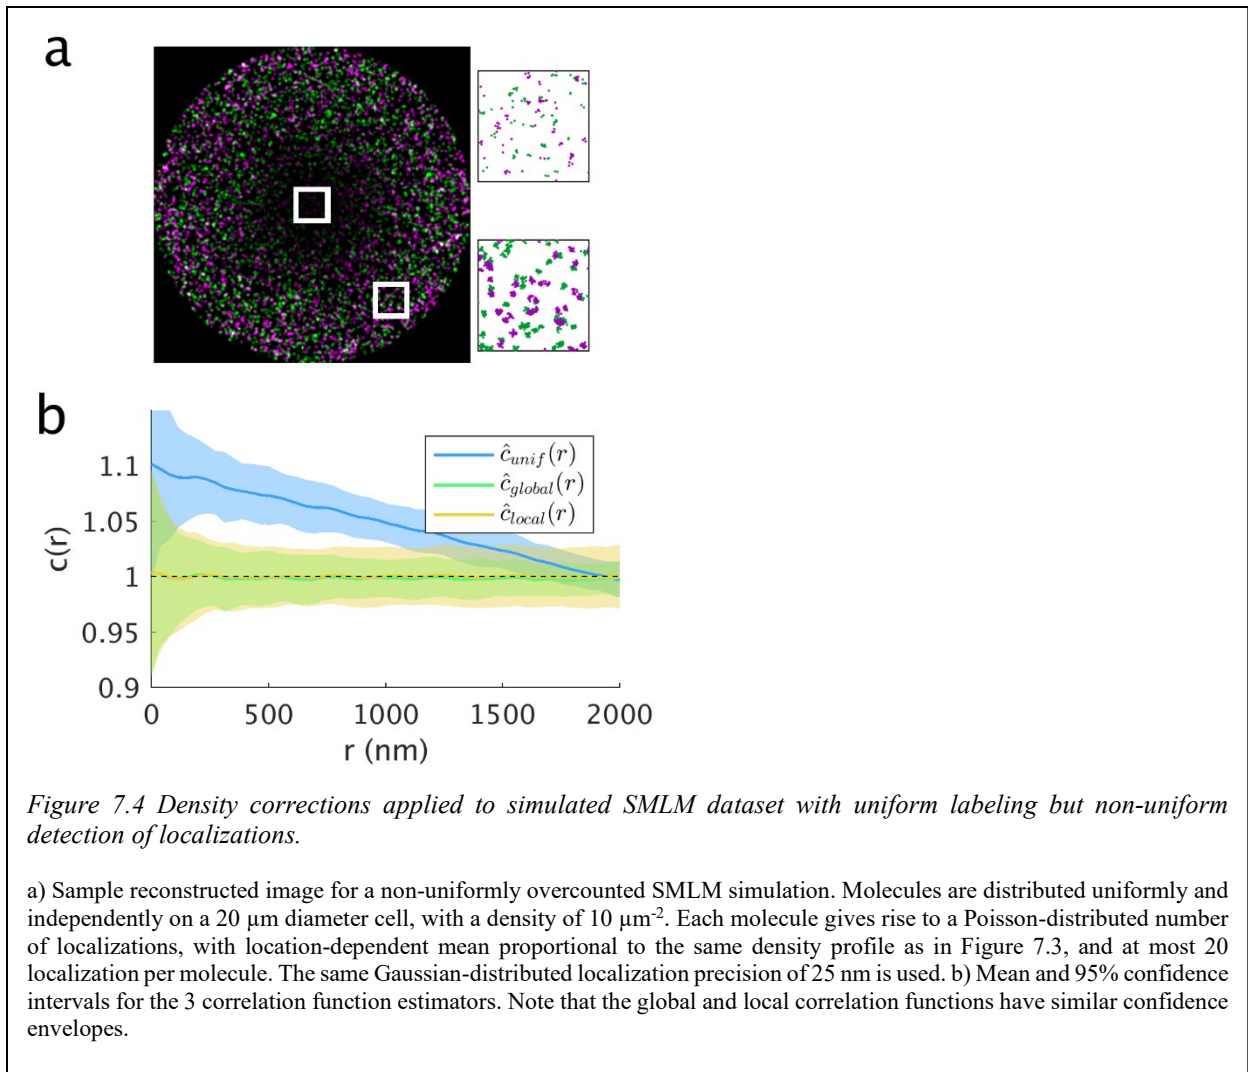
Imaging was performed using an Olympus IX83-XDC inverted microscope. TIRF laser angles were achieved using a 100X UAPO TIRF objective (NA = 1.50), and active Z-drift correction (ZDC) (Olympus America). AlexaFluor647 was excited using a 647 nm solid state laser (OBIS, 150 mW, Coherent) and AlexaFluor532 was excited using a 532nm solid state laser (OBIS, 120 mW , Coherent), both coupled in free-space through the back aperture of the microscope. Fluorescence emission was detected on an EMCCD camera (Ultra 897, Andor). Samples were imaged in a buffer containing 100mM Tris, 10mM NaCl, 550mM glucose, 1% (v/v) β -mercaptoethanol, 500 μ g/ml glucose oxidase (Sigma) and 40 μ g/ml catalase (Sigma).

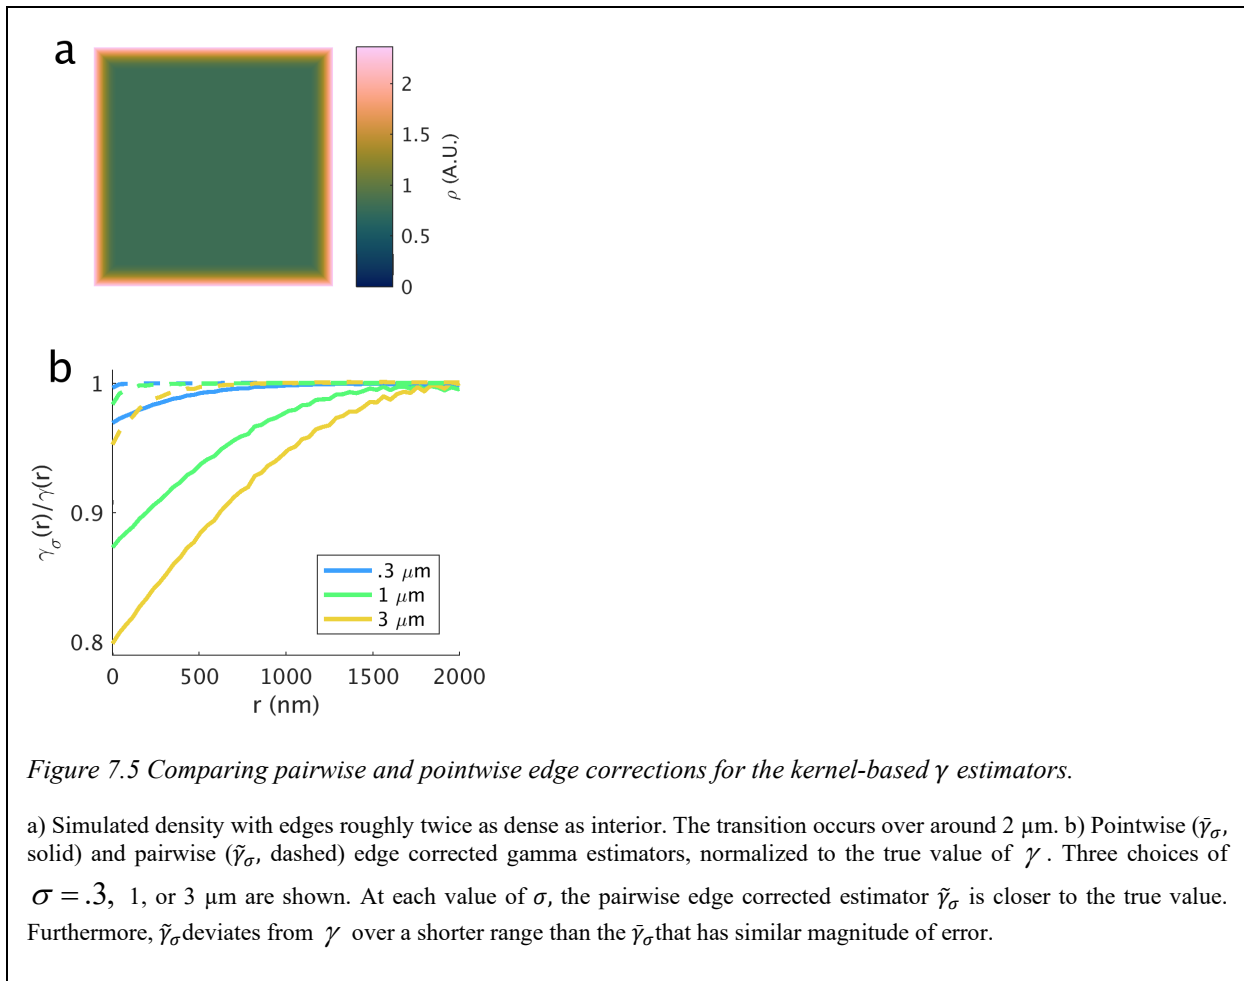
Localization of single molecules was performed using custom MATLAB software as described previously (271). Briefly, peaks were segmented using a standard wavelet algorithm (55) and segmented peaks were then fit on GPUs using previously described algorithms based maximum likelihood estimation (58). After localization, points were culled to remove outliers. Finally, stage drift was corrected by applying the method of (271). Images were rendered by generating 2D histograms from localizations followed by convolution with a Gaussian for display purposes.

