



CARDIFF UNIVERSITY

SCHOOL OF PHYSICS AND ASTRONOMY

OPTICAL IMAGING OF LIPID BILAYERS
AND ITS APPLICATIONS TO NEUROLOGY

ALEXANDER NAHMAD-ROHEN

a thesis submitted in fulfilment of the requirements
for the degree of Doctor of Philosophy

PROJECT SUPERVISORS:
PROFESSOR WOLFGANG LANGBEIN
PROFESSOR PAOLA BORRI

2019

to ♡ Claudia,

*without whose infinite patience and unconditional
support none of this would have been possible;*

thank you for always believing in me

Acknowledgements

As always, my biggest thanks go to my wife, Claudia Torregrosa-Flores, who has been nothing short of wonderful. She's been open when I've wanted to tell her about my work, supportive when I've felt hopeless, and insightful when I've needed advice, which I often haven't realised until she's offered it. Claudia, I have everything to be grateful for; thank you.

Thanks, obviously, to my supervisors, Wolfgang and Paola, who guided me through this long and winding process and, especially in the case of Wolfgang, didn't let me get away with doing a poor job at anything.

I would also like to thank the people whose help proved invaluable to this work: Iestyn Pope, Francesco Masia, David Regan, Paulina Carriba-Domínguez, Ollie Williams, Bledyn Williams, Paul Kemp, Georgiy Zoriniants, Lukas Payne, Thomas Heimburg, Tian Wang and Rima Budvytyte.

Thanks to my parents for all their moral support and their help with clarifications and ideas.

While on the subject of moral support, I would like to thank the people who lent me their time and general advice and believed in me: my aunts Catalina and Hilde, Mags Lyons, Louise Winter, Richard Lewis, Fernanda Torregrosa-Flores, Ruth Nasrat, Claudia Carrillo and Gloria Flores-Vizcarra.

Finally, I would like to thank the Mexican National Council for Science and Technology, CONACYT, for funding my doctorate.

Summary

Despite numerous advances in the field, the exact mechanism behind neural communication remains elusive. This work is an attempt to answer the question of whether the propagation of an action potential along a neuron axon involves a localised phase transition of the neuron membrane. The work consists of the development of optical techniques for the study of supported lipid bilayers as models for the cell membrane. The final aim of the work was to perform simultaneous optical and electrophysiological measurements on the membranes of living neurons in order to answer the aforementioned question.

We first present the scientific theory required to understand this work, as well as a brief summary of cell membrane biology. We then describe in detail the Hodgkin-Huxley and Heimburg-Jackson models of action potential propagation and examine how well each model is able to explain the vast body of experimental findings of mechanical, optical, thermodynamic and chemical phenomena concomitant with the action potential. From this examination we determine the properties required of a technique capable of observing action potentials in a way which can determine whether or not they involve a phase transition of the cell membrane. Two such techniques are interferometric reflectometry and stimulated Raman scattering. We then describe the procedures which were used to prepare the samples we performed experiments on, as well as the optical and electrophysiological experimental setups and methodology we used in our experiments. Afterwards, we explain the theoretical work we have done for interferometric reflectometry, followed by our attempts at reducing experimental noise and a discussion of our experimental results with this technique and of electrophysiology experiments. We finally describe our data analysis procedure for stimulated Raman scattering microscopy and discuss our results with this technique.

Introduction

The question of how neurons communicate with each other and with the rest of the body has preoccupied scientists for centuries. From Galvani's first investigations into bioelectricity in the 1700s to experiments still being carried out today, scientists have attempted to understand the mechanisms involved and have proposed multiple models to explain their observations, yet no model has been able to provide a complete description of neural communication consistent with the results of all the relevant experiments.

The best-known of these models was proposed by A L Hodgkin and A F Huxley in 1952. Inspired by evidence that neural activity involved travelling electric potentials — called “action potentials”— across the cell membrane, this model viewed the membrane and its surrounding medium as an RC circuit with variable resistance. For decades it was hailed as the answer biophysicists were looking for, and even today it is widely accepted, yet it fails to correctly account for many of the non-electrical phenomena observed during action potential propagation.

In the last decade, a new model was proposed by T R Heimburg and A D Jackson. This model is thermodynamic rather than electrical, suggesting a travelling phase transition in the cell membrane is responsible for the wide variety of phenomena caused by neural activity, one of which is the action potential. While this model appears to be more complete than the Hodgkin-Huxley model (in the sense that it can explain some of the findings the previous model cannot), it has received little attention from the scientific community, so whether a phase transition is indeed the mechanism of neural activity remains unknown. This problem is what motivates this work; the work is an attempt to answer the question of whether or not the cell membrane undergoes a phase transition during an action potential.

Cells are very complicated objects, and neurons are no exception. The membrane is full of embedded proteins and other molecules, both the interior and the exterior contain positive and negative ions, and the interior is packed with an actin skeleton and many different kinds of organelles. Furthermore, they are living things, which means they can die in the course of an experiment.

Depending on what an experiment's goals are, sometimes it is appropriate to study a

simplified system instead; if chosen carefully, such a model system can remove many of the complications associated with working with living cells without invalidating the results. One example of a model which is often used in lieu of a cell membrane is a supported lipid bilayer, an artificial membrane without any of the ions, organelles and embedded molecules. This can be used to study phase transitions in lipid membranes, for example, but not, of course, travelling electric potentials.

Chapter 1 of this work first presents the knowledge —separated by field: thermodynamics, electronics, optics, and membrane biology— necessary to understand the remainder of the work, after which it describes the Hodgkin-Huxley and Heimburg-Jackson models for the action potential in detail.

Chapter 2 is a review of the results of some of the most important experiments carried out in the last century related to the action potential and of how well each of the models is able to explain said results. It concludes by determining the traits which a measurement technique suitable for optical imaging of action potentials should have if the technique is to shed light on the question of the presence or absence of an action-potential-related membrane phase transition.

Chapter 3 describes the methodology employed by the author of this work to study both supported lipid bilayers and live cells. These methods were not developed as part of this work. The chapter begins with the preparation of the samples to be studied and then describes the different experimental setups used: quantitative differential interference contrast, interferometric reflectometry, stimulated Raman scattering and electrophysiology. The theory and analysis procedure for interferometric reflectometry and part of the analysis process for stimulated Raman scattering data, which *were* developed as part of this work, are not included, as they are the subject matter of later chapters.

Chapter 4 presents the interferometric reflectometry theory, starting with the simplest case —that of light reflected by a single layer of a homogeneous material at an angle of 0° with respect to the direction of propagation of the light— and then generalising it to the case of oblique incidence. After this mathematical description of the total reflection coefficient of such a system and the signal detected by the interferometric reflectometry setup described in chapter 3, the various attempts at dealing with noise and the results obtained after data analysis are presented and discussed. It was hoped the technique could be used to study action potential propagation in live axons, but the issues generated by the extremely low signal-to-noise ratio of the data prevented this. Chapter 4 also contains the results of the preliminary electrophysiology experiments conducted. These served both to create familiarity with electrophysiological techniques and to determine the optimal properties of the cells and the cell-medium evaporation rate. Electrophysiological measurements of action potentials must be undertaken simultaneously with the (new) optical measurements if the latter are to be confirmed as indeed being appropriate for action potential imaging, since electrophysiology remains the established technique to observe action potentials.

Chapter 5 describes the data analysis procedure employed for the stimulated Raman scattering data, as well as the results obtained. It shows how stimulated Raman scattering

microscopy can be employed to distinguish different lipid bilayer phases even when the signal-to-noise ratio is low due to the small thickness of a single lipid bilayer. The chapter concludes with a discussion of the significance of the results presented therein and their relation to other results which have been obtained in the field.

Finally, chapter 6 provides a brief summary of this work and discusses future plans to continue working towards a better understanding of the cell membrane and whether phase transitions play a role in neural communication.

Contents

I	Background	I
1.1	Thermodynamics	2
1.1.1	Diffusion	2
1.1.2	Thermodynamic potentials	3
1.1.3	Phase transitions	5
1.2	Electronics	7
1.2.1	RC circuits	7
1.2.2	Circuits with multiple RC units	11
1.3	Optics	14
1.3.1	Polarisation and birefringence	14
1.3.2	Raman scattering	16
1.4	The cell membrane	21
1.4.1	The structure of the cell membrane	21
1.4.2	Flexoelectricity	23
1.4.3	The Nernst potential and the Goldman-Hodgkin-Katz equation	24
1.5	The Hodgkin-Huxley model	26
1.5.1	The membrane equation	27
1.5.2	The ionic current	28
1.5.3	The membrane current	32
1.5.4	The action potential in the Hodgkin-Huxley model	34
1.6	The Heimburg-Jackson model	35
1.6.1	Lipid phase transitions	36
1.6.2	Compression-wave propagation in membranes	37
1.6.3	The action potential in the Heimburg-Jackson model	38
1.7	References	42
2	Current understanding of action potentials and how to increase it	47
2.1	Mechanical aspects of action potentials	47
2.1.1	Axon radius, length and pressure changes	47
2.1.2	Electron-microscope evidence of membrane thickness changes	53
2.1.3	Action potential generation by mechanical stimuli	54

2.2	Optical aspects of action potentials	56
2.2.1	Membrane opacity changes	56
2.2.2	Membrane birefringence changes	57
2.2.3	Membrane scattering changes	60
2.3	Thermodynamic aspects of action potentials	62
2.3.1	Heat release	62
2.3.2	Action potential termination by heating	65
2.3.3	Action potential generation by cooling	67
2.4	Chemical aspects of action potentials	67
2.4.1	Action of anaesthetics and poisons	67
2.4.2	Anaesthetic potency reduction by pH	69
2.4.3	Action potential generation by chemical stimuli	71
2.5	The current state of our understanding of action potentials	72
2.6	Requirements for increasing our understanding of action potentials	74
2.7	References	75
3	Methods	81
3.1	Sample preparation	82
3.1.1	Supported lipid bilayers	82
3.1.2	Mouse neuron culture	86
3.1.3	Lobster giant axon and leg nerves	87
3.2	Quantitative differential interference contrast	93
3.2.1	Differential interference contrast	93
3.2.2	Making DIC quantitative	97
3.2.3	Obtaining the optical thickness of a sample	98
3.3	Interferometric reflectometry	100
3.3.1	Microscope and optics	100
3.3.2	Detection	103
3.3.3	Image acquisition	103
3.3.4	Regularisation	104
3.4	Coherent Raman scattering microscopy	105
3.4.1	Microscope and optics	105
3.4.2	Image acquisition	106
3.5	Electrophysiology	107
3.5.1	Pipette tip pulling	107
3.5.2	Intracellular and extracellular solution preparation	107
3.5.3	Lobster ventral nerve cord	108
3.5.4	Mouse neurons	108
3.6	References	113

4	Towards standalone optical measurement of neural activity with reflectometry	ii7
4.1	Interferometric reflectometry theory	117
4.1.1	Reflection by a thin layer	117
4.1.2	Obtaining thickness and refractive index from the reflection coefficient	121
4.1.3	Non-normal incidence	123
4.2	Dealing with noise	128
4.2.1	Phase filtering	129
4.2.2	Referencing	131
4.3	Line trace analysis	149
4.4	Electrophysiology	157
4.4.1	Lobster ventral nerve cord	157
4.4.2	Mouse neurons	158
4.5	Conclusions	159
4.6	References	162
5	Lipid domain discrimination by stimulated Raman scattering microscopy .	165
5.1	Data analysis	165
5.1.1	Spectral factorisation	165
5.1.2	PBS signal subtraction	167
5.1.3	Pump-Stokes overlap correction	168
5.1.4	PBS spectrum subtraction	170
5.1.5	Local PBS subtraction	173
5.1.6	Spatial binning	174
5.1.7	Simultaneous spectral factorisation	175
5.1.8	Denoising	177
5.1.9	Gradient correction by background fit	178
5.1.10	Remarks	180
5.2	Analysis of the 2:2:1 sample	181
5.2.1	Pre-factorisation analysis	181
5.2.2	Spectrum retrieval	182
5.2.3	Post-factorisation corrections to the spectra	182
5.2.4	Results from LO and DOPC samples	183
5.2.5	Results from unguided analysis of the ternary sample	186
5.2.6	Results from self-guided analysis of the ternary sample	187
5.2.7	Domain composition studies	189
5.3	Conclusions	190
5.4	References	192

6	Conclusions	195
A	Additional reflectometry theory	199
A.1	Remarks on some mathematical properties of the reflection coefficient	199
A.2	Reflection by multiple layers	201
A.2.1	Two layers	201
A.2.2	An arbitrary number of layers	203
B	Correlation algorithm for interferometric reflectometry	205
B.1	The code	205
B.2	Results with simulated data	209
C	Gradient descent algorithm for interferometric reflectometry	215
C.1	The code	215
C.2	Results with simulated data	223
D	Genetic algorithm for interferometric reflectometry	227
D.1	The code	227
D.2	Results with experimental data	242
E	Line trace analysis for interferometric reflectometry	245
F	Full results of FSC³ of SRS images	253
F.1	Figure 5.2	253
F.2	Figure 5.5	254
F.3	Figure 5.8	255
G	Parameters used for background fitting and FSC³	257
H	Publications and conference presentations	259
H.1	Publications	259
H.2	Conference presentations	259

Chapter I

Background

Scientists have been fascinated by how the nervous system works for centuries. Luigi Galvani discovered in the late eighteenth century that frog nerves reacted to electrical currents even after the frog's death, which led him to conclude that nerve activity involves travelling electricity.¹ His studies of nerve electricity were followed by those of Emil Heinrich duBois-Reymond, who half a century later discovered that such electrical activity was in the form of transient changes in the electric potential of the nerve,² which are today known as action potentials. A few decades later, Santiago Ramón y Cajal showed that nerves are made of individual cells³ (neurons).