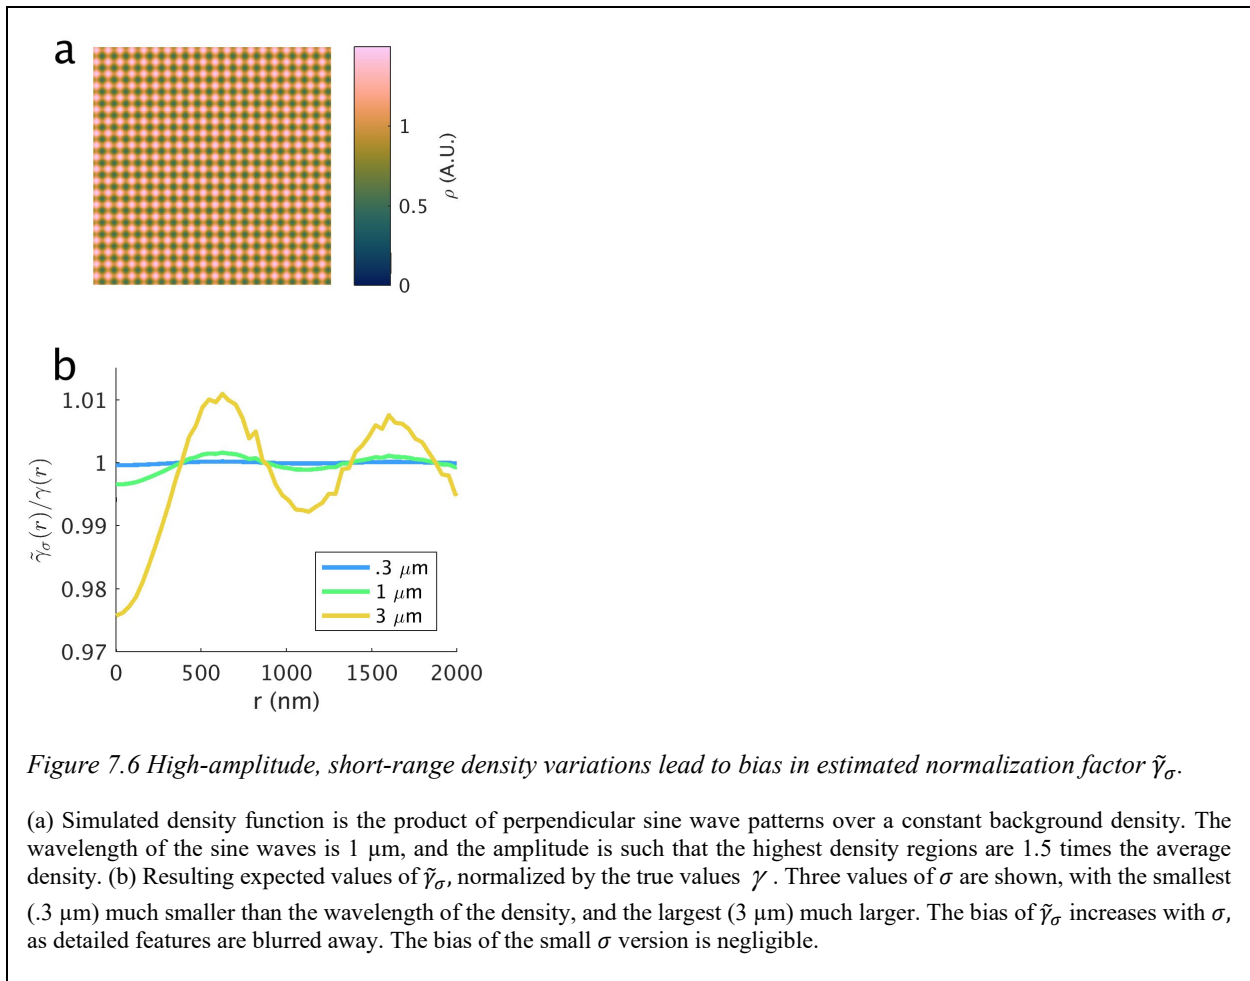












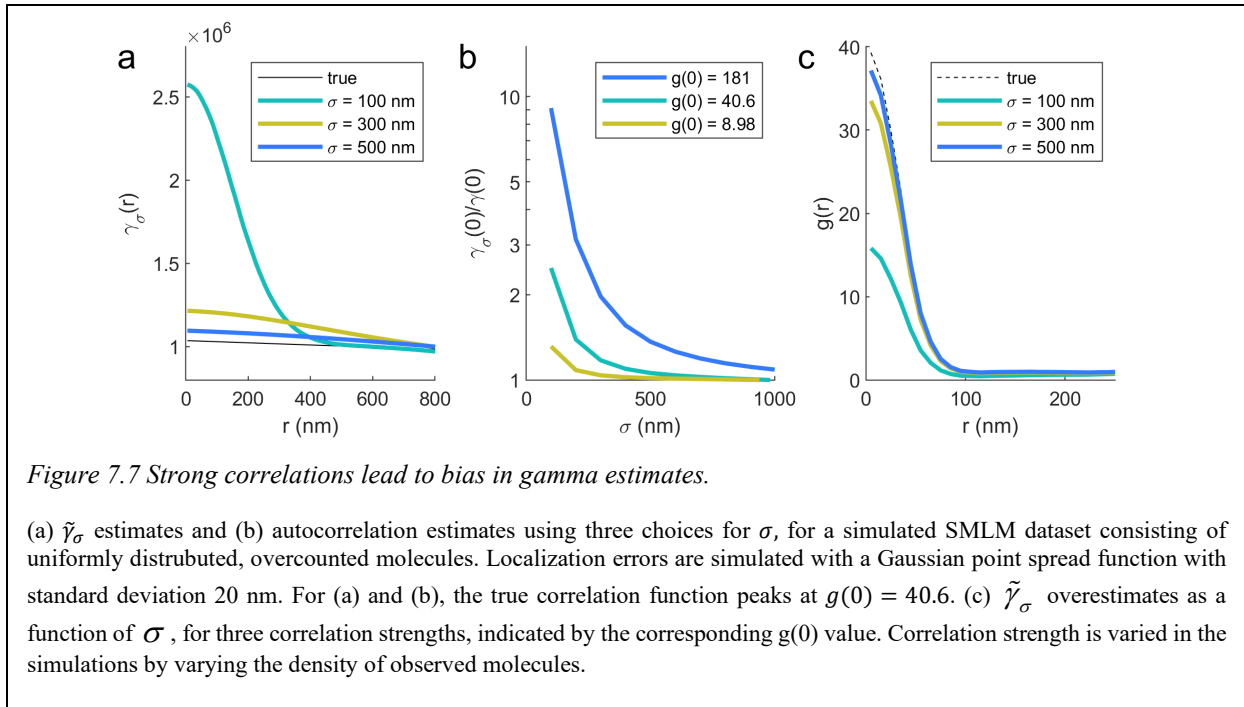


Figure 7.7 Strong correlations lead to bias in gamma estimates.

(a) $\tilde{\gamma}_\sigma$ estimates and (b) autocorrelation estimates using three choices for σ , for a simulated SMLM dataset consisting of uniformly distributed, overcounted molecules. Localization errors are simulated with a Gaussian point spread function with standard deviation 20 nm. For (a) and (b), the true correlation function peaks at $g(0) = 40.6$. (c) $\tilde{\gamma}_\sigma$ overestimates as a function of σ , for three correlation strengths, indicated by the corresponding $g(0)$ value. Correlation strength is varied in the simulations by varying the density of observed molecules.

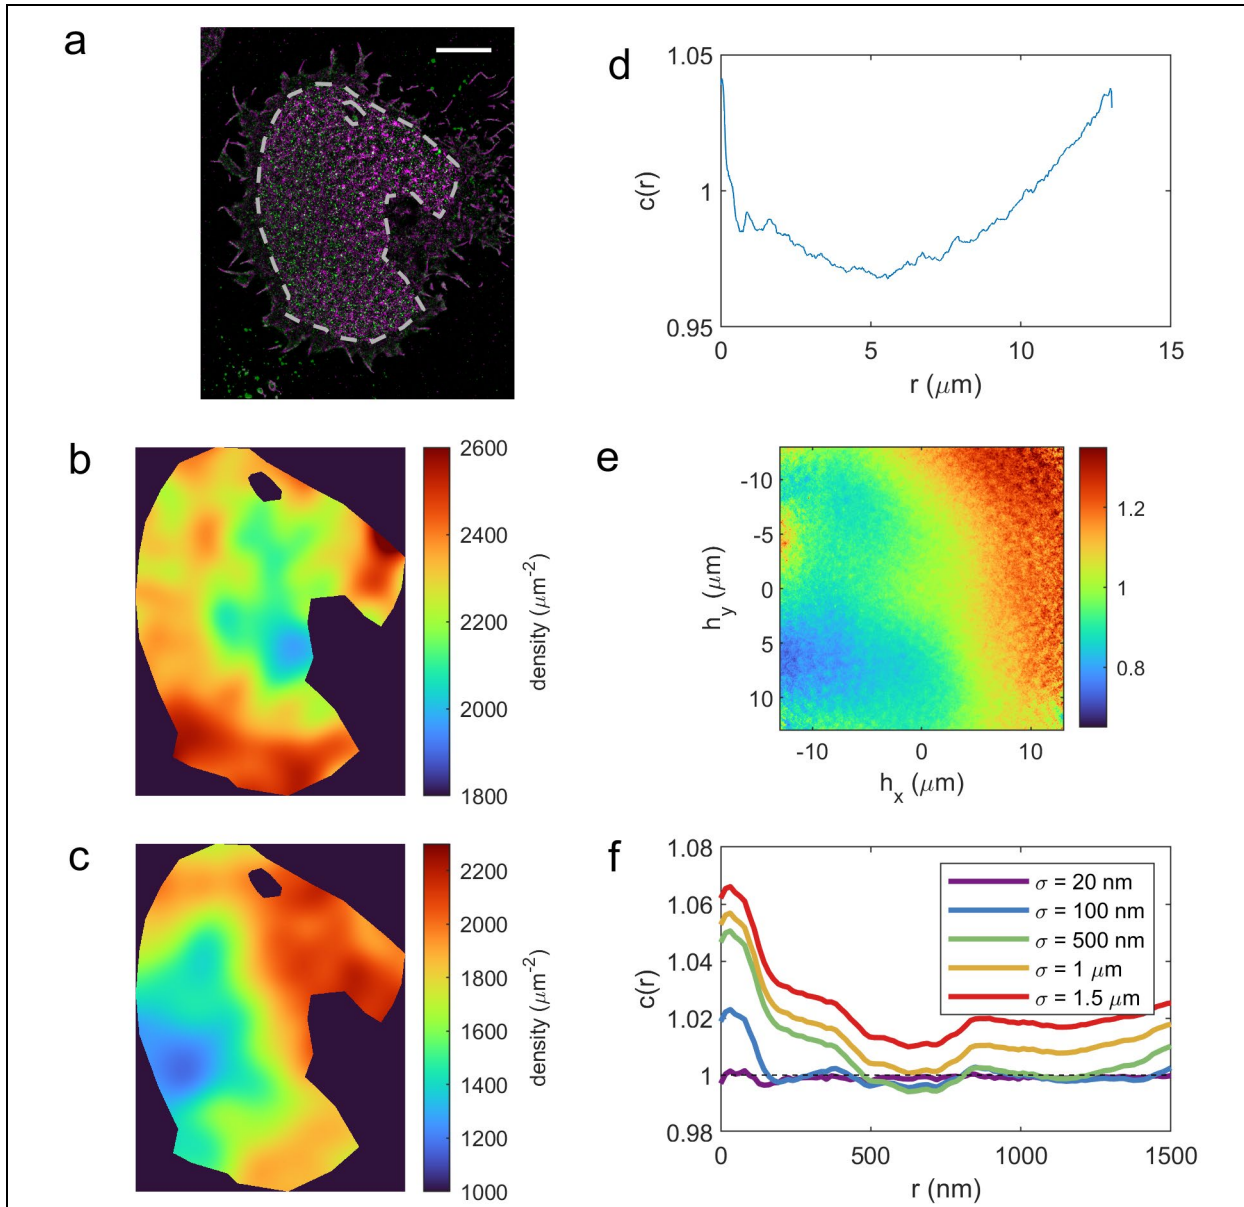


Figure 7.8 Density corrected cross-correlations of phalloidin-AlexaFluor647 and cholera toxin B subunit-AlexaFluor532 labeling a CH27 B cell.

(a) Two color reconstructed TIRF dSTORM image of cholera toxin B subunit (CTxB), which binds to the ganglioside GM1 at the cell surface (green) and phalloidin, which labels actin (magenta). Scale bar is $5 \mu\text{m}$. (b,c) Kernel-based density estimates for the densities in the (b) CTxB and (c) phalloidin channel, respectively, using a smoothing length-scale of $\sigma = 1 \mu\text{m}$. In both cases, regions of high and low density are apparent. (d) Uncorrected cross-correlation $\hat{c}_{\text{unif}}(r)$ shows strong deviations from 1 at long distances. (e) The uncorrected 2d cross-correlation $\hat{c}_{\text{unif}}(\mathbf{h})$ shows strong and anisotropic deviations from 1 and long distances. (f) Globally density-corrected cross-correlations $\hat{c}_{\text{global}}(r)$, using the pairwise estimator $\tilde{\gamma}$, with σ values as indicated.

Chapter 8 Conclusions

In the foregoing chapters I have presented a collection of projects that lead toward a clearer picture of the chemical and organizational properties of biological membranes. I will conclude with a few remarks on the main results of each project, the direct implications for membrane biophysics, and the future work that is enabled by the methods developed here.

In chapter 3, I present measurements of the chemical potential of cholesterol in synthetic membranes with a range of compositions. These compositions span a wide range of cholesterol concentrations so that we may draw direct conclusions about non-idealities of the concentration dependence in these membranes. Indeed, we observe that mixtures of cholesterol with different phospholipids reveal qualitatively different concentration dependence of the chemical potential. Mixtures with a fully unsaturated phospholipid, for example, are close to the ideal mixing behavior, but mixtures with a fully saturated phospholipid reveal very favorable interactions at low cholesterol concentrations, rapidly transitioning to unfavorable interactions at higher cholesterol concentrations. A third phospholipid with one saturated and one unsaturated acyl tail shows the opposite behavior, with interactions becoming more favorable as cholesterol concentration is increased. Notably, we find that the different phospholipid interactions lead to meaningful differences in cholesterol chemical potential, which can be as large as $1.5 k_B T$ at a fixed cholesterol concentration. To contextualize the size of this difference, it could correspond to a factor of 5 difference in cholesterol binding to a low-affinity binding site. Our measurements can be compared directly to measurements of cholesterol chemical potential in red blood cells (223). We also present measurements in three-component mixtures of cholesterol with two phospholipids. These measurements include compositions that phase separate, and we observe constant chemical potential along tie lines in the two-phase coexistence region, as required for equilibrium. Furthermore, contours of the chemical potential outside of the coexistence region are roughly parallel to the tie lines, suggesting that the interactions that determine the chemical potential in that regime are similar to those that cause phase separation in these mixtures. Future experiments could directly characterize this connection by measuring the cholesterol chemical potential in conditions such as alkanol treatments that perturb the boundary of the two phase region (333), with

the expectation that cholesterol chemical potential is well predicted by the distance in composition from the two-phase region. Because cell plasma membranes reside near a similar region of two phase coexistence, we also hypothesize that the chemical potential of cholesterol in cells will follow the qualitative trends observed in these simpler systems. The close connection between the chemical potential of one component and the abundance of other species in the membrane suggests that many common membrane-active treatments could be acting by modulating the chemical potentials of other species and thus affecting the allosteric modulation of some protein by that species. It will be important to develop experiments to distinguish these indirect effects from direct modulation by the treatment in question.

We explored a family of simple thermodynamic models based on regular solutions theory, and found them to be largely consistent with the measured chemical potentials. Fitting these models to the data led to useful conclusions about the energy scales of interactions implied by the best fit parameters. In particular, we can conclude that, for a popular model that involves stoichiometric binding of cholesterol to saturated phospholipids, the binding affinity of cholesterol to the lipid is constrained by this data to be fairly weak, in contrast to the expectations from other kinds of data (255, 257, 258).

One further avenue for extending these measurements could address the determinants of the distribution of cholesterol between the two leaflets of asymmetric membranes. Asymmetric membranes have received renewed attention in recent years with increased appreciation for the fact that lipid compositions and protein biophysical properties are quite different in the two leaflets of mammalian plasma membranes (261). It is difficult to directly measure the asymmetric distribution of cholesterol because it can flip between the two leaflets much faster than experimental timescales, but it has been noted that the phospholipids of the outer leaflet provide a more favorable environment for cholesterol than the inner leaflet, suggesting that cholesterol mostly resides in the outer leaflet. The methods described here can be readily extended to give insight into this question, by performing separate measurements of cholesterol chemical potential in membrane compositions that mimic the inner and outer leaflet, respectively. An estimate for the distribution of cholesterol between the leaflets can then be obtained by finding the chemical potential such that the average of the corresponding inner and outer leaflet cholesterol concentrations coincides with the cholesterol concentration overall in the asymmetric membrane.

(Note that this is not the whole story, since it neglects the contribution of differential stress induced by having more lipids in one leaflet than the other (262)).

The remaining chapters present method developments for single molecule localization microscopy. Although these methods are broadly applicable to improving the quality and analysis of SMLM datasets, they are undertaken with the primary goal of increasing the accessibility of spatially detailed and subtle, low-contrast organization of proteins on the plasma membrane.

In chapter 4, I describe a new method for correcting microscope stage drift in SMLM datasets. This method relies on a simple and robust mean shift algorithm for estimating the stage drift from subsets of the SMLM data. It is shown to be in some cases more robust than existing methods (61), so that drift can be estimated from smaller subsets of the data, with higher time resolution. Good time resolution is essential for correcting nonlinear drift, which we do observe in our microscopes. In addition, the reduced computational requirements of this method enable 3d drift correction, for which previously reported methods either required vast computational resources or degraded precision. Together, these developments improve the precision of drift correction, and therefore the ultimate precision of the SMLM dataset. This improved precision implies that finer details of biological samples can be distinguished, including weak colocalization of membrane components.