In the beginning of the twentieth century, Julius Bernstein proposed that the action potential occurred due to an exchange of ions through the cell membrane due to changes in membrane permeability.⁴ This later served as the basis for Alan Lloyd Hodgkin and Andrew Fielding Huxley's theory of 1952,⁵ which modelled the cell membrane and the ion channels embedded in it as an electric circuit. Though it invokes mysterious "gating particles" which by turns activate and inactivate different ion channels but have never actually been detected, and though it contains numerous parameters which can be freely adjusted to fit experimental observations, the model was a breakthrough in the scientific understanding of neural activity: there was now a quantitative, if largely phenomenological, description of the action potential and its propagation along a neuron's axon.

Throughout the twentieth century, scientists conducted experiments of increasing complexity to elucidate the nature of the action potential: the mechanical, optical and thermodynamic changes the cell undergoes during activity. The objects of these experiments ranged from single cells to nerve bundles and from tiny myelinated mammalian neurons to the giant axons of crustaceans and cephalopods. Some of the articles published after 1952 attempted to explain the results presented in them under the light of the Hodgkin-Huxley model to varying degrees of success. It eventually became clear that the model suffered from several shortcomings in its ability to explain many of these findings, particularly the mechanical and thermodynamic ones.

The year 2005 saw the emergence of an alternative model to explain action potential propagation. This theory, proposed by Thomas Rainer Heimburg and Andrew D Jack-

son,⁶ views the cell membrane as a thermodynamic system near a phase transition; this (localised) transition drives a pressure pulse which, together with the action potential (the electric potential change), travels in a solitonic fashion down the axon; the action potential is seen as merely part of a broader phenomenon with thermodynamic, mechanical and electrical components.

In the fourteen years since its appearance, the thermodynamic model has received little attention from scientists outside of Heimburg and Jackson's research group. Perhaps this is due to its multidisciplinary nature or to its increased mathematical complexity compared to the Hodgkin-Huxley model; perhaps it is merely a consequence of the electrical model having been well established for decades among the biomedical community. Nevertheless, it is considerably more robust than its older counterpart in terms of being able to explain the wide variety of non-electrical experimental results of the previous century — but it is far from perfect, as there remain observations the model cannot account for.

Consequently, although our understanding of action potentials has come a long way since Galvani's experiments, there is still much work to be done. More-refined experiments would have the potential to elucidate more of the mechanism of neural activity, shed light on the limits of applicability of the Hodgkin-Huxley and Heimburg-Jackson theories, and ultimately, perhaps, allow us to construct a theory that is closer to the truth.

Chapter 2 will explore the ability of each of the two models to explain the more important experimental findings related to the action potential and outline the characteristics an experiment should have in order to provide the missing information alluded to in the above paragraph; the remainder of this work will deal with two such experiments. In order to fully understand all of this, however, some scientific background is necessary. That is the purpose of this chapter. The background theory will be presented by discipline: thermodynamics, electronics, optics and membrane biology. The models will then be described in detail.

1.1. Thermodynamics

1.1.1. Diffusion

The diffusion equation, derived by Adolf Fick in 1855,⁷ describes the diffusion of some quantity as a function of time and distance when there is an imbalance of said quantity between two regions; for example, it can describe the change in the temperature of a cylinder or the varying local density of a fluid in a chamber. For one-dimensional systems, it is

$$\frac{\partial u}{\partial t} = D \frac{\partial^2 u}{\partial x^2},$$

where u is the density of a quantity, t is time, x is distance, and D is the diffusion coefficient of whatever is quantified by u and has units of distance²/time. D is dependent on both what is quantified by u and what the medium in which it diffuses is.

1.1.2. Thermodynamic potentials

The first and second laws of thermodynamics

Consider a system with internal energy U on which some amount of work dW is performed by an external agent and/or into which some amount of heat dQ enters. The first law of thermodynamics states that the change of internal energy of the system is⁸

$$dU = dW + dQ.$$

Thus, a positive work (i.e. work performed on the system, as opposed to work performed by the system on its surroundings, which would be negative) causes a positive change of energy, as does a “flow” of heat into the system.

The nature of the work performed on the system is unimportant for the above relationship but depends on the macroscopic variables which describe the system. Thus, the work can be mechanical, electrical, chemical, and so on. The variables relevant to the mechanical case depend on the number of affected dimensions: the work to change the volume v of the system is

$$dW = -Pdv,$$

where P is the bulk pressure on the system; the work to change its surface area A is

$$dW = -\Pi dA,$$

where Π is the lateral pressure on the system; and the work to change the length l of a spring is

$$dW = -fdl,$$

where f is the force on the spring along the direction in which the spring expands and contracts. An example of an electrostatic case is the work to charge a capacitor,

$$dW = Vdq,$$

where V is the voltage difference between the capacitor's plates and dq is the change in the capacitor's charge.

A useful quantity in thermodynamics is the entropy S of a system, which characterises the amount of disorder in the system. It is, like volume and charge, an extensive variable, meaning that its total value for a system is equal to the sum of the values for any set of partitions of the system; contrast this with the intensive variables (such as pressure and voltage), which do not depend on the size of the system or how it is partitioned.

The second law of thermodynamics states that the entropy change in any spontaneous process is non-negative:

$$dS \geq 0.$$

Many processes in biology, however, decrease the entropy of an organism — at the cost of increasing that of its surroundings. The entropy is related to the temperature T of a system and to the amount of heat that enters it by⁹

$$dQ = TdS.$$

We thus have, for a three-dimensional system,

$$dU = -Pdv + Vdq + \dots + TdS, \quad (1.1)$$

where the ellipsis denotes any other contributions to dW . In general, we may write

$$dU = \sum_j x_j dy_j,$$

where the x_j are the intensive variables of the system with the appropriate sign (in the case of pressure, for instance, $x = -P$) and the y_j are their associated extensive variables.

Thermodynamic potentials

U is an example of a function called a thermodynamic potential. Other such functions are the enthalpy H , the Helmholtz free energy F and the Gibbs free energy G , which are defined as⁸

$$\begin{aligned} H &= U + Pv, \\ F &= U - Vq - TS, \\ G &= U + Pv - Vq - TS = F + Pv. \end{aligned}$$

Thus,

$$\begin{aligned} dH &= dU + Pdv + vdP = vdP + Vdq + \dots + TdS, \\ dF &= dU - Vdq - qdV - TdS - SdT = -Pdv - qdV + \dots - SdT, \\ dG &= dF + Pdv + vdP = vdP - qdV + \dots - SdT. \end{aligned}$$

The Gibbs free energy is particularly important because it is the energy available to the system to do work on its surroundings at constant pressure and temperature.

Dilation and compression

If we think of the volume of a substance as dependent on pressure and temperature, then its differential is

$$dv = \frac{\partial v}{\partial P} dP + \frac{\partial v}{\partial T} dT.$$

The quantities

$$\kappa = -\frac{1}{v} \frac{\partial v}{\partial P}$$

and

$$\beta = \frac{1}{v} \frac{\partial v}{\partial T}$$

are called, respectively, the substance's isothermal compressibility and dilation coefficient; they quantify the decrease in the substance's volume if pressure is applied to it and the increase in its volume if the temperature is increased, respectively. For most liquids, these vary slowly with pressure and temperature, so they may usually be taken to be constant when used to describe a liquid. With this assumption, we may divide dv by v and integrate to obtain a useful equation of state for liquids:

$$\ln\left(\frac{v_2}{v_1}\right) = -\kappa(P_2 - P_1) + \beta(T_2 - T_1), \quad (1.2)$$

where the subindices 1 and 2 refer to two equilibrium states of the liquid.

1.1.3. Phase transitions

Thermodynamic phases

A phase is a homogeneous constituent of a system. What is meant by “homogeneous” here is that the (macroscopic) physical properties of the constituent are uniform; another way of saying this is that the intensive variables describing the constituent (such as temperature, pressure, voltage and chemical potential — the latter of which depends on the particles' chemical composition) change continuously.^{10,11} For example, a salt dissolved in a homogeneous liquid forms a single phase, two immiscible liquids in a container form two different phases, and a single substance which is solid in some regions and liquid in others also forms two phases. The physical properties of a phase, then, depend only on the intensive variables.¹²

Two phases can be in thermodynamic equilibrium with each other only if their intensive variables have the same values. Some derivatives of the Gibbs and Helmholtz free energies with respect to temperature and/or pressure may be discontinuous at the boundary between the two phases; if this is the case, the transition is called an n th-order phase transition, where n is the order of the lowest-order discontinuous derivative; if the derivatives are all continuous, it is instead called a continuous phase transition.¹³ Phase transitions such as melting and vaporisation are first-order transitions.¹⁴

A phase transition can occur at multiple combinations of the values of the intensive variables describing the system. Thus, there is a region of dimension $N - 1$ in the system's phase space —where N is the number of extensive/intensive variable pairs (pressure and volume, temperature and entropy, and so on) describing the system— in which two phases can coexist in the same system. A well-known example of this phase coexistence is water at a pressure of 1 atm and a temperature of 0 °C; liquid water and ice can coexist, and, unless there is a temperature or pressure fluctuation, neither phase (liquid or solid) will undergo a transition to the other phase.

The Clausius-Clapeyron relation and similar relations

According to the second law of thermodynamics, any spontaneous process either increases the entropy of a system or does not alter it. Thus, equilibrium —the state in which the system undergoes no further changes— occurs when entropy has reached a maximum. At equilibrium, the variables describing the system do not change, so their derivatives with respect to time are all zero; therefore, the differentials of U , H , F and G are also zero: the thermodynamic potentials and the entropy reach their extremal points concomitantly. Because the free energies are proportional to $-S$, when the entropy reaches a maximum they reach a minimum.

The fact that at equilibrium G reaches its minimum means that two phases can only coexist in equilibrium if their Gibbs free energies are equal (otherwise only the phase with lower Gibbs free energy would exist).

Consider a change (brought about, for example, by some external agent) from a state in the phase-coexistence region with pressure P , voltage V , temperature T and Gibbs free energy $G(P, V, T)$ to another state in the coexistence region with pressure $P + \Delta_P$, voltage $V + \Delta_V$, temperature $T + \Delta_T$ and Gibbs free energy $G(P, V, T) + \Delta_G$. If Δ_P , Δ_V and Δ_T are small, Δ_G will be correspondingly small, so we may expand it to first order about zero:

$$\Delta_G \approx \Delta_P \left. \frac{\partial G}{\partial P} \right|_{\Delta_P=\Delta_V=\Delta_T=0} + \Delta_V \left. \frac{\partial G}{\partial V} \right|_{\Delta_P=\Delta_V=\Delta_T=0} + \Delta_T \left. \frac{\partial G}{\partial T} \right|_{\Delta_P=\Delta_V=\Delta_T=0}.$$

Let us now recall that

$$\begin{aligned} \frac{\partial G}{\partial P} &= v, \\ \frac{\partial G}{\partial V} &= -q, \\ \frac{\partial G}{\partial T} &= -S. \end{aligned}$$

Substituting this into the first-order expansion of Δ_G yields

$$\Delta_G \approx v\Delta_P - q\Delta_V - S\Delta_T.$$

Now, Δ_G is the same for both phases; using the subscripts 1 and 2 to differentiate between the phases, we may write

$$v_1\Delta_P - q_1\Delta_V - S_1\Delta_T = v_2\Delta_P - q_2\Delta_V - S_2\Delta_T.$$

Holding any one of the intensive variables constant and reducing our changes in the other two to infinitesimal ones, we obtain the slope of the coexistence curve on the corresponding two-dimensional phase diagram:

$$\begin{aligned} \frac{\partial P}{\partial V} &= \frac{q_2 - q_1}{v_2 - v_1}, \\ \frac{\partial P}{\partial T} &= \frac{S_2 - S_1}{v_2 - v_1}, \\ \frac{\partial V}{\partial T} &= -\frac{S_2 - S_1}{q_2 - q_1}. \end{aligned}$$

The second of these equations is called the Clausius-Clapeyron relation; we might call the other two “Clausius-Clapeyron-type relations”.

I.2. Electronics

I.2.1. RC circuits

Resistors and capacitors

The electric current flowing through a resistor is

$$I_R = \frac{V_R}{R},$$

where V_R is the voltage across the resistor and R is its resistance; this is known as Ohm’s law. The current flowing through a capacitor is

$$I_C = \frac{dq}{dt} = C \frac{dV_C}{dt},$$

where q is (as in section 1.1) the capacitor’s charge, C is its capacitance (here assumed constant) and V_C is the voltage across it. Kirchhoff’s law states that the total current leaving any given point in any circuit —taken, by convention, as positive current— equals the total current entering that point — taken as negative current;¹⁵ in other words, it states that charge is conserved in the circuit.

Two resistors with resistances R_1 and R_2 connected in parallel are equivalent to a single resistor with effective resistance R_{eff} . According to Ohm’s law, the current flowing through R_j is

$$I_j = \frac{V_R}{R_j}$$

with $j \in \{1, 2\}$. By Kirchhoff’s law, the total current flowing from one side of the two-resistor system to its other side is $I = I_1 + I_2$. However, by replacing the two resistors in parallel with the single resistor R_{eff} , we obtain

$$I = \frac{V_R}{R_{\text{eff}}}.$$

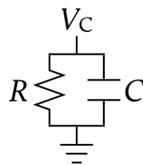


Figure 1.1: An RC circuit.