Chapter 5 defines an objective metric called the localization spread function for evaluating the quality of SMLM datasets. By analogy with the PSF of a microscope, the LSF describes the distribution of localization errors with respect to the true location of the fluorophore that produces them, including the contributions from all of the steps of the SMLM analysis process. The LSF is also the primary determinant of the spatial scales that can be accessed by pair correlation analysis of SMLM data. We demonstrate how the LSF can be estimated directly from an SMLM dataset by exploiting the temporal characteristics of switching that characterize SMLM probes. Notably, our method for estimating the LSF can specifically distinguish the instantaneous localization precision from processes that induce additional spatial errors over a time delay, including imprecise correction of microscope stage drift. While it is standard practice to evaluate new methods on simulated datasets where the ground truth is known, the ability to determine the LSF directly and objectively from biological samples will be essential to evaluating the ways new methods interact with the idiosyncrasies of the sample of interest, as well as other aspects of the existing analysis pipeline. Notably the LSF is sensitive to the entire SMLM process: sample

preparation, microscope properties, imaging parameters and all of the steps of analysis, so can be used as an objective when improving any of those steps. Additionally, an objective metric is useful in training microscopists, by giving concrete feedback on how to choose imaging parameters for best image quality.

Finally, chapters 6 and 7 develop a method for estimating the pair correlation functions of target molecules that are observed through a sampling process that is not uniform in space. This situation can arise in SMLM, especially in biological samples that present excessive constraints on the optimization of the labeling procedure. We show that the new estimator has improved bias and variance compared to an existing estimator that applies to this situation (88), and demonstrate its application to simple point processes in chapter 6, and to SMLM datasets with overcounting in chapter 7. We anticipate that these methods will be invaluable for analyzing the organization of some fixed-cell SMLM datasets that are difficult to label uniformly. In addition, these methods have already been extended to spatiotemporal pair correlation functions in live cell SMLM datasets, where non-uniform sampling is frequently due to the photobleaching of diffusing probes. Probes that have diffused to the center of the illuminated region of the cell have on average been exposed to illumination for a longer time, and therefore are more likely to have photobleached. These methods are actively being applied to live cell SMLM in the Veatch lab (23, 91).

Together, the improvements to SMLM imaging represented by Chapters 4-7 will make way for more precise measurements of more subtle organization on biomembranes. This type of measurement has already been essential for establishing a link between the phase behavior of membranes and the organization of membrane proteins. In the future, more precise measurements will enable this link to be extended to new systems, especially systems where the interactions between the proteins of interest and the membrane are weak. In particular, I hope that SMLM will soon be able to detect correlations between single, unclustered proteins and markers of membrane phases.

More generally, this dissertation strengthens the connections between membrane phase behavior and biological function. This includes introducing a new link between the chemical potentials of membrane components and membrane phase behavior, with a focus on cholesterol, and also paving the way for future measurements of the organization of proteins on biomembranes.

Bibliography

1. Singer, S.J., and G.L. Nicolson. 1972. The Fluid Mosaic Model of the Structure of Cell Membranes. *Science*. 175:720–731.
2. Janiak, M.J., D.M. Small, and G.G. Shipley. 1979. Temperature and compositional dependence of the structure of hydrated dimyristoyl lecithin. *J. Biol. Chem.* 254:6068–6078.
3. Dietrich, C., L.A. Bagatolli, Z.N. Volovyk, N.L. Thompson, M. Levi, K. Jacobson, and E. Gratton. 2001. Lipid Rafts Reconstituted in Model Membranes. *Biophysical Journal*. 80:1417–1428.
4. Samsonov, A.V., I. Mihalyov, and F.S. Cohen. 2001. Characterization of Cholesterol-Sphingomyelin Domains and Their Dynamics in Bilayer Membranes. *Biophysical Journal*. 81:1486–1500.
5. Veatch, S.L., and S.L. Keller. 2002. Organization in Lipid Membranes Containing Cholesterol. *Phys. Rev. Lett.* 89:268101.
6. Veatch, S.L., and S.L. Keller. 2003. Separation of Liquid Phases in Giant Vesicles of Ternary Mixtures of Phospholipids and Cholesterol. *Biophys J.* 85:3074–3083.
7. Veatch, S.L., K. Gawrisch, and S.L. Keller. 2006. Closed-Loop Miscibility Gap and Quantitative Tie-Lines in Ternary Membranes Containing Diphytanoyl PC. *Biophys J.* 90:4428–4436.
8. Ionova, I.V., V.A. Livshits, and D. Marsh. 2012. Phase Diagram of Ternary Cholesterol/Palmitoylsphingomyelin/Palmitoyloleoyl-Phosphatidylcholine Mixtures: Spin-Label EPR Study of Lipid-Raft Formation. *Biophysical Journal*. 102:1856–1865.
9. Veatch, S.L., and S.L. Keller. 2005. Miscibility Phase Diagrams of Giant Vesicles Containing Sphingomyelin. *Phys. Rev. Lett.* 94:148101.
10. Konyakhina, T.M., J. Wu, J.D. Mastroianni, F.A. Heberle, and G.W. Feigenson. 2013. Phase diagram of a 4-component lipid mixture: DSPC/DOPC/POPC/chol. *Biochimica et Biophysica Acta (BBA) - Biomembranes*. 1828:2204–2214.
11. Heberle, F.A., J. Wu, S.L. Goh, R.S. Petruzielo, and G.W. Feigenson. 2010. Comparison of Three Ternary Lipid Bilayer Mixtures: FRET and ESR Reveal Nanodomains. *Biophysical Journal*. 99:3309–3318.

12. Hjort Ipsen, J., G. Karlström, O.G. Mourtsen, H. Wennerström, and M.J. Zuckermann. 1987. Phase equilibria in the phosphatidylcholine-cholesterol system. *Biochimica et Biophysica Acta (BBA) - Biomembranes*. 905:162–172.
13. Vist, M.R., and J.H. Davis. 1990. Phase equilibria of cholesterol/dipalmitoylphosphatidylcholine mixtures: deuterium nuclear magnetic resonance and differential scanning calorimetry. *Biochemistry*. 29:451–464.
14. Almeida, P.F., W.L. Vaz, and T.E. Thompson. 1992. Lateral diffusion in the liquid phases of dimyristoylphosphatidylcholine/cholesterol lipid bilayers: a free volume analysis. *Biochemistry*. 31:6739–6747.
15. Khadka, N.K., C.S. Ho, and J. Pan. 2015. Macroscopic and Nanoscopic Heterogeneous Structures in a Three-Component Lipid Bilayer Mixtures Determined by Atomic Force Microscopy. *Langmuir*. 31:12417–12425.
16. Sezgin, E., I. Levental, M. Grzybek, G. Schwarzmann, V. Mueller, A. Honigmann, V.N. Belov, C. Eggeling, Ü. Coskun, K. Simons, and P. Schwille. 2012. Partitioning, diffusion, and ligand binding of raft lipid analogs in model and cellular plasma membranes. *Biochimica et Biophysica Acta (BBA) - Biomembranes*. 1818:1777–1784.
17. Sodt, A.J., M.L. Sandar, K. Gawrisch, R.W. Pastor, and E. Lyman. 2014. The Molecular Structure of the Liquid-Ordered Phase of Lipid Bilayers. *J. Am. Chem. Soc.* 136:725–732.
18. Sodt, A.J., R.W. Pastor, and E. Lyman. 2015. Hexagonal Substructure and Hydrogen Bonding in Liquid-Ordered Phases Containing Palmitoyl Sphingomyelin. *Biophysical Journal*. 109:948–955.
19. Levental, I., F.J. Byfield, P. Chowdhury, F. Gai, T. Baumgart, and P.A. Janmey. 2009. Cholesterol-dependent phase separation in cell-derived giant plasma-membrane vesicles. *Biochem J*. 424:163–167.
20. Veatch, S.L., P. Cicuta, P. Sengupta, A. Honerkamp-Smith, D. Holowka, and B. Baird. 2008. Critical Fluctuations in Plasma Membrane Vesicles. *ACS Chem. Biol.* 3:287–293.
21. Burns, M., K. Wisser, J. Wu, I. Levental, and S.L. Veatch. 2017. Miscibility Transition Temperature Scales with Growth Temperature in a Zebrafish Cell Line. *Biophysical Journal*. 113:1212–1222.
22. Chaikin, P.M., and T.C. Lubensky. 2010. Principles of condensed matter physics. 1. paperback ed. (with corr.), 5. print. Cambridge: Cambridge Univ. Press.
23. Shelby, S.A., I. Castello-Serrano, K.C. Wisser, I. Levental, and S.L. Veatch. 2023. Membrane phase separation drives responsive assembly of receptor signaling domains. *Nat Chem Biol*. 1–9.

24. Brannigan, G., J. Hénin, R. Law, R. Eckenhoff, and M.L. Klein. 2008. Embedded cholesterol in the nicotinic acetylcholine receptor. *Proceedings of the National Academy of Sciences*. 105:14418–14423.
25. Baenziger, J.E., J.A. Domville, and J.P.D. Therien. 2017. The Role of Cholesterol in the Activation of Nicotinic Acetylcholine Receptors. *Curr Top Membr*. 80:95–137.
26. Fantini, J., and F.J. Barrantes. 2009. Sphingolipid/cholesterol regulation of neurotransmitter receptor conformation and function. *Biochimica et Biophysica Acta (BBA) - Biomembranes*. 1788:2345–2361.
27. Cang, X., Y. Du, Y. Mao, Y. Wang, H. Yang, and H. Jiang. 2013. Mapping the Functional Binding Sites of Cholesterol in β 2-Adrenergic Receptor by Long-Time Molecular Dynamics Simulations. *The Journal of Physical Chemistry B*. 117:1085–1094.
28. Budelier, M.M., W.W.L. Cheng, Z.-W. Chen, J.R. Bracamontes, Y. Sugasawa, K. Kathiresan, L. Mydock-McGrane, D.F. Covey, and A.S. Evers. 2019. Common binding sites for cholesterol and neurosteroids on a pentameric ligand-gated ion channel. *Biochim Biophys Acta Mol Cell Biol Lipids*. 1864:128–136.
29. Hénin, J., R. Salari, S. Murlidaran, and G. Brannigan. 2014. A Predicted Binding Site for Cholesterol on the GABAA Receptor. *Biophysical Journal*. 106:1938–1949.
30. Ufret-Vincenty, C.A., R.M. Klein, L. Hua, J. Angueyra, and S.E. Gordon. 2011. Localization of the PIP2 Sensor of TRPV1 Ion Channels. *J. Biol. Chem*. 286:9688–9698.
31. Stansfeld, P.J., R. Hopkinson, F.M. Ashcroft, and M.S.P. Sansom. 2009. PIP2-Binding Site in Kir Channels: Definition by Multiscale Biomolecular Simulations. *Biochemistry*. 48:10926–10933.
32. Seebohm, G., E. Wrobel, M. Pusch, M. Dicks, J. Terhag, V. Matschke, I. Rothenberg, O.N. Ursu, F. Hertel, L. Pott, F. Lang, E. Schulze-Bahr, M. Hollmann, R. Stoll, and N. Strutz-Seebohm. 2014. Structural basis of PI(4,5)P2-dependent regulation of GluA1 by phosphatidylinositol-5-phosphate 4-kinase, type II, alpha (PIP5K2A). *Pflugers Arch - Eur J Physiol*. 466:1885–1897.
33. Yeagle, P.L., and J.E. Young. 1986. Factors contributing to the distribution of cholesterol among phospholipid vesicles. *Journal of Biological Chemistry*. 261:8175–8181.
34. Niu, S.-L., and B.J. Litman. 2002. Determination of Membrane Cholesterol Partition Coefficient Using a Lipid Vesicle–Cyclodextrin Binary System: Effect of Phospholipid Acyl Chain Unsaturation and Headgroup Composition. *Biophysical Journal*. 83:3408–3415.
35. Tsamaloukas, A., H. Szadkowska, P.J. Slotte, and H. Heerklotz. 2005. Interactions of Cholesterol with Lipid Membranes and Cyclodextrin Characterized by Calorimetry. *Biophysical Journal*. 89:1109–1119.

36. Pike, L.J. 2006. Rafts defined: a report on the Keystone Symposium on Lipid Rafts and Cell Function. *J Lipid Res.* 47:1597–8.
37. Betzig, E., G.H. Patterson, R. Sougrat, O.W. Lindwasser, S. Olenych, J.S. Bonifacino, M.W. Davidson, J. Lippincott-Schwartz, and H.F. Hess. 2006. Imaging intracellular fluorescent proteins at nanometer resolution. *Science.* 313:1642–1645.
38. Rust, M.J., M. Bates, and X. Zhuang. 2006. Sub-diffraction-limit imaging by stochastic optical reconstruction microscopy (STORM). *Nat Methods.* 3:793–795.
39. Hess, S.T., T.P.K. Girirajan, and M.D. Mason. 2006. Ultra-high resolution imaging by fluorescence photoactivation localization microscopy. *Biophys J.* 91:4258–4272.
40. Heilemann, M., S. van de Linde, M. Schüttpelz, R. Kasper, B. Seefeldt, A. Mukherjee, P. Tinnefeld, and M. Sauer. 2008. Subdiffraction-Resolution Fluorescence Imaging with Conventional Fluorescent Probes. *Angewandte Chemie International Edition.* 47:6172–6176.
41. Heilemann, M., S. van de Linde, A. Mukherjee, and M. Sauer. 2009. Super-Resolution Imaging with Small Organic Fluorophores. *Angewandte Chemie International Edition.* 48:6903–6908.
42. Lukinavičius, G., K. Umezawa, N. Olivier, A. Honigmann, G. Yang, T. Plass, V. Mueller, L. Reymond, I.R.C. Jr, Z.-G. Luo, C. Schultz, E.A. Lemke, P. Heppenstall, C. Eggeling, S. Manley, and K. Johnsson. 2013. A near-infrared fluorophore for live-cell super-resolution microscopy of cellular proteins. *Nature Chemistry.* 5:132.
43. Pillai, V.N.R. 1980. Photoremovable Protecting Groups in Organic Synthesis. *Synthesis.* 1980:1–26.
44. Belov, V.N., C.A. Wurm, V.P. Boyarskiy, S. Jakobs, and S.W. Hell. 2010. Rhodamines NN: A Novel Class of Caged Fluorescent Dyes. *Angewandte Chemie International Edition.* 49:3520–3523.
45. Patterson, G.H., and J. Lippincott-Schwartz. 2002. A Photoactivatable GFP for Selective Photolabeling of Proteins and Cells. *Science.* 297:1873–1877.
46. Subach, F.V., G.H. Patterson, S. Manley, J.M. Gillette, J. Lippincott-Schwartz, and V.V. Verkhusha. 2009. Photoactivatable mCherry for high-resolution two-color fluorescence microscopy. *Nat Methods.* 6:153–159.
47. Sharonov, A., and R.M. Hochstrasser. 2006. Wide-field subdiffraction imaging by accumulated binding of diffusing probes. *PNAS.* 103:18911–18916.
48. Jungmann, R., C. Steinhauer, M. Scheible, A. Kuzyk, P. Tinnefeld, and F.C. Simmel. 2010. Single-Molecule Kinetics and Super-Resolution Microscopy by Fluorescence Imaging of Transient Binding on DNA Origami. *Nano Lett.* 10:4756–4761.

49. Small, A.R. 2009. Theoretical limits on errors and acquisition rates in localizing switchable fluorophores. *Biophysical journal*. 96:L16–L18.
50. Saxton, M.J., and K. Jacobson. 1997. SINGLE-PARTICLE TRACKING: Applications to Membrane Dynamics. *Annual Review of Biophysics and Biomolecular Structure*. 26:373–399.
51. Shelby, S.A., D. Holowka, B. Baird, and S.L. Veatch. 2013. Distinct Stages of Stimulated FcεRI Receptor Clustering and Immobilization Are Identified through Superresolution Imaging. *Biophysical Journal*. 105:2343–2354.
52. Chen, Z., L. Geffroy, and J.S. Biteen. 2021. NOBIAS: Analyzing Anomalous Diffusion in Single-Molecule Tracks With Nonparametric Bayesian Inference. *Frontiers in Bioinformatics*. 1.
53. Stone, M.B., and S.L. Veatch. 2015. Steady-state cross-correlations for live two-colour super-resolution localization data sets. *Nat Commun*. 6.
54. Small, A., and S. Stahlheber. 2014. Fluorophore localization algorithms for super-resolution microscopy. *Nature Methods*. 11:267–279.
55. Izeddin, I., J. Boulanger, V. Racine, C.G. Specht, A. Kechkar, D. Nair, A. Triller, D. Choquet, M. Dahan, and J.B. Sibarita. 2012. Wavelet analysis for single molecule localization microscopy. *Opt. Express, OE*. 20:2081–2095.
56. Ober, R.J., S. Ram, and E.S. Ward. 2004. Localization Accuracy in Single-Molecule Microscopy. *Biophysical Journal*. 16.
57. Abraham, A.V., S. Ram, J. Chao, E.S. Ward, and R.J. Ober. 2009. Quantitative study of single molecule location estimation techniques. *Optics Express*. 17:23352.
58. Smith, C.S., N. Joseph, B. Rieger, and K.A. Lidke. 2010. Fast, single-molecule localization that achieves theoretically minimum uncertainty. *Nature Methods*. 7:373–375.
59. Mlodzianoski, M.J., J.M. Schreiner, S.P. Callahan, K. Smolková, A. Dlasková, J. Šantorová, P. Ježek, and J. Bewersdorf. 2011. Sample drift correction in 3D fluorescence photoactivation localization microscopy. *Opt. Express, OE*. 19:15009–15019.
60. Geisler, C., T. Hotz, A. Schönle, S.W. Hell, A. Munk, and A. Egner. 2012. Drift estimation for single marker switching based imaging schemes. *Opt. Express*. 20:7274.
61. Wang, Y., J. Schnitzbauer, Z. Hu, X. Li, Y. Cheng, Z.-L. Huang, and B. Huang. 2014. Localization events-based sample drift correction for localization microscopy with redundant cross-correlation algorithm. *Opt. Express, OE*. 22:15982–15991.
62. Colomb, W., J. Czerski, J.D. Sau, and S.K. Sarker. 2017. Estimation of microscope drift using fluorescent nanodiamonds as fiducial markers: ESTIMATION OF MICROSCOPE DRIFT. *Journal of Microscopy*. 266:298–306.

63. Balinovic, A., D. Albrecht, and U. Endesfelder. 2019. Spectrally red-shifted fluorescent fiducial markers for optimal drift correction in localization microscopy. *J. Phys. D: Appl. Phys.* 52:204002.
64. Elmokadem, A., and J. Yu. 2015. Optimal Drift Correction for Superresolution Localization Microscopy with Bayesian Inference. *Biophysical Journal.* 109:1772–1780.
65. Schlangen, I., J. Franco, J. Houssineau, W.T.E. Pitkeathly, D. Clark, I. Smal, and C. Rickman. 2016. Marker-Less Stage Drift Correction in Super-Resolution Microscopy Using the Single-Cluster PHD Filter. *IEEE J. Sel. Top. Signal Process.* 10:193–202.
66. Churchman, L.S., Z. Ökten, R.S. Rock, J.F. Dawson, and J.A. Spudich. 2005. Single molecule high-resolution colocalization of Cy3 and Cy5 attached to macromolecules measures intramolecular distances through time. *Proceedings of the National Academy of Sciences.* 102:1419–1423.
67. Carrington, G., D. Tomlinson, and M. Peckham. 2019. Exploiting nanobodies and Affimers for superresolution imaging in light microscopy. *MBoC.* 30:2737–2740.
68. Götzke, H., M. Kilisch, M. Martínez-Carranza, S. Sograte-Idrissi, A. Rajavel, T. Schlichthaerle, N. Engels, R. Jungmann, P. Stenmark, F. Opazo, and S. Frey. 2019. The ALFA-tag is a highly versatile tool for nanobody-based bioscience applications. *Nat Commun.* 10:4403.
69. Wagner, T.R., and U. Rothbauer. 2021. Nanobodies – Little helpers unravelling intracellular signaling. *Free Radical Biology and Medicine.* 176:46–61.
70. Los, G.V., L.P. Encell, M.G. McDougall, D.D. Hartzell, N. Karassina, C. Zimprich, M.G. Wood, R. Learish, R.F. Ohana, M. Urh, D. Simpson, J. Mendez, K. Zimmerman, P. Otto, G. Vidugiris, J. Zhu, A. Darzins, D.H. Klaubert, R.F. Bulleit, and K.V. Wood. 2008. HaloTag: A Novel Protein Labeling Technology for Cell Imaging and Protein Analysis. *ACS Chem. Biol.* 3:373–382.
71. Grimm, J.B., B.P. English, J. Chen, J.P. Slaughter, Z. Zhang, A. Revyakin, R. Patel, J.J. Macklin, D. Normanno, R.H. Singer, T. Lionnet, and L.D. Lavis. 2015. A general method to improve fluorophores for live-cell and single-molecule microscopy. *Nat Methods.* 12:244–250.
72. Grimm, J.B., T.A. Brown, B.P. English, T. Lionnet, and L.D. Lavis. 2017. Synthesis of Janelia Fluor HaloTag and SNAP-Tag Ligands and Their Use in Cellular Imaging Experiments. In: Erfle H, editor. *Super-Resolution Microscopy: Methods and Protocols.* New York, NY: Springer New York. pp. 179–188.
73. Zheng, Y., Z. Ye, Z. Liu, W. Yang, X. Zhang, Y. Yang, and Y. Xiao. 2021. Nitroso-Caged Rhodamine: A Superior Green Light-Activatable Fluorophore for Single-Molecule Localization Super-Resolution Imaging. *Anal. Chem.* 93:7833–7842.

74. Pavani, S.R.P., M.A. Thompson, J.S. Biteen, S.J. Lord, N. Liu, R.J. Twieg, R. Piestun, and W.E. Moerner. 2009. Three-dimensional, single-molecule fluorescence imaging beyond the diffraction limit by using a double-helix point spread function. *Proceedings of the National Academy of Sciences*. 106:2995–2999.
75. Huang, F., G. Sirinakis, E.S. Allgeyer, L.K. Schroeder, W.C. Duim, E.B. Kromann, T. Phan, F.E. Rivera-Molina, J.R. Myers, I. Irnov, M. Lessard, Y. Zhang, M.A. Handel, C. Jacobs-Wagner, C.P. Lusk, J.E. Rothman, D. Toomre, M.J. Booth, and J. Bewersdorf. 2016. Ultra-High Resolution 3D Imaging of Whole Cells. *Cell*. 166:1028–1040.
76. Mlodzianoski, M.J., N.M. Curthoys, M.S. Gunewardene, S. Carter, and S.T. Hess. 2016. Super-Resolution Imaging of Molecular Emission Spectra and Single Molecule Spectral Fluctuations. *PLOS ONE*. 11:e0147506.
77. Fazel, M., M.J. Wester, H. Mazloom-Farsibaf, M.B.M. Meddens, A.S. Eklund, T. Schlichthaerle, F. Schueder, R. Jungmann, and K.A. Lidke. 2019. Bayesian Multiple Emitter Fitting using Reversible Jump Markov Chain Monte Carlo. *Sci Rep*. 9:13791.
78. Fazel, M., M.J. Wester, D.J. Schodt, S.R. Cruz, S. Strauss, F. Schueder, T. Schlichthaerle, J.M. Gillette, D.S. Lidke, B. Rieger, R. Jungmann, and K.A. Lidke. 2022. High-precision estimation of emitter positions using Bayesian grouping of localizations. *Nat Commun*. 13:7152.
79. Speiser, A., L.-R. Müller, P. Hoess, U. Matti, C.J. Obara, W.R. Legant, A. Kreshuk, J.H. Macke, J. Ries, and S.C. Turaga. 2021. Deep learning enables fast and dense single-molecule localization with high accuracy. *Nat Methods*. 18:1082–1090.
80. Banterle, N., K.H. Bui, E.A. Lemke, and M. Beck. 2013. Fourier ring correlation as a resolution criterion for super-resolution microscopy. *Journal of Structural Biology*. 183:363–367.
81. Nieuwenhuizen, R.P.J., K.A. Lidke, M. Bates, D.L. Puig, D. Grünwald, S. Stallinga, and B. Rieger. 2013. Measuring image resolution in optical nanoscopy. *Nature Methods*. 10:557–562.
82. Endesfelder, U., S. Malkusch, F. Fricke, and M. Heilemann. 2014. A simple method to estimate the average localization precision of a single-molecule localization microscopy experiment. *Histochem Cell Biol*. 141:629–638.
83. Stone, M.B., S.A. Shelby, M.F. Núñez, K. Wisser, and S.L. Veatch. 2017. Protein sorting by lipid phase-like domains supports emergent signaling function in B lymphocyte plasma membranes. *eLife Sciences*. 6:e19891.
84. Núñez, M.F., K. Wisser, and S.L. Veatch. 2019. Synergistic factors control kinase–phosphatase organization in B-cells engaged with supported bilayers. *MBoC*. 31:667–682.
85. Møller, J., and R.P. Waagepetersen. 2003. Statistical Inference and Simulation for Spatial Point Processes. CRC Press.

86. Illian, J., A. Penttinen, H. Stoyan, and D. Stoyan. 2008. Statistical analysis and modelling of spatial point patterns. Chichester, England ; Hoboken, NJ: John Wiley.
87. Veatch, S.L., B.B. Machta, S.A. Shelby, E.N. Chiang, D.A. Holowka, and B.A. Baird. 2012. Correlation Functions Quantify Super-Resolution Images and Estimate Apparent Clustering Due to Over-Counting. *PLOS ONE*. 7:e31457.
88. Baddeley, A.J., J. Møller, and R. Waagepetersen. 2000. Non- and semi-parametric estimation of interaction in inhomogeneous point patterns. *Statistica Neerlandica*. 54:329–350.
89. Gabriel, E., and P.J. Diggle. 2009. Second-order analysis of inhomogeneous spatio-temporal point process data. *Statistica Neerlandica*. 63:43–51.
90. Møller, J., and M. Ghorbani. 2012. Aspects of second-order analysis of structured inhomogeneous spatio-temporal point processes. *Statistica Neerlandica*. 66:472–491.
91. Shelby, S.A., T.R. Shaw, and S.L. Veatch. 2023. Measuring the co-localization and dynamics of mobile proteins in live cells undergoing signaling responses. In: Baldari CT, ML Dustin, editors. *The Immune Synapse: Methods and Protocols*. New York, NY: Humana.
92. Shaw, T.R., S. Ghosh, and S.L. Veatch. 2021. Critical Phenomena in Plasma Membrane Organization and Function. *Annual Review of Physical Chemistry*. 72:51–72.
93. Simons, K., and E. Ikonen. 1997. Functional rafts in cell membranes. *Nature*. 387:569–572.
94. Baumgart, T., A.T. Hammond, P. Sengupta, S.T. Hess, D.A. Holowka, B.A. Baird, and W.W. Webb. 2007. Large-scale fluid/fluid phase separation of proteins and lipids in giant plasma membrane vesicles. *PNAS*. 104:3165–3170.
95. Veatch, S.L., and S.L. Keller. 2005. Seeing spots: Complex phase behavior in simple membranes. *Biochimica et Biophysica Acta (BBA) - Molecular Cell Research*. 1746:172–185.
96. Oldfield, E., and D. Chapman. 1971. Effects of cholesterol and cholesterol derivatives on hydrocarbon chain mobility in lipids. *Biochemical and Biophysical Research Communications*. 43:610–616.
97. Oldfield, E., and D. Chapman. 1972. Dynamics of lipids in membranes: Heterogeneity and the role of cholesterol. *FEBS Lett*. 23:285–297.
98. Engelman, D.M., and J.E. Rothman. 1972. The planar organization of lecithin-cholesterol bilayers. *J. Biol. Chem*. 247:3694–3697.
99. Shimshick, E.J., and H.M. McConnell. 1973. Lateral phase separations in binary mixtures of cholesterol and phospholipids. *Biochemical and Biophysical Research Communications*. 53:446–451.