Thus,

$$\frac{1}{R_{\text{eff}}} = \frac{1}{R_1} + \frac{1}{R_2}.$$

Similarly, two capacitors connected in parallel are equivalent to a single capacitor with effective capacitance C_{eff} . As before, we have

$$I_j = C_j \frac{dV_C}{dt}$$

and $I = I_1 + I_2$; on the other hand, replacing the two capacitors in parallel with the single capacitor C_{eff} ,

$$I = C_{\text{eff}} \frac{dV_C}{dt}.$$

Hence,

$$C_{\text{eff}} = C_1 + C_2.$$

RC circuits without a voltage source

An RC circuit (figure 1.1) is a circuit formed by a resistor and a capacitor connected in series without any external current sources. Using Kirchhoff's law for a point between the resistor and the capacitor and noting that $V_R = V_C$, we have

$$\frac{V_C}{R} + C \frac{dV_C}{dt} = 0,$$

which has the solution

$$V_C = V_0 e^{-\frac{t}{RC}},$$

where $V_0 = V_C(t = 0)$ is the capacitor's initial charge. This means that the voltage across the capacitor takes a time $\tau = RC$ to reach $1/e$ times its original value. τ is called the time constant of the RC circuit, as it determines how quickly the capacitor's plates are charged.

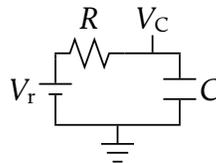


Figure 1.2: An RC circuit with a voltage source.

RC circuits with a constant voltage source

If the circuit has a source of constant voltage V_r connected in series to the resistor (figure 1.2), we have $V_r + V_R = V_C$ and

$$\frac{V_C - V_r}{R} + C \frac{dV_C}{dt} = 0.$$

Making the substitution $V = V_C - V_r = V_R$, we obtain

$$\frac{V}{R} + C \frac{dV}{dt} = 0,$$

which is identical to the equation for an RC circuit without a voltage source; thus, the solution is

$$V_C = V_r + V_0 e^{-\frac{t}{RC}},$$

where V_0 is now the excess of V_C with respect to V_r at time $t = 0$. If the capacitor is initially charged to a voltage greater than V_r , it will discharge down to V_r ; otherwise, it will charge up to V_r ; in the special case in which it is initially discharged, $V_0 = -V_r$. Note that the time constant of the capacitor has remained the same.

Pressure on the plates of a charged capacitor

The charge on the plates of a capacitor produces an electric field \vec{E} between them. If the surface area A of the plates is much greater than the distance between them, we may consider the plates to be infinite in extent when performing calculations; in this approximation, the electric field is perpendicular to the plates everywhere. Taking a cylindrical surface \mathfrak{S} of base area Δ_A and height $2\Delta_z$ surrounding an area Δ_A of one plate (figure 1.3) and using Gauss's law in its integral form,¹⁶ we have

$$\frac{\Delta q}{\epsilon} = \int_{\mathfrak{S}} \vec{E} \cdot d\vec{A} = 2\Delta_A E,$$

where $\epsilon = n^2 \epsilon_0$ is the electric permittivity of the (dielectric) material between the plates, n is the material's refractive index, ϵ_0 is the permittivity of vacuum and Δq is the charge

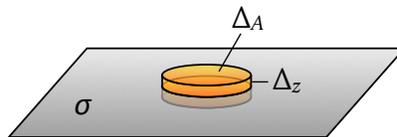


Figure 1.3: An imaginary cylindrical surface (yellow) surrounding a portion of a plate with uniform charge density σ .

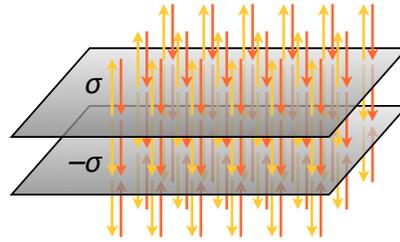


Figure 1.4: The electric fields produced by the top (yellow arrows) and bottom (orange arrows) plates of a capacitor.

of the enclosed portion of the plate (\mathcal{S} is assumed to be small enough for E to be approximately constant throughout it). Note that this does not depend on the height of \mathcal{S} , indicating that E is indeed uniform on either side of the plate and we could have taken an arbitrarily sized cylindrical surface.

The electric field due to a single charged plate, then, is

$$E = \frac{\Delta q}{2\Delta_A n^2 \epsilon_0} = \frac{\sigma}{2n^2 \epsilon_0},$$

where σ is the charge density. Now, a capacitor consists of two plates with opposite charges, so the second plate produces a field $-\vec{E}$ (figure 1.4). Note that the fields outside the capacitor cancel each other out and the fields inside the capacitor add up, yielding

$$E = \frac{\sigma}{n^2 \epsilon_0}$$

between the plates and $E = 0$ everywhere else.

Because the capacitor's plates are charged, this electric field exerts a force on them. Each plate thus feels an attractive force equal to its charge multiplied by the other plate's electric field:

$$F = \frac{\sigma q}{2n^2 \epsilon_0} = \frac{\sigma^2 A}{2n^2 \epsilon_0}.$$

Hence, the electrostatic pressure on each plate is

$$P = \frac{\sigma^2}{2n^2 \epsilon_0}. \quad (1.3)$$

For example, suppose a capacitor of $1 \mu\text{F}$ whose plates have an area of 0.3 cm^2 and are separated by a dielectric of refractive index 1.5 is charged to 100 mV. The charge on each plate is then 10^{-7} C , meaning that the charge density is $3.33 \times 10^{-3} \text{ C/m}^2$. This generates a pressure of about $2.79 \times 10^5 \text{ Pa}$ or 2.75 atm.

1.2.2. Circuits with multiple RC units

Two RC units

Suppose we have a circuit formed by two identical units, each of which is an RC circuit like that of figure 1.2, connected in parallel (figure 1.5, left); we will assume that V_r is constant. Thévenin's theorem states that any circuit formed by two terminals (the voltage across which we are interested) and any number of resistors, voltage sources and/or current sources is equivalent to a circuit formed by the same two terminals with a single voltage source and a single resistor connected in series between them;¹⁵ can we do the same for a circuit formed by resistors, capacitors and voltage sources (figure 1.5, right)?

To show that we can, we take into consideration the currents flowing through each loop of the circuit (figure 1.6, left) and of a circuit formed by exchanging the positions of the second resistor and the first capacitor-voltage subunit (figure 1.6, right).

In the former case, we have

$$\begin{aligned}\frac{V_C + V_r}{R} &= I_1 = -I_2 + I_3, \\ C \frac{dV_C}{dt} &= I_1 - I_2 = I_3,\end{aligned}$$

whereby $I_2 = 0$ and $I_1 = I_3$. Each resistor, then, has a current I_3 flowing through it, as does each capacitor and each voltage source. Because $I_2 = 0$, this circuit is equivalent to two isolated RC circuits with time constant $\tau = RC$.

In the latter case, we have

$$\begin{aligned}\frac{V_C + V_r}{R} &= I_1 = I_2 - I_3, \\ C \frac{dV_C}{dt} &= I_2 - I_3 = I_3,\end{aligned}$$

whereby $I_2 = 2I_1 = 2I_3$. Each component still has a current I_3 flowing through it. Now, as we saw at the beginning of section 1.2.1, this circuit is equivalent to a single RC circuit like the one on the right side of figure 1.5 with $R_{\text{eff}} = R/2$ and $C_{\text{eff}} = 2C$ (to see that the capacitor-voltage subunits in parallel are equivalent to a single capacitor-voltage subunit, we note that, because the two capacitors are connected to identical voltage sources, we

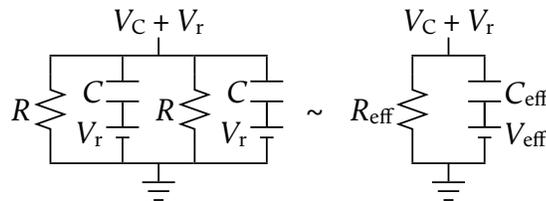


Figure 1.5: A circuit formed by two identical RC circuits in parallel and a possible equivalent circuit.

can replace the two subunits with a single voltage source V_r connected in series with two capacitors C in parallel with each other). For said circuit, we have

$$I = \frac{V_C + V_r}{R_{\text{eff}}} = C_{\text{eff}} \frac{dV_C}{dt},$$

which leads to $R_{\text{eff}}I = RI_1$ and

$$\frac{I}{C_{\text{eff}}} = \frac{I_1}{C}.$$

Thus,

$$\frac{R}{R_{\text{eff}}} = \frac{C_{\text{eff}}}{C},$$

which finally yields $\tau_{\text{eff}} = R_{\text{eff}}C_{\text{eff}} = RC = \tau$.

Many RC units

By showing that the time constants and the current through each of the components of the two circuits on figure 1.6 are the same, we have proven that the circuits are equivalent to each other and thus that the equivalence of figure 1.5 holds. In what follows we will prove that, if the same is true for a circuit formed by an arbitrary number N of RC units, it is also true for one formed by $N + 1$ units.

If the equivalence holds for N units, we can rearrange the resistors and the capacitor-voltage subunits of the first N units so that all the resistors are together and all the capacitor-voltage subunits are together (figure 1.7, centre). We then note that this is equivalent to a circuit formed by only two units, one of which has a resistor R/N and a capacitor

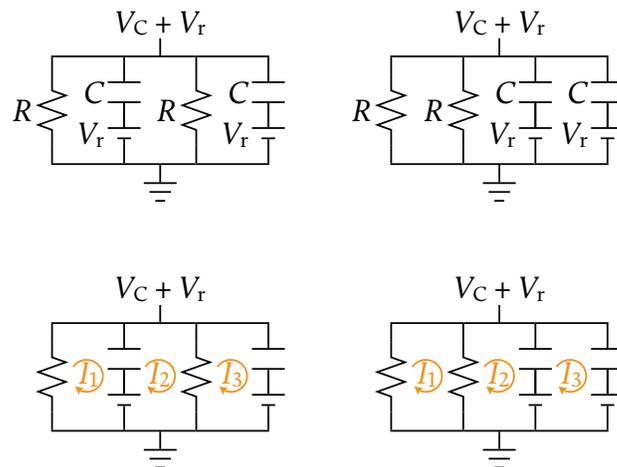


Figure 1.6: A circuit formed by two identical RC circuits in parallel and a circuit with the resistors and the capacitor-voltage subunits rearranged.

NC and the other of which has a resistor R and a capacitor C (figure 1.7, bottom). As before, we look at the current in each loop.

In the case shown at the bottom of figure 1.7, we have

$$\begin{aligned} \frac{V_C - V_r}{R} &= \frac{I_1}{N} = -I_2 + I_3, \\ C \frac{dV_C}{dt} &= \frac{I_1 - I_2}{N} = I_3, \end{aligned}$$

which implies that once again $I_2 = 0$ and that $I_1 = NI_3$. Thus, the resistor R/N and the capacitor NC have a current NI_3 flowing through them while the resistor R and the capacitor C have a current I_3 flowing through them. The time constants of the two isolated RC circuits (to which this circuit is again equivalent on account of I_2 being equal to zero) are $\tau = RC$, as before.

In the case analogous to the one shown on the right of figure 1.6 (obtained by exchanging the two central columns of the former case with each other), we have

$$\begin{aligned} \frac{V_C - V_r}{R} &= \frac{I_1}{N} = I_2 - I_1, \\ C \frac{dV_C}{dt} &= \frac{I_2 - I_3}{N} = I_3, \end{aligned}$$

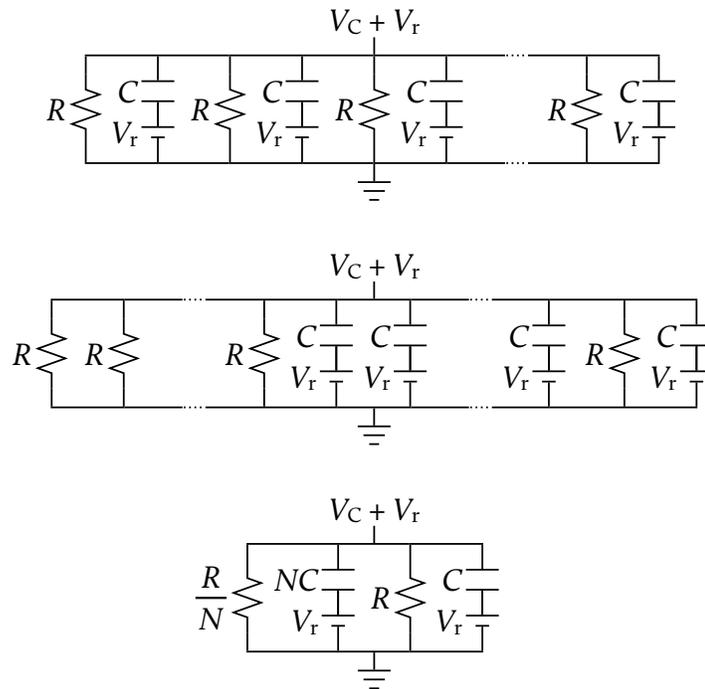


Figure 1.7: A circuit formed by $N + 1$ identical RC circuits (top), a rearrangement of the first N units (centre) and an equivalent circuit of the rearrangement (bottom).

whereby

$$I_2 = \frac{N+1}{N} I_1 = (N+1)I_3$$

and again $I_1 = NI_3$. Once again, the resistor R/N and the capacitor NC have a current equal to NI_3 while the resistor R and the capacitor C have a current equal to I_3 flowing through them. As before, this circuit is equivalent to a single RC circuit with

$$\frac{1}{R_{\text{eff}}} = \frac{N}{R} + \frac{1}{R},$$

which yields $R_{\text{eff}} = R/(N+1)$, and $C_{\text{eff}} = NC + C = (N+1)C$. The time constant is then $\tau_{\text{eff}} = R_{\text{eff}}C_{\text{eff}} = RC = \tau$.

This proves that the equivalence remains true for any number of RC units.

Continuous limit

In the limit in which $N \rightarrow \infty$ with finite R_{eff} and C_{eff} , it is useful to use the conductance G , defined as

$$G = \frac{1}{R},$$

instead of the resistance; in this case, we can define a conductance per unit length (or per unit area if the array is two-dimensional) and a capacitance per unit length or unit area; these are called the specific conductance (denoted by g) and the specific capacitance (denoted by c). Note that the specific resistance r , defined as the inverse of g , has units of resistance times length. Note also that $\tau = RC = rc$.