100. Chapman, D., and S.A. Penkett. 1966. Nuclear magnetic resonance spectroscopic studies of the interaction of phospholipids with cholesterol. *Nature*. 211:1304–1305.
101. Ladbroke, B.D., R.M. Williams, and D. Chapman. 1968. Studies on lecithin-cholesterol-water interactions by differential scanning calorimetry and X-ray diffraction. *Biochimica et Biophysica Acta (BBA) - Biomembranes*. 150:333–340.
102. Opella, S.J., J.P. Yesinowski, and J.S. Waugh. 1976. Nuclear magnetic resonance description of molecular motion and phase separations of cholesterol in lecithin dispersions. *PNAS*. 73:3812–3815.
103. Lentz, B.R., D.A. Barrow, and M. Hoehli. 1980. Cholesterol-phosphatidylcholine interactions in multilamellar vesicles. *Biochemistry*. 19:1943–1954.
104. Rubenstein, J.L.R., J.C. Owicki, and H.M. McConnell. 1980. Dynamic properties of binary mixtures of phosphatidylcholines and cholesterol. *Biochemistry*. 19:569–573.
105. Recktenwald, D.J., and H.M. McConnell. 1981. Phase equilibria in binary mixtures of phosphatidylcholine and cholesterol. *Biochemistry*. 20:4505–4510.
106. Alecio, M.R., D.E. Golan, W.R. Veatch, and R.R. Rando. 1982. Use of a fluorescent cholesterol derivative to measure lateral mobility of cholesterol in membranes. *Proc. Natl. Acad. Sci. U.S.A.* 79:5171–5174.
107. Huang, T.H., C.W. Lee, S.K. Das Gupta, A. Blume, and R.G. Griffin. 1993. A ¹³C and ²H nuclear magnetic resonance study of phosphatidylcholine/cholesterol interactions: characterization of liquid-gel phases. *Biochemistry*. 32:13277–13287.
108. McMullen, T.P.W., and R.N. McElhaney. 1995. New aspects of the interaction of cholesterol with dipalmitoylphosphatidylcholine bilayers as revealed by high-sensitivity differential scanning calorimetry. *Biochimica et Biophysica Acta (BBA) - Biomembranes*. 1234:90–98.
109. Ahmed, S.N., D.A. Brown, and E. London. 1997. On the Origin of Sphingolipid/Cholesterol-Rich Detergent-Insoluble Cell Membranes: Physiological Concentrations of Cholesterol and Sphingolipid Induce Formation of a Detergent-Insoluble, Liquid-Ordered Lipid Phase in Model Membranes. *Biochemistry*. 36:10944–10953.
110. Feigenson, G.W., and J.T. Buboltz. 2001. Ternary Phase Diagram of Dipalmitoyl-PC/Dilauroyl-PC/Cholesterol: Nanoscopic Domain Formation Driven by Cholesterol. *Biophysical Journal*. 80:2775–2788.
111. van Meer, G., and K. Simons. 1982. Viruses budding from either the apical or the basolateral plasma membrane domain of MDCK cells have unique phospholipid compositions. *EMBO J.* 1:847–852.
112. Heiniger, H.-J., A.A. Kandutsch, and H.W. Chen. 1976. Depletion of L-cell sterol depresses endocytosis. *Nature*. 263:515–517.

113. Alderson, J.C.E., and C. Green. 1975. Enrichment of lymphocytes with cholesterol and its effect on lymphocyte activation. *FEBS Letters*. 52:208–211.
114. Lohr, K.M., and R. Snyderman. 1982. Amphotericin B alters the affinity and functional activity of the oligopeptide chemotactic factor receptor on human polymorphonuclear leukocytes. *The Journal of Immunology*. 129:1594–1599.
115. Dahl, C., H.P. Biemann, and J. Dahl. 1987. A protein kinase antigenically related to pp60v-src possibly involved in yeast cell cycle control: positive in vivo regulation by sterol. *Proc. Natl. Acad. Sci. U.S.A.* 84:4012–4016.
116. Brown, D.A., and J.K. Rose. 1992. Sorting of GPI-anchored proteins to glycolipid-enriched membrane subdomains during transport to the apical cell surface. *Cell*. 68:533–544.
117. Fiedler, K., T. Kobayashi, T.V. Kurzchalia, and K. Simons. 1993. Glycosphingolipid-enriched, detergent-insoluble complexes in protein sorting in epithelial cells. *Biochemistry*. 32:6365–6373.
118. Harder, T., P. Scheiffele, P. Verkade, and K. Simons. 1998. Lipid Domain Structure of the Plasma Membrane Revealed by Patching of Membrane Components. *J Cell Biol*. 141:929–942.
119. Ge, M., A. Gidwani, H.A. Brown, D. Holowka, B. Baird, and J.H. Freed. 2003. Ordered and Disordered Phases Coexist in Plasma Membrane Vesicles of RBL-2H3 Mast Cells. An ESR Study. *Bioophysical Journal*. 85:1278–1288.
120. Gidwani, A., D. Holowka, and B. Baird. 2001. Fluorescence Anisotropy Measurements of Lipid Order in Plasma Membranes and Lipid Rafts from RBL-2H3 Mast Cells. *Biochemistry*. 40:12422–12429.
121. Kenworthy, A.K., and M. Edidin. 1998. Distribution of a Glycosylphosphatidylinositol-anchored Protein at the Apical Surface of MDCK Cells Examined at a Resolution of <math>< 100 \text{ \AA}</math> Using Imaging Fluorescence Resonance Energy Transfer. *J Cell Biol*. 142:69–84.
122. Sheets, E.D., D. Holowka, and B. Baird. 1999. Critical Role for Cholesterol in Lyn-mediated Tyrosine Phosphorylation of Fc ϵ RI and Their Association with Detergent-resistant Membranes. *J Cell Biol*. 145:877–887.
123. Pierce, S.K. 2002. Lipid rafts and B-cell activation. *Nature Reviews Immunology*. 2:nri726.
124. Plowman, S.J., C. Muncke, R.G. Parton, and J.F. Hancock. 2005. H-ras, K-ras, and inner plasma membrane raft proteins operate in nanoclusters with differential dependence on the actin cytoskeleton. *Proc Natl Acad Sci U S A*. 102:15500–5.
125. Brown, D.A., and E. London. 1998. Functions of lipid rafts in biological membranes. *Annual review of cell and developmental biology*. 14:111–136.
126. Munro, S. 2003. Lipid Rafts: Elusive or Illusive? *Cell*. 115:377–388.

127. Veatch, S.L., O. Soubias, S.L. Keller, and K. Gawrisch. 2007. Critical fluctuations in domain-forming lipid mixtures. *PNAS*. 104:17650–17655.
128. Zhao, J., J. Wu, F.A. Heberle, T.T. Mills, P. Klawitter, G. Huang, G. Costanza, and G.W. Feigenson. 2007. Phase studies of model biomembranes: Complex behavior of DSPC/DOPC/Cholesterol. *Biochimica et Biophysica Acta (BBA) - Biomembranes*. 1768:2764–2776.
129. Petruzielo, R.S., F.A. Heberle, P. Drazba, J. Katsaras, and G.W. Feigenson. 2013. Phase behavior and domain size in sphingomyelin-containing lipid bilayers. *Biochimica et Biophysica Acta (BBA) - Biomembranes*. 1828:1302–1313.
130. Connell, S.D., G. Heath, P.D. Olmsted, and A. Kisil. 2013. Critical point fluctuations in supported lipid membranes. *Faraday Discuss*. 161:91–111.
131. Mills, T.T., S. Tristram-Nagle, F.A. Heberle, N.F. Morales, J. Zhao, J. Wu, G.E.S. Toombes, J.F. Nagle, and G.W. Feigenson. 2008. Liquid-Liquid Domains in Bilayers Detected by Wide Angle X-Ray Scattering. *Biophysical Journal*. 95:682–690.
132. Konyakhina, T.M., and G.W. Feigenson. 2016. Phase diagram of a polyunsaturated lipid mixture: Brain sphingomyelin/1-stearoyl-2-docosahexaenoyl-sn-glycero-3-phosphocholine/cholesterol. *Biochimica et Biophysica Acta (BBA) - Biomembranes*. 1858:153–161.
133. de Almeida, R.F.M., J. Borst, A. Fedorov, M. Prieto, and A.J.W.G. Visser. 2007. Complexity of Lipid Domains and Rafts in Giant Unilamellar Vesicles Revealed by Combining Imaging and Microscopic and Macroscopic Time-Resolved Fluorescence. *Biophysical Journal*. 93:539–553.
134. Sibold, J., V.E. Tewaag, T. Vagedes, I. Mey, and C. Steinem. 2020. Phase separation in pore-spanning membranes induced by differences in surface adhesion. *Phys. Chem. Chem. Phys.* 22:9308–9315.
135. Kahya, N., D. Scherfeld, K. Bacia, B. Poolman, and P. Schwille. 2003. Probing lipid mobility of raft-exhibiting model membranes by fluorescence correlation spectroscopy. *J. Biol. Chem.* 278:28109–28115.
136. Aufderhorst-Roberts, A., U. Chandra, and S.D. Connell. 2017. Three-Phase Coexistence in Lipid Membranes. *Biophysical Journal*. 112:313–324.
137. Goñi, F.M., A. Alonso, L.A. Bagatolli, R.E. Brown, D. Marsh, M. Prieto, and J.L. Thewalt. 2008. Phase diagrams of lipid mixtures relevant to the study of membrane rafts. *Biochimica et Biophysica Acta (BBA) - Molecular and Cell Biology of Lipids*. 1781:665–684.
138. Heberle, F.A., and G.W. Feigenson. 2011. Phase Separation in Lipid Membranes. *Cold Spring Harb Perspect Biol*. 3:a004630.

139. Ladbrooke, B.D., and D. Chapman. 1969. Thermal analysis of lipids, proteins and biological membranes a review and summary of some recent studies. *Chemistry and Physics of Lipids*. 3:304–356.
140. Pathak, P., and E. London. 2011. Measurement of Lipid Nanodomain (Raft) Formation and Size in Sphingomyelin/POPC/Cholesterol Vesicles Shows TX-100 and Transmembrane Helices Increase Domain Size by Coalescing Preexisting Nanodomains But Do Not Induce Domain Formation. *Biophysical Journal*. 101:2417–2425.
141. Pathak, P., and E. London. 2015. The Effect of Membrane Lipid Composition on the Formation of Lipid Ultrananodomains. *Biophysical Journal*. 109:1630–1638.
142. Konyakhina, T.M., J. Wu, J.D. Mastroianni, F.A. Heberle, and G.W. Feigenson. 2013. Phase Diagram of a 4-Component Lipid Mixture: DSPC/DOPC/POPC/chol. *Biochimica et Biophysica Acta (BBA)-Biomembranes*.
143. Konyakhina, T.M., S.L. Goh, J. Amazon, F.A. Heberle, J. Wu, and G.W. Feigenson. 2011. Control of a Nanoscopic-to-Macroscopic Transition: Modulated Phases in Four-Component DSPC/DOPC/POPC/Chol Giant Unilamellar Vesicles. *Biophysical Journal*. 101:L8–L10.
144. Collado, M.I., F.M. Goñi, A. Alonso, and D. Marsh. 2005. Domain formation in sphingomyelin/cholesterol mixed membranes studied by spin-label electron spin resonance spectroscopy. *Biochemistry*. 44:4911–4918.
145. de Almeida, R.F.M., A. Fedorov, and M. Prieto. 2003. Sphingomyelin/Phosphatidylcholine/Cholesterol Phase Diagram: Boundaries and Composition of Lipid Rafts. *Biophysical Journal*. 85:2406–2416.
146. Doktorova, M., F.A. Heberle, B. Eicher, R.F. Standaert, J. Katsaras, E. London, G. Pabst, and D. Marquardt. 2018. Preparation of asymmetric phospholipid vesicles for use as cell membrane models. *Nature Protocols*. 13:2086–2101.
147. Lin, Q., and E. London. 2015. Ordered Raft Domains Induced by Outer Leaflet Sphingomyelin in Cholesterol-Rich Asymmetric Vesicles. *Biophysical Journal*. 108:2212–2222.
148. Honerkamp-Smith, A.R., P. Cicuta, M.D. Collins, S.L. Veatch, M. den Nijs, M. Schick, and S.L. Keller. 2008. Line Tensions, Correlation Lengths, and Critical Exponents in Lipid Membranes Near Critical Points. *Biophysical Journal*. 95:236–246.
149. Leibler, S., and D. Andelman. 1987. Ordered and curved meso-structures in membranes and amphiphilic films. *J. Phys. France*. 48:2013–2018.
150. Palmieri, B., T. Yamamoto, R.C. Brewster, and S.A. Safran. 2014. Line active molecules promote inhomogeneous structures in membranes: Theory, simulations and experiments. *Advances in Colloid and Interface Science*. 208:58–65.