1.3. Optics

1.3.1. Polarisation and birefringence

Ordinary and extraordinary refractive indices of a uniaxial material

When an electromagnetic field \vec{E} passes through a material, it polarises its constituent atoms, thereby forming small dipoles which create an electric field in the opposite direction as \vec{E} . If the material has permittivity ϵ , then the degree to which this polarisation occurs is characterised by the electric displacement vector¹⁷ $\vec{D} = \epsilon\vec{E}$.

In the preceding section, ϵ was treated as a scalar. In general, however, it is a non-singular symmetric 3-dimensional tensor,¹⁸

$$\epsilon = \begin{pmatrix} \epsilon_{11} & \epsilon_{12} & \epsilon_{13} \\ \epsilon_{12} & \epsilon_{22} & \epsilon_{23} \\ \epsilon_{13} & \epsilon_{23} & \epsilon_{33} \end{pmatrix},$$

and the components of the electric displacement vector are given by

$$D_j = \sum_{k=1}^3 \epsilon_{jk} E_k.$$

Because ϵ is symmetric, there always exists a coordinate system in which it is diagonal. The axes in this coordinate system are called the material's optic axes. Throughout this section, we shall only use such coordinate systems and write the components of the diagonal of ϵ as $\epsilon_x = \epsilon_{11}$, $\epsilon_y = \epsilon_{22}$ and $\epsilon_z = \epsilon_{33}$. If $\epsilon_x = \epsilon_y = \epsilon_z$, then ϵ is a scalar. Materials for which this is not so are called birefringent. Birefringence is a consequence of a material being polarisable (meaning its positive and negative centres of charge are displaced) to different degrees in different directions, which itself is due to an anisotropic molecular structure. Examples of birefringent materials are non-cubic crystals;¹⁸ biological membranes;¹⁹ DNA;²⁰ and colloids formed by non-spherical particles oriented non-randomly in a homogeneous liquid, such as blood.²¹

Let us consider a material with $\epsilon_x \neq \epsilon_y = \epsilon_z$. We write $\epsilon_e = \epsilon_x$ and $\epsilon_o = \epsilon_y$. The equality of the y and z components of ϵ means that a rotation by any angle ψ about the x axis will leave ϵ invariant:

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos(\psi) & -\sin(\psi) \\ 0 & \sin(\psi) & \cos(\psi) \end{pmatrix} \begin{pmatrix} \epsilon_e & 0 & 0 \\ 0 & \epsilon_o & 0 \\ 0 & 0 & \epsilon_o \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos(\psi) & \sin(\psi) \\ 0 & -\sin(\psi) & \cos(\psi) \end{pmatrix} = \begin{pmatrix} \epsilon_e & 0 & 0 \\ 0 & \epsilon_o & 0 \\ 0 & 0 & \epsilon_o \end{pmatrix}.$$

This means that the material only has one optic axis, which is the x axis. Such materials are called uniaxial.

In section 1.2 we assumed that the materials the circuit components were made of were non-magnetic (i.e. that $\mu \approx \mu_0$) and thus their permittivities and refractive indices were related by $\epsilon = n^2 \epsilon_0$. We shall make the same assumption here and, indeed, throughout this work, as biological materials are in general non-magnetic.²²

The fact that ϵ is a tensor means that a birefringent material will have different refractive indices for different polarisations of the electric field: the x component of \vec{E} (that parallel to the material's optic axis) will "see" a refractive index

$$n_e = \sqrt{\frac{\epsilon_e}{\epsilon_0}},$$

while any component perpendicular to the optic axis will "see" a refractive index

$$n_o = \sqrt{\frac{\epsilon_o}{\epsilon_0}}.$$

The subscripts e and o stand for "extraordinary" and "ordinary", respectively. The quantity $\Delta_n = n_e - n_o$ is called the birefringence of the material.

Phase change due to a uniaxial material

Consider a plane wave travelling in the z direction through a material with refractive index n . The electric field \vec{E} of this wave is then contained in the xy plane. The polarisation of \vec{E} depends on the relative phase φ between its components; for example, $\varphi = \pm\pi/2$ corresponds to circular polarisation if the x and y components have the same magnitude. If the wave travels a distance a through the material, its polarisation remains unchanged, as the phases of both components evolve in the same way.

If the material is birefringent and its optic axis lies on the xy plane (for simplicity, let us assume it is the x axis), then the relative phase changes by an amount $\Delta\varphi = (k_x - k_y)a$, where k_x is the wave number of E_x and similarly for the y component. The reason φ —and thus the polarisation—changes is that k_x and k_y are different: $k_x = n_e k_0$ and $k_y = n_o k_0$, where k_0 is the wave number of \vec{E} in vacuum. We may thus rewrite the change of relative phase as

$$\Delta\varphi = \frac{2\pi\Delta_n a}{\lambda},$$

where λ is the wavelength of \vec{E} in vacuum.

If the material's optic axis is the z axis, both components of the polarisation see the same refractive index and there is again no phase change.

1.3.2. Raman scattering

Spontaneous Raman scattering

A material's polarisation density, which to first order is given by

$$\vec{P} = (\epsilon - \epsilon_0)\vec{E} = \epsilon_0\chi_1\vec{E},$$

describes its reaction to an external electric field \vec{E} . Here, χ_1 is the material's first-order electric susceptibility, which, like the permittivity, is a 3-dimensional tensor acting on \vec{E} . An oscillating electric field

$$\vec{E} = \vec{E}_0 e^{i(\vec{k}\cdot\vec{r} - \omega t)}$$

will induce a phonon—a localised vibration wave—of amplitude A_0 , wave vector \vec{a} and frequency $\omega_0 < \omega$ in a material it is incident on. This oscillation may be written as

$$A = A_0 \left(e^{i(\vec{a}\cdot\vec{r} - \omega_0 t)} + e^{-i(\vec{a}\cdot\vec{r} - \omega_0 t)} \right).$$

The material's susceptibility is a function of A .²³ For small A (which occurs if \vec{E} is sufficiently weaker than the fields holding the nuclei of the material's molecules together²⁴), we may expand $\chi_1(A)$ to first order about $\chi_1(0)$:

$$\chi_1(A) \approx \chi_1(0) + \left. \frac{d\chi_1}{dA} \right|_{A=0} A.$$

Putting all of this together, we obtain

$$\begin{aligned} \vec{P} \approx & \epsilon_0 \chi_1(0) \vec{E}_0 e^{i(\vec{k} \cdot \vec{r} - \omega t)} \\ & + \epsilon_0 A_0 \left. \frac{d\chi_1}{dA} \right|_{A=0} \vec{E}_0 e^{i((\vec{k}-\vec{a}) \cdot \vec{r} - (\omega-\omega_0)t)} \\ & + \epsilon_0 A_0 \left. \frac{d\chi_1}{dA} \right|_{A=0} \vec{E}_0 e^{i((\vec{k}+\vec{a}) \cdot \vec{r} - (\omega+\omega_0)t)}. \end{aligned}$$

This shows that \vec{P} has three components: one with the original angular frequency ω , one with angular frequency $\omega - \omega_0$ and one with angular frequency $\omega + \omega_0$. The material may release the energy acquired in its interaction with \vec{E} by emitting light of any of these frequencies.