151. Brodbek, L., and F. Schmid. 2016. Interplay of curvature-induced micro- and nanodomain structures in multicomponent lipid bilayers. *Int J Adv Eng Sci Appl Math*. 8:111–120.
152. Cornell, C.E., A.D. Skinkle, S. He, I. Levental, K.R. Levental, and S.L. Keller. 2018. Tuning Length Scales of Small Domains in Cell-Derived Membranes and Synthetic Model Membranes. *Biophysical Journal*. 115:690–701.
153. Fridriksson, E.K., P.A. Shipkova, E.D. Sheets, D. Holowka, B. Baird, and F.W. McLafferty. 1999. Quantitative Analysis of Phospholipids in Functionally Important Membrane Domains from RBL-2H3 Mast Cells Using Tandem High-Resolution Mass Spectrometry. *Biochemistry*. 38:8056–8063.
154. Lingwood, D., J. Ries, P. Schwille, and K. Simons. 2008. Plasma membranes are poised for activation of raft phase coalescence at physiological temperature. *PNAS*. 105:10005–10010.
155. Levental, I., M. Grzybek, and K. Simons. 2011. Raft domains of variable properties and compositions in plasma membrane vesicles. *PNAS*. 108:11411–11416.
156. Zhao, J., J. Wu, and S.L. Veatch. 2013. Adhesion Stabilizes Robust Lipid Heterogeneity in Supercritical Membranes at Physiological Temperature. *Biophysical Journal*. 104:825–834.
157. Gray, E.M., G. Díaz-Vázquez, and S.L. Veatch. 2015. Growth Conditions and Cell Cycle Phase Modulate Phase Transition Temperatures in RBL-2H3 Derived Plasma Membrane Vesicles. *PLOS ONE*. 10:e0137741.
158. Levental, K.R., J.H. Lorent, X. Lin, A.D. Skinkle, M.A. Surma, E.A. Stockenbojer, A.A. Gorfe, and I. Levental. 2016. Polyunsaturated Lipids Regulate Membrane Domain Stability by Tuning Membrane Order. *Biophysical Journal*. 110:1800–1810.
159. Levental, K.R., M.A. Surma, A.D. Skinkle, J.H. Lorent, Y. Zhou, C. Klose, J.T. Chang, J.F. Hancock, and I. Levental. 2017. ω -3 polyunsaturated fatty acids direct differentiation of the membrane phenotype in mesenchymal stem cells to potentiate osteogenesis. *Science Advances*. 3:eaa01193.
160. Cammarota, E., C. Soriani, R. Taub, F. Morgan, J. Sakai, S.L. Veatch, C.E. Bryant, and P. Cicuta. 2020. Criticality of plasma membrane lipids reflects activation state of macrophage cells. *Journal of The Royal Society Interface*. 17:20190803.
161. Symons, J.L., K.-J. Cho, J.T. Chang, G. Du, M.N. Waxham, J.F. Hancock, I. Levental, and K.R. Levental. 2020. Lipidomic atlas of mammalian cell membranes reveals hierarchical variation induced by culture conditions, subcellular membranes, and cell lineages. *Soft Matter*.
162. Kaiser, H.-J., D. Lingwood, I. Levental, J.L. Sampaio, L. Kalvodova, L. Rajendran, and K. Simons. 2009. Order of lipid phases in model and plasma membranes. *PNAS*. 106:16645–16650.

163. Lorent, J.H., B. Diaz-Rohrer, X. Lin, K. Spring, A.A. Gorfe, K.R. Levental, and I. Levental. 2017. Structural determinants and functional consequences of protein affinity for membrane rafts. *Nature Communications*. 8:1–10.
164. Menon, A.K., D. Holowka, W.W. Webb, and B. Baird. 1986. Cross-linking of receptor-bound IgE to aggregates larger than dimers leads to rapid immobilization. *J Cell Biol*. 102:541–550.
165. Machta, B.B., S. Papanikolaou, J.P. Sethna, and S.L. Veatch. 2011. Minimal Model of Plasma Membrane Heterogeneity Requires Coupling Cortical Actin to Criticality. *Biophysical Journal*. 100:1668–1677.
166. Keller, H., M. Lorzate, and P. Schwille. 2009. PI(4,5)P2 Degradation Promotes the Formation of Cytoskeleton-Free Model Membrane Systems. *ChemPhysChem*. 10:2805–2812.
167. McLaughlin, S., J. Wang, A. Gambhir, and D. Murray. 2002. PIP2 and Proteins: Interactions, Organization, and Information Flow. *Annu. Rev. Biophys. Biomol. Struct.* 31:151–175.
168. Skinkle, A.D., K.R. Levental, and I. Levental. 2020. Cell-Derived Plasma Membrane Vesicles Are Permeable to Hydrophilic Macromolecules. *Biophysical Journal*. 118:1292–1300.
169. Li, G., Q. Wang, S. Kakuda, and E. London. 2020. Nanodomains can persist at physiologic temperature in plasma membrane vesicles and be modulated by altering cell lipids. *J. Lipid Res.* jlr.RA119000565.
170. Toulmay, A., and W.A. Prinz. 2013. Direct imaging reveals stable, micrometer-scale lipid domains that segregate proteins in live cells. *J Cell Biol*. 202:35–44.
171. Rayermann, S.P., G.E. Rayermann, C.E. Cornell, A.J. Merz, and S.L. Keller. 2017. Hallmarks of Reversible Separation of Living, Unperturbed Cell Membranes into Two Liquid Phases. *Biophysical Journal*. 113:2425–2432.
172. Lee, I.-H., S. Saha, A. Polley, H. Huang, S. Mayor, M. Rao, and J.T. Groves. 2015. Live Cell Plasma Membranes Do Not Exhibit a Miscibility Phase Transition over a Wide Range of Temperatures. *J. Phys. Chem. B*. 119:4450–4459.
173. Sevcsik, E., M. Brameshuber, M. Fölser, J. Weghuber, A. Honigmann, and G.J. Schütz. 2015. GPI-anchored proteins do not reside in ordered domains in the live cell plasma membrane. *Nature Communications*. 6:1–10.
174. Levental, I., K.R. Levental, and F.A. Heberle. 2020. Lipid Rafts: Controversies Resolved, Mysteries Remain. *Trends in Cell Biology*. 30:341–353.

175. Kusumi, A., T.K. Fujiwara, T.A. Tsunoyama, R.S. Kasai, A.-A. Liu, K.M. Hirose, M. Kinoshita, N. Matsumori, N. Komura, H. Ando, and K.G.N. Suzuki. 2020. Defining raft domains in the plasma membrane. *Traffic*. 21:106–137.
176. Castello-Serrano, I., J.H. Lorent, R. Ippolito, K.R. Levental, and I. Levental. 2020. Myelin-Associated MAL and PLP Are Unusual among Multipass Transmembrane Proteins in Preferring Ordered Membrane Domains. *J. Phys. Chem. B*.
177. Marinko, J.T., A.K. Kenworthy, and C.R. Sanders. 2020. Peripheral myelin protein 22 preferentially partitions into ordered phase membrane domains. *PNAS*.
178. Yang, S.-T., A.J.B. Kreutzberger, V. Kiessling, B.K. Ganser-Pornillos, J.M. White, and L.K. Tamm. 2017. HIV virions sense plasma membrane heterogeneity for cell entry. *Science Advances*. 3:e1700338.
179. Anderson, R.G.W., and K. Jacobson. 2002. A Role for Lipid Shells in Targeting Proteins to Caveolae, Rafts, and Other Lipid Domains. *Science*. 296:1821–1825.
180. Corradi, V., E. Mendez-Villuendas, H.I. Ingólfsson, R.-X. Gu, I. Siuda, M.N. Melo, A. Moussatova, L.J. DeGagné, B.I. Sejdiu, G. Singh, T.A. Wassenaar, K. Delgado Magnero, S.J. Marrink, and D.P. Tieleman. 2018. Lipid–Protein Interactions Are Unique Fingerprints for Membrane Proteins. *ACS Cent. Sci.* 4:709–717.
181. Gray, E., J. Karlslake, B.B. Machta, and S.L. Veatch. 2013. Liquid General Anesthetics Lower Critical Temperatures in Plasma Membrane Vesicles. *Biophysical Journal*. 105:2751–2759.
182. Machta, B.B., E. Gray, M. Nouri, N.L.C. McCarthy, E.M. Gray, A.L. Miller, N.J. Brooks, and S.L. Veatch. 2016. Conditions that Stabilize Membrane Domains Also Antagonize n-Alcohol Anesthesia. *Biophysical Journal*. 111:537–545.
183. Grinstein, G., and S. Ma. 1982. Roughening and Lower Critical Dimension in the Random-Field Ising Model. *Phys. Rev. Lett.* 49:685–688.
184. Yethiraj, A., and J.C. Weisshaar. 2007. Why Are Lipid Rafts Not Observed In Vivo? *Biophysical Journal*. 93:3113–3119.
185. Honigsmann, A., S. Sadeghi, J. Keller, S.W. Hell, C. Eggeling, and R. Vink. 2014. A lipid bound actin meshwork organizes liquid phase separation in model membranes. *eLife*. 3.
186. Kwik, J., S. Boyle, D. Fooksman, L. Margolis, M.P. Sheetz, and M. Edidin. 2003. Membrane cholesterol, lateral mobility, and the phosphatidylinositol 4,5-bisphosphate-dependent organization of cell actin. *PNAS*. 100:13964–13969.
187. Liu, A.P., and D.A. Fletcher. 2006. Actin Polymerization Serves as a Membrane Domain Switch in Model Lipid Bilayers. *Biophysical Journal*. 91:4064–4070.

188. Shelby, S.A., S.L. Veatch, D.A. Holowka, and B.A. Baird. 2016. Functional nanoscale coupling of Lyn kinase with IgE-FcεRI is restricted by the actin cytoskeleton in early antigen-stimulated signaling. *MBoC*. 27:3645–3658.
189. Gudheti, M.V., N.M. Curthoys, T.J. Gould, D. Kim, M.S. Gunewardene, K.A. Gabor, J.A. Gosse, C.H. Kim, J. Zimmerberg, and S.T. Hess. 2013. Actin Mediates the Nanoscale Membrane Organization of the Clustered Membrane Protein Influenza Hemagglutinin. *Biophysical Journal*. 104:2182–2192.
190. Chichili, G.R., and W. Rodgers. 2009. Cytoskeleton–membrane interactions in membrane raft structure. *Cell. Mol. Life Sci*. 66:2319–2328.
191. Kinoshita, M., K.G.N. Suzuki, N. Matsumori, M. Takada, H. Ano, K. Morigaki, M. Abe, A. Makino, T. Kobayashi, K.M. Hirosawa, T.K. Fujiwara, A. Kusumi, and M. Murata. 2017. Raft-based sphingomyelin interactions revealed by new fluorescent sphingomyelin analogs. *J Cell Biol*. 216:1183–1204.
192. Allender, D.W., H. Giang, and M. Schick. 2020. Model Plasma Membrane Exhibits a Microemulsion in Both Leaves Providing a Foundation for “Rafts.” *Biophysical Journal*. 118:1019–1031.
193. Foret, L. 2005. A simple mechanism of raft formation in two-component fluid membranes. *EPL*. 71:508.
194. Gowrishankar, K., S. Ghosh, S. Saha, R. C., S. Mayor, and M. Rao. 2012. Active Remodeling of Cortical Actin Regulates Spatiotemporal Organization of Cell Surface Molecules. *Cell*. 149:1353–1367.
195. Schmid, F. 2017. Physical mechanisms of micro- and nanodomain formation in multicomponent lipid membranes. *Biochimica et Biophysica Acta (BBA) - Biomembranes*. 1859:509–528.
196. Goldenfeld, N. 1992. Lectures on phase transitions and the renormalization group. Reading, Mass: Addison-Wesley, Advanced Book Program.
197. Cardy, J. 1996. Scaling and Renormalization in Statistical Physics. Cambridge: Cambridge University Press.
198. Kardar, M. 2007. Statistical physics of fields. Cambridge ; New York: Cambridge University Press.
199. Honerkamp-Smith, A.R., S.L. Veatch, and S.L. Keller. 2009. An introduction to critical points for biophysicists; observations of compositional heterogeneity in lipid membranes. *Biochimica et Biophysica Acta (BBA) - Biomembranes*. 1788:53–63.
200. Cagniard de la Tour, C. 1822. Exposé de quelques résultats obtenu par l’action combinée de la chaleur et de la compression sur certains liquides, tels que l’eau, l’alcool, l’éther sulfurique et l’essence de pétrole rectifiée. *Annales de Chimie et de Physique*. 21:127–132.

201. van der Waals, J.D., and J.S. Rowlinson. 2004. On the continuity of the gaseous and liquid states. Mineola, N.Y: Dover Publications.
202. Ornstein, L., and F. Zernike. 1918. Die linearen dimensionen der dichteschwankungen. *Phys. Zeit.* 19:134–7.
203. Guggenheim, E.A. 1945. The Principle of Corresponding States. *J. Chem. Phys.* 13:253–261.
204. Fisher, M.E. 1967. The theory of equilibrium critical phenomena. *Rep. Prog. Phys.* 30:615.
205. Onsager, L. 1944. Crystal Statistics. I. A Two-Dimensional Model with an Order-Disorder Transition. *Physical Review.* 65:117–149.
206. Lee, T.D., and C.N. Yang. 1952. Statistical Theory of Equations of State and Phase Transitions. II. Lattice Gas and Ising Model. *Physical Review.* 87:410–419.
207. Widom, B. 1965. Equation of State in the Neighborhood of the Critical Point. *J. Chem. Phys.* 43:3898–3905.
208. Kadanoff, L.P. 1966. Scaling laws for ising models near T_c . *Physics Physique Fizika.* 2:263–272.
209. Wilson, K.G. 1971. Renormalization Group and Critical Phenomena. I. Renormalization Group and the Kadanoff Scaling Picture. *Phys. Rev. B.* 4:3174–3183.
210. Mora, T., and W. Bialek. 2011. Are Biological Systems Poised at Criticality? *Journal of Statistical Physics.* 144:268–302.
211. Honerkamp-Smith, A.R., B.B. Machta, and S.L. Keller. 2012. Experimental Observations of Dynamic Critical Phenomena in a Lipid Membrane. *Phys. Rev. Lett.* 108:265702.
212. Frazier, M.L., J.R. Wright, A. Pokorny, and P.F.F. Almeida. 2007. Investigation of Domain Formation in Sphingomyelin/Cholesterol/POPC Mixtures by Fluorescence Resonance Energy Transfer and Monte Carlo Simulations. *Biophysical Journal.* 92:2422–2433.
213. Manley, S., M.R. Horton, S. Leczynski, and A.P. Gast. 2008. Sorting of Streptavidin Protein Coats on Phase-Separating Model Membranes. *Biophysical Journal.* 95:2301–2307.
214. Kimchi, O., S.L. Veatch, and B.B. Machta. 2018. Ion channels can be allosterically regulated by membrane domains near a de-mixing critical point. *The Journal of General Physiology.* 150:1769–1777.
215. Lingwood, D., and K. Simons. 2010. Lipid rafts as a membrane-organizing principle. *Science.* 327:46–50.
216. Levental, I., and S.L. Veatch. 2016. The Continuing Mystery of Lipid Rafts. *Journal of Molecular Biology.* 428:4749–4764.

217. Sezgin, E., I. Levental, S. Mayor, and C. Eggeling. 2017. The mystery of membrane organization: composition, regulation and roles of lipid rafts. *Nature Reviews Molecular Cell Biology*. 18:361–374.
218. Casimir, H.B.G. 1948. On the Attraction Between Two Perfectly Conducting Plates. *Indag.Math.* 10:261–263.
219. Machta, B.B., S.L. Veatch, and J.P. Sethna. 2012. Critical Casimir Forces in Cellular Membranes. *Phys. Rev. Lett.* 109:138101.
220. Reynwar, B.J., and M. Deserno. 2008. Membrane composition-mediated protein-protein interactions. *Biointerphases*. 3:FA117–FA124.
221. Saheki, Y., and P. De Camilli. 2017. Endoplasmic Reticulum–Plasma Membrane Contact Sites. *Annual Review of Biochemistry*. 86:659–684.
222. Zeng, M., X. Chen, D. Guan, J. Xu, H. Wu, P. Tong, and M. Zhang. 2018. Reconstituted Postsynaptic Density as a Molecular Platform for Understanding Synapse Formation and Plasticity. *Cell*. 174:1172-1187.e16.
223. Ayuyan, A.G., and F.S. Cohen. 2018. The Chemical Potential of Plasma Membrane Cholesterol: Implications for Cell Biology. *Biophysical Journal*. 114:904–918.
224. Fantini, J., R.M. Epanand, and F.J. Barrantes. 2019. Cholesterol-Recognition Motifs in Membrane Proteins. *Adv. Exp. Med. Biol.* 1135:3–25.
225. De Craene, J.-O., D.L. Bertazzi, S. Bär, and S. Friant. 2017. Phosphoinositides, Major Actors in Membrane Trafficking and Lipid Signaling Pathways. *Int J Mol Sci*. 18.
226. Shimizu, T. 2009. Lipid mediators in health and disease: enzymes and receptors as therapeutic targets for the regulation of immunity and inflammation. *Annu. Rev. Pharmacol. Toxicol.* 49:123–150.
227. Eggeling, C., C. Ringemann, R. Medda, G. Schwarzmann, K. Sandhoff, S. Polyakova, V.N. Belov, B. Hein, C. von Middendorff, A. Schönle, and S.W. Hell. 2009. Direct observation of the nanoscale dynamics of membrane lipids in a living cell. *Nature*. 457:1159–1162.
228. Rao, M., and S. Mayor. 2014. Active organization of membrane constituents in living cells. *Current Opinion in Cell Biology*. 29:126–132.
229. Sorre, B., A. Callan-Jones, J.-B. Manneville, P. Nassoy, J.-F. Joanny, J. Prost, B. Goud, and P. Bassereau. 2009. Curvature-driven lipid sorting needs proximity to a demixing point and is aided by proteins. *PNAS*. 106:5622–5626.
230. Mitra, E.D., S.C. Whitehead, D. Holowka, B. Baird, and J.P. Sethna. 2018. Computation of a Theoretical Membrane Phase Diagram and the Role of Phase in Lipid-Raft-Mediated Protein Organization. *J. Phys. Chem. B*. 122:3500–3513.

231. Shaw, T.R., K.C. Wisser, T.A. Schaffner, A.D. Gaffney, B.B. Machta, and S.L. Veatch. 2023. Chemical potential measurements constrain models of cholesterol-phosphatidylcholine interactions. *Biophysical Journal*. 122:1105–1117.
232. Heberle, F.A., M. Doktorova, S.L. Goh, R.F. Standaert, J. Katsaras, and G.W. Feigenson. 2013. Hybrid and Nonhybrid Lipids Exert Common Effects on Membrane Raft Size and Morphology. *J. Am. Chem. Soc.* 135:14932–14935.
233. Huster, D., K. Arnold, and K. Gawrisch. 1998. Influence of Docosahexaenoic Acid and Cholesterol on Lateral Lipid Organization in Phospholipid Mixtures. *Biochemistry*. 37:17299–17308.
234. Veatch, S.L., I.V. Polozov, K. Gawrisch, and S.L. Keller. 2004. Liquid Domains in Vesicles Investigated by NMR and Fluorescence Microscopy. *Biophysical Journal*. 86:2910–2922.
235. Leventis, R., and J.R. Silvius. 2001. Use of Cyclodextrins to Monitor Transbilayer Movement and Differential Lipid Affinities of Cholesterol. *Biophysical Journal*. 81:2257–2267.
236. Tsamaloukas, A., H. Szadkowska, and H. Heerklotz. 2006. Thermodynamic Comparison of the Interactions of Cholesterol with Unsaturated Phospholipid and Sphingomyelins. *Biophysical Journal*. 90:4479–4487.
237. Tsamaloukas, A., H. Szadkowska, and H. Heerklotz. 2006. Nonideal mixing in multicomponent lipid/detergent systems. *J. Phys.: Condens. Matter*. 18:S1125–S1138.
238. Sokolov, A., and A. Radhakrishnan. 2010. Accessibility of Cholesterol in Endoplasmic Reticulum Membranes and Activation of SREBP-2 Switch Abruptly at a Common Cholesterol Threshold. *Journal of Biological Chemistry*. 285:29480–29490.
239. Lange, Y., S.M.A. Tabei, J. Ye, and T.L. Steck. 2013. Stability and Stoichiometry of Bilayer Phospholipid–Cholesterol Complexes: Relationship to Cellular Sterol Distribution and Homeostasis. *Biochemistry*. 52:6950–6959.
240. Nyholm, T.K.M., S. Jaikishan, O. Engberg, V. Hautala, and J.P. Slotte. 2019. The Affinity of Sterols for Different Phospholipid Classes and Its Impact on Lateral Segregation. *Biophysical Journal*. 116:296–307.
241. Scott, H.L., A. Skinkle, E.G. Kelley, M.N. Waxham, I. Levental, and F.A. Heberle. 2019. On the Mechanism of Bilayer Separation by Extrusion, or Why Your LUVs Are Not Really Unilamellar. *Biophysical Journal*. 117:1381–1386.
242. Huang, Z., and E. London. 2013. Effect of Cyclodextrin and Membrane Lipid Structure upon Cyclodextrin–Lipid Interaction. *Langmuir*. 29:14631–14638.
243. Anderson, T.G., A. Tan, P. Ganz, and J. Seelig. 2004. Calorimetric Measurement of Phospholipid Interaction with Methyl- β -Cyclodextrin. *Biochemistry*. 43:2251–2261.

244. Fenyvesi, É., J. Szemán, K. Csabai, M. Malanga, and L. Szente. 2014. Methyl-Beta-Cyclodextrins: The Role of Number and Types of Substituents in Solubilizing Power. *Journal of Pharmaceutical Sciences*. 103:1443–1452.
245. Jandacek, R.J., M.R. Webb, and F.H. Mattson. 1977. Effect of an aqueous phase on the solubility of cholesterol in an oil phase. *Journal of Lipid Research*. 18:203–210.
246. Wolff, J., C.M. Marques, and F. Thalmann. 2011. Thermodynamic approach to phase coexistence in ternary phospholipid-cholesterol mixtures. *Phys Rev Lett*. 106:128104.
247. Huang, J., J.T. Buboltz, and G.W. Feigenson. 1999. Maximum solubility of cholesterol in phosphatidylcholine and phosphatidylethanolamine bilayers. *Biochimica et Biophysica Acta (BBA) - Biomembranes*. 1417:89–100.
248. Ali, M.R., K.H. Cheng, and J. Huang. 2007. Assess the nature of cholesterol–lipid interactions through the chemical potential of cholesterol in phosphatidylcholine bilayers. *Proceedings of the National Academy of Sciences*. 104:5372–5377.
249. Stevens, M.M., A.R. Honerkamp-Smith, and S.L. Keller. 2010. Solubility Limits of Cholesterol, Lanosterol, Ergosterol, Stigmasterol, and β -Sitosterol in Electroformed Lipid Vesicles. *Soft Matter*. 6:5882–5890.
250. Epand, R.M., D. Bach, and E. Wachtel. 2016. In vitro determination of the solubility limit of cholesterol in phospholipid bilayers. *Chemistry and Physics of Lipids*. 199:3–10.
251. Heerklotz, H., and A. Tsamaloukas. 2006. Gradual Change or Phase Transition: Characterizing Fluid Lipid-Cholesterol Membranes on the Basis of Thermal Volume Changes. *Biophysical Journal*. 91:600–607.
252. Reyes Mateo, C., A. Ulises Acuña, and J.C. Brochon. 1995. Liquid-crystalline phases of cholesterol/lipid bilayers as revealed by the fluorescence of trans-parinaric acid. *Biophysical Journal*. 68:978–987.
253. Widom, B. 2002. *Statistical mechanics: a concise introduction for chemists*. Cambridge ; New York: Cambridge University Press.
254. Veatch, S.L., S.S.W. Leung, R.E.W. Hancock, and J.L. Thewalt. 2007. Fluorescent Probes Alter Miscibility Phase Boundaries in Ternary Vesicles. *J. Phys. Chem. B*. 111:502–504.
255. Radhakrishnan, A., and H.M. McConnell. 1999. Condensed Complexes of Cholesterol and Phospholipids. *Biophysical Journal*. 77:1507–1517.
256. Radhakrishnan, A., and H.M. McConnell. 2000. Chemical activity of cholesterol in membranes. *Biochemistry*. 39:8119–8124.
257. Radhakrishnan, A., and H. McConnell. 2005. Condensed complexes in vesicles containing cholesterol and phospholipids. *Proceedings of the National Academy of Sciences*. 102:12662–12666.

258. Radhakrishnan, A. 2010. Phase Separations in Binary and Ternary Cholesterol-Phospholipid Mixtures. *Biophysical Journal*. 98:L41–L43.
259. Bell, G.M. 1953. Statistical thermodynamics of regular ternary solutions. *Transactions of the Faraday Society*. 49:122.
260. Lönnfors, M., J.P.F. Doux, J.A. Killian, T.K.M. Nyholm, and J.P. Slotte. 2011. Sterols Have Higher Affinity for Sphingomyelin than for Phosphatidylcholine Bilayers even at Equal Acyl-Chain Order. *Biophysical Journal*. 100:2633–2641.
261. Lorent, J.H., K.R. Levental, L. Ganesan, G. Rivera-Longsworth, E. Sezgin, M. Doktorova, E. Lyman, and I. Levental. 2020. Plasma membranes are asymmetric in lipid unsaturation, packing and protein shape. *Nat Chem Biol*. 16:644–652.
262. Varma, M., and M. Deserno. 2022. Distribution of cholesterol in asymmetric membranes driven by composition and differential stress. *Biophysical Journal*.
263. Kinnebrew, M., G. Luchetti, R. Sircar, S. Frigui, L.V. Viti, T. Naito, F. Beckert, Y. Saheki, C. Siebold, A. Radhakrishnan, and R. Rohatgi. 2021. Patched 1 reduces the accessibility of cholesterol in the outer leaflet of membranes. *eLife*. 10:e70504.
264. Haynes, M.P., M.C. Phillips, and G.H. Rothblat. 2000. Efflux of Cholesterol from Different Cellular Pools. *Biochemistry*. 39:4508–4517.
265. Litz, J.P., N. Thakkar, T. Portet, and S.L. Keller. 2016. Depletion with Cyclodextrin Reveals Two Populations of Cholesterol in Model Lipid Membranes. *Biophysical Journal*. 110:635–645.
266. Ahn, K., and N.S. Sampson. 2004. Cholesterol Oxidase Senses Subtle Changes in Lipid Bilayer Structure. *Biochemistry*. 43:827–836.
267. Johnson, K.A., and A. Radhakrishnan. 2021. Chapter Eighteen - The use of anthrolysin O and ostreolysin A to study cholesterol in cell membranes. In: Heuck AP, editor. *Methods in Enzymology*. Academic Press. pp. 543–566.
268. Johnson, B.B., P.C. Moe, D. Wang, K. Rossi, B.L. Trigatti, and A.P. Heuck. 2012. Modifications in Perfringolysin O Domain 4 Alter the Cholesterol Concentration Threshold Required for Binding. *Biochemistry*. 51:3373–3382.
269. Gay, A., D. Rye, and A. Radhakrishnan. 2015. Switch-like responses of two cholesterol sensors do not require protein oligomerization in membranes. *Biophys J*. 108:1459–1469.
270. Chakrabarti, R.S., S.A. Ingham, J. Kozlitina, A. Gay, J.C. Cohen, A. Radhakrishnan, and H.H. Hobbs. 2017. Variability of cholesterol accessibility in human red blood cells measured using a bacterial cholesterol-binding toxin. *Elife*. 6.
271. Fazekas, F.J., T.R. Shaw, S. Kim, R.A. Bogucki, and S.L. Veatch. 2021. A mean shift algorithm for drift correction in localization microscopy. *Biophysical Reports*. 1:100008.

272. Kao, H.P., and A.S. Verkman. 1994. Tracking of single fluorescent particles in three dimensions: use of cylindrical optics to encode particle position. *Biophysical Journal*. 67:1291–1300.
273. Holtzer, L., T. Meckel, and T. Schmidt. 2007. Nanometric three-dimensional tracking of individual quantum dots in cells. *Appl. Phys. Lett.* 90:053902.
274. Juette, M.F., T.J. Gould, M.D. Lessard, M.J. Mlodzianoski, B.S. Nagpure, B.T. Bennett, S.T. Hess, and J. Bewersdorf. 2008. Three-dimensional sub-100 nm resolution fluorescence microscopy of thick samples. *Nat Methods*. 5:527–529.
275. Carter, A.R., G.M. King, T.A. Ulrich, W. Halsey, D. Alchenberger, and T.T. Perkins. 2007. Stabilization of an optical microscope to 0.1 nm in three dimensions. *Appl. Opt.* 46:421.
276. Grover, G., W. Mohrman, and R. Piestun. 2015. Real-time adaptive drift correction for super-resolution localization microscopy. *Opt. Express*. 23:23887.
277. Tafteh, R., L. Abraham, D. Seo, H.Y. Lu, M.R. Gold, and K.C. Chou. 2016. Real-time 3D stabilization of a super-resolution microscope using an electrically tunable lens. *Opt. Express*. 24:22959.
278. Schmidt, P.D., B.H. Reichert, J.G. Lajoie, and S. Sivasankar. 2018. Method for high frequency tracking and sub-nm sample stabilization in single molecule fluorescence microscopy. *Sci Rep*. 8:13912.
279. Coelho, S., J. Baek, J. Walsh, J.J. Gooding, and K. Gaus. 2021. 3D active stabilization for single-molecule imaging. *Nat Protoc*. 16:497–515.
280. Huang, B., W. Wang, M. Bates, and X. Zhuang. 2008. Three-Dimensional Super-Resolution Imaging by Stochastic Optical Reconstruction Microscopy. *Science*. 319:810–813.
281. Bon, P., N. Bourg, S. Lécart, S. Monneret, E. Fort, J. Wenger, and S. Lévêque-Fort. 2015. Three-dimensional nanometre localization of nanoparticles to enhance super-resolution microscopy. *Nat Commun*. 6:7764.
282. Ma, H., J. Xu, J. Jin, Y. Huang, and Y. Liu. 2017. A Simple Marker-Assisted 3D Nanometer Drift Correction Method for Superresolution Microscopy. *Biophysical Journal*. 112:2196–2208.
283. Han, R., L. Wang, F. Xu, Y. Zhang, M. Zhang, Z. Liu, F. Ren, and F. Zhang. 2015. Drift correction for single-molecule imaging by molecular constraint field, a distance minimum metric. *BMC Biophys*. 8:1.
284. Wester, M.J., D.J. Schodt, H. Mazloom-Farsibaf, M. Fazel, S. Pallikkuth, and K.A. Lidke. 2021. Robust, Fiducial-Free Drift Correction for Super-resolution Imaging. *bioRxiv*. 2021.03.26.437196.

285. Cnossen, J., T.J. Cui, C. Joo, and C. Smith. 2021. Drift correction in localization microscopy using entropy minimization. *bioRxiv*. 2021.03.30.437682.
286. Fukunaga, K., and L. Hostetler. 1975. The estimation of the gradient of a density function, with applications in pattern recognition. *IEEE Trans. Inform. Theory*. 21:32–40.
287. Yizong Cheng. 1995. Mean shift, mode seeking, and clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 17:790–799.
288. Comaniciu, D., and P. Meer. 2002. Mean shift: a robust approach toward feature space analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 24:603–619.
289. Baddeley, A., E. Rubak, and R. Turner. 2016. Spatial point patterns: methodology and applications with R. Boca Raton ; London ; New York: CRC Press, Taylor & Francis Group.
290. Hilgenberg, L.G.W., and M.A. Smith. 2007. Preparation of Dissociated Mouse Cortical Neuron Cultures. *J Vis Exp*.
291. Pappas, S.S., C.-C. Liang, S. Kim, C.O. Rivera, and W.T. Dauer. 2018. TorsinA dysfunction causes persistent neuronal nuclear pore defects. *Human Molecular Genetics*. 27:407–420.
292. Haughton, G., L.W. Arnold, G.A. Bishop, and T.J. Mercolino. 1986. The CH Series of Murine B Cell Lymphomas: Neoplastic Analogues of Ly-1+ Normal B Cells. *Immunological Reviews*. 93:35–52.
293. Li, Y., M. Mund, P. Hoess, J. Deschamps, U. Matti, B. Nijmeijer, V.J. Sabinina, J. Ellenberg, I. Schoen, and J. Ries. 2018. Real-time 3D single-molecule localization using experimental point spread functions. *Nature Methods*. 15:367–369.
294. Shaw, T.R., F.J. Fazekas, S. Kim, J.C. Flanagan-Natoli, E.R. Sumrall, and S.L. Veatch. 2022. Estimating the localization spread function of static single-molecule localization microscopy images. *Biophysical Journal*. 121:2906–2920.
295. Annibale, P., S. Vanni, M. Scarselli, U. Rothlisberger, and A. Radenovic. 2011. Quantitative Photo Activated Localization Microscopy: Unraveling the Effects of Photoblinking. *PLOS ONE*. 6:e22678.
296. Nieuwenhuizen, R.P.J., M. Bates, A. Szymborska, K.A. Lidke, B. Rieger, and S. Stallinga. 2015. Quantitative Localization Microscopy: Effects of Photophysics and Labeling Stoichiometry. *PLOS ONE*. 10:e0127989.
297. Patel, L., D. Williamson, D.M. Owen, and E.A.K. Cohen. 2021. Blinking statistics and molecular counting in direct stochastic reconstruction microscopy (dSTORM). *Bioinformatics*. 37:2730–2737.
298. Shang, M., Z. Huang, and Y. Wang. 2022. Influence of drift correction precision on super-resolution localization microscopy. *Appl. Opt., AO*. 61:3516–3522.

299. Ni, J., B. Cao, G. Niu, D. Chen, G. Liang, T. Xia, H. Li, C. Xu, J. Wang, W. Zhang, Y. Zhang, X. Yuan, and Y. Ni. 2022. Improved localization precision via restricting confined biomolecule stochastic motion in single-molecule localization microscopy. *Nanophotonics*. 11:53–65.
300. Patel, L., N. Gustafsson, Y. Lin, R. Ober, R. Henriques, and E. Cohen. 2019. A hidden Markov model approach to characterizing the photo-switching behavior of fluorophores. *The Annals of Applied Statistics*. 13:1397–1429.
301. Dana. 2022. Simulate Continuous-Time Markov Chains. .
302. Chan, J.R., S.J. Hyduk, and M.I. Cybulsky. 2000. $\alpha 4\beta 1$ Integrin/VCAM-1 Interaction Activates $\alpha L\beta 2$ Integrin-Mediated Adhesion to ICAM-1 in Human T Cells. *The Journal of Immunology*. 164:746–753.
303. Rodgers, W. 2002. Making membranes green: construction and characterization of GFP-fusion proteins targeted to discrete plasma membrane domains. *Biotechniques*. 32:1044–1046, 1048, 1050–1051.
304. Zhang, M., H. Chang, Y. Zhang, J. Yu, L. Wu, W. Ji, J. Chen, B. Liu, J. Lu, Y. Liu, J. Zhang, P. Xu, and T. Xu. 2012. Rational design of true monomeric and bright photoactivatable fluorescent proteins. *Nat Methods*. 9:727–729.
305. Ovesný, M., P. Křížek, J. Borkovec, Z. Švindrych, and G.M. Hagen. 2014. ThunderSTORM: a comprehensive ImageJ plug-in for PALM and STORM data analysis and super-resolution imaging. *Bioinformatics*. 30:2389–2390.
306. Ester, M., H.-P. Kriegel, J. Sander, and X. Xu. 1996. A density-based algorithm for discovering clusters in large spatial databases with noise. AAAI Press. pp. 226–231.
307. Shaw, T.R., F.J. Fazekas, and S.L. Veatch. 2022. VeatchLab / SMLM Spacetime Resolution. .
308. Dai, M., R. Jungmann, and P. Yin. 2016. Optical imaging of individual biomolecules in densely packed clusters. *Nat Nanotechnol*. 11:798–807.
309. Niekamp, S., J. Sung, W. Huynh, G. Bhabha, R.D. Vale, and N. Stuurman. 2019. Nanometer-accuracy distance measurements between fluorophores at the single-molecule level. *Proceedings of the National Academy of Sciences*. 116:4275–4284.
310. Andersen, I.T., U. Hahn, E.C. Arnspang, L.N. Nejsum, and E.B.V. Jensen. 2018. Double Cox cluster processes — with applications to photoactivated localization microscopy. *Spatial Statistics*. 27:58–73.
311. Diggle, P. 2014. Statistical analysis of spatial and spatio-temporal point patterns. Boca Raton: CRC Press.

312. Shaw, T., J. Møller, and R.P. Waagepetersen. 2021. Globally intensity-reweighted estimators for K- and pair correlation functions. *Australian & New Zealand Journal of Statistics*. 63:93–118.
313. Stoyan, D., W.S. Kendall, S.N. Chiu, and J. Mecke. 2013. *Stochastic Geometry and Its Applications*. John Wiley & Sons.
314. Møller, J., and R.P. Waagepetersen. 2007. Modern Statistics for Spatial Point Processes. *Scandinavian Journal of Statistics*. 0:070927154002001-???
315. Lawrence, T., A. Baddeley, R.K. Milne, and G. Nair. 2016. Point pattern analysis on a region of a sphere. *Stat*. 5:144–157.
316. Møller, J., and E. Rubak. 2016. Functional summary statistics for point processes on the sphere with an application to determinantal point processes. *Spatial Statistics*. 18:4–23.
317. van Lieshout, M.N.M. 2011. A J–function for inhomogeneous point processes. *Statistica Neerlandica*. 65:183–201.
318. Waagepetersen, R., and Y. Guan. 2009. Two-step estimation for inhomogeneous spatial point processes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 71:685–702.
319. Ohser, J., and D. Stoyan. 1981. On the Second-Order and Orientation Analysis of Planar Stationary Point Processes. *Biometrical Journal*. 23:523–533.
320. Ripley, B.D. 1988. *Statistical Inference for Spatial Processes*. Cambridge University Press.
321. Lang, G., and E. Marcon. 2013. Testing randomness of spatial point patterns with the Ripley statistic. *ESAIM: Probability and Statistics*. 17:767–788.
322. Diggle, P. 1985. A Kernel Method for Smoothing Point Process Data. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*. 34:138–147.
323. van Lieshout, M.-C.N.M. 2012. On Estimation of the Intensity Function of a Point Process. *Methodol Comput Appl Probab*. 14:567–578.
324. Liao, J.G., and A. Berg. 2019. Sharpening Jensen’s Inequality. *The American Statistician*. 73:278–281.
325. Coeurjolly, J.-F., J. Møller, and R. Waagepetersen. 2017. A Tutorial on Palm Distributions for Spatial Point Processes: A Tutorial on Palm Distributions for Spatial Point Processes. *International Statistical Review*. 85:404–420.
326. Møller, J., A.R. Syversveen, and R.P. Waagepetersen. 1998. Log Gaussian Cox Processes. *Scandinavian Journal of Statistics*. 25:451–482.

327. Lavancier, F., J. Møller, and E. Rubak. 2015. Determinantal point process models and statistical inference. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*. 77:853–877.
328. R Core Team. 2022. R: A Language and Environment for Statistical Computing. .
329. Baddeley, A., E. Rubak, and R. Turner. 2016. Spatial point patterns: methodology and applications with R. Boca Raton London New York: CRC Press.
330. Shaw, T., E. Rubak, A. Baddeley, and R. Turner. 2023. globalKinhom: Inhomogeneous K-And Pair Correlation Functions Using Global Estimators. .
331. Cronie, O., and M.N.M. Van Lieshout. 2018. A non-model-based approach to bandwidth selection for kernel estimators of spatial intensity functions. *Biometrika*. 105:455–462.
332. Arnspang, E.C., P. Sengupta, K.I. Mortensen, H.H. Jensen, U. Hahn, E.B.V. Jensen, J. Lippincott-Schwartz, and L.N. Nejsun. 2018. Regulation of Plasma Membrane Nanodomains of the Water Channel Aquaporin-3 Revealed by Fixed and Live Photoactivated Localization Microscopy. *Nano Lett*.
333. Cornell, C.E., N.L.C. McCarthy, K.R. Levental, I. Levental, N.J. Brooks, and S.L. Keller. 2017. n-Alcohol Length Governs Shift in Lo-Ld Mixing Temperatures in Synthetic and Cell-Derived Membranes. *Biophysical Journal*. 113:1200–1211.