To find the value of ω_0 , it is necessary to view the interaction between the material and the electric field from a microscopic perspective.

The energy of an electron in an atom or molecule is quantised, which means it can only have certain values. The same is true for atomic nuclei in molecules. The energies in the electronic case are related to the orbitals the electrons occupy, while those in the nuclear case are related to the vibrational and rotational modes of the nuclei with respect to each other. The lowest-energy state of an atom or molecule is called the ground state, and all other states are called excited states. In the case of the excitation of a phonon in the material by \vec{E} , the material's molecules are undergoing transitions from one vibrational mode or state to another. Thus, $\omega_0 \hbar$ must equal the energy difference between two vibrational states.

If \vec{E} comes from incident light, the emitted light with angular frequency $\omega \mp \omega_0$ is the result of inelastic scattering of the incident light (figure 1.8). This type of inelastic light scattering is called Raman scattering in honour of Chandrasekhara Venkata Raman, who reported the phenomenon in 1928,²⁵ about four and a half years after its prediction by Adolf Smekal.²⁶ If the energy difference between a scattered photon and an incident photon is $\Delta_U = \hbar(\omega_r - \omega_i) < 0$, where ω_r is the angular frequency of the scattered light and ω_i is that of the incident light, it is called Stokes Raman scattering; this is in honour

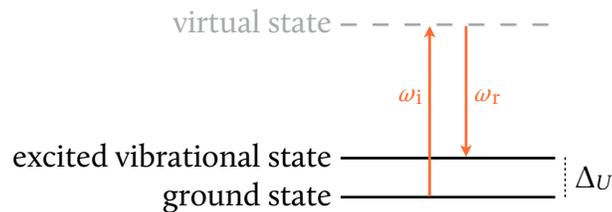


Figure 1.8: Diagram of spontaneous Stokes Raman scattering. An incident beam excites the transition between vibrational energy states separated by an energy Δ_U and is inelastically scattered, producing a scattered beam with angular frequency $\omega_r = \omega_i - \Delta_U/\hbar$.

of George Gabriel Stokes, who in 1852 extensively described fluorescence,²⁷ which also consists of a shift to a lower frequency; if $\Delta_U > 0$, it is accordingly called anti-Stokes Raman scattering.²⁸

By measuring and graphing the amount of photons of each frequency which are scattered by a material, a Raman spectrum of the material can be constructed. The fact that each type of molecule has a unique collection of vibrational energy levels—and thus a unique Raman spectrum—makes Raman scattering a powerful tool for noninvasively determining the chemical composition of a sample.

Stokes Raman scattering has a higher probability of occurring than anti-Stokes Raman scattering because of the energy distribution of molecules. The probability that a given molecule in a sample will be in a state with energy U is proportional to $e^{-U/k_B T}$, where k_B is Boltzmann's constant and T is the temperature of the sample;²⁹ therefore, given two energy states separated by an energy difference $\Delta_U > 0$, the ratio of molecules in the higher state to molecules in the lower state is $e^{-\Delta_U/k_B T} < 1$. As a result, while the positions of the peaks in any Raman spectrum are symmetric about $\Delta_U = 0$ (because photons scattered on the Stokes side and those scattered on the anti-Stokes side correspond to the same transitions between molecular vibrational states), the Stokes-side peaks are usually stronger and give a better signal in practice. However, using the Stokes-side signal makes Raman scattering and fluorescence (which is also Stokes-side, but it involves a transition to a higher energy state followed by a multi-step return to the lower-energy state, while Raman scattering involves a transition to a “virtual” higher energy state followed by a transition to a state different from the one the molecule started at) difficult to separate experimentally without post-detection data processing. This problem is diminished by the fact that, while what matters in fluorescence is the wavelength of the incident light, what matters in Raman scattering is the energy difference between the scattered light and the incident light.

Raman shifts—that is, the energy differences between the scattered photons and the incident ones—are usually calculated as

$$\Delta_k = \frac{1}{\lambda_r} - \frac{1}{\lambda_i},$$

where λ_r is the wavelength of the scattered light and λ_i is that of the incident light (the factor of 2π is omitted out of convention), and are measured in cm^{-1} . For example, if the incident light has a wavelength of 550.0 nm, then scattered light with a wavelength of 652.2 nm corresponds to a Raman peak at about $2,850 \text{ cm}^{-1}$ on the Stokes side.

In spontaneous Raman scattering, the vibrations (and thus the scattered light) are created by a combination of the external field and the material itself; because the material's molecules vibrate thermally (that is, incoherently), the resulting vibrations are incoherent.

Let us write the electric field of the light scattered by a single molecule in some direction as $E = E_0 e^{i\varphi}$. The intensity of the light scattered by that molecule is proportional to $|E|^2 = E_0^2$. The total intensity of the light scattered by N identical molecules in that

direction is then

$$I = \frac{\epsilon_0 c}{2} \left| \sum_{j=1}^N E_j \right|^2 = \frac{\epsilon_0 c E_0^2}{2} \sum_{j=1}^N \left(1 + \sum_{\ell \neq j} e^{i(\varphi_j - \varphi_\ell)} \right).$$

For large N ,

$$\sum_{j=1}^N \sum_{\ell \neq j} e^{i(\varphi_j - \varphi_\ell)} \approx 0$$

because the molecules scatter light incoherently (i.e. the φ_j are randomly distributed). The intensity is thus

$$I = \frac{\epsilon_0 c N E_0^2}{2} \sim N.$$

Coherent anti-Stokes Raman scattering

For stronger electric fields, the first-order expression for the polarisation density is no longer valid. We write

$$\vec{P} = \epsilon_0 \sum_j \chi_j \vec{E}^j = \sum_j P_j,$$

where each χ_j is a tensor which acts on j vectors (hence the exponent of \vec{E}). For a given j , χ_j is called the j th-order electric susceptibility of the material. The second-order susceptibility of centrally symmetric materials is zero, so we shall focus on the first universally-nonzero nonlinear term, χ_3 . When χ_2 is not zero, it is generally of the order of 10^{-12} m/V; χ_3 is generally of the order of $10^{-24} \chi_1$ m²/V².³⁰

Let us now consider two high-power incident fields

$$\begin{aligned} \vec{E}_p &= \frac{1}{2} \vec{E}_{p_0} \left(e^{i(\vec{k}_p \cdot \vec{r} - \omega_p t)} + e^{-i(\vec{k}_p \cdot \vec{r} - \omega_p t)} \right), \\ \vec{E}_s &= \frac{1}{2} \vec{E}_{s_0} \left(e^{i(\vec{k}_s \cdot \vec{r} - \omega_s t)} + e^{-i(\vec{k}_s \cdot \vec{r} - \omega_s t)} \right) \end{aligned}$$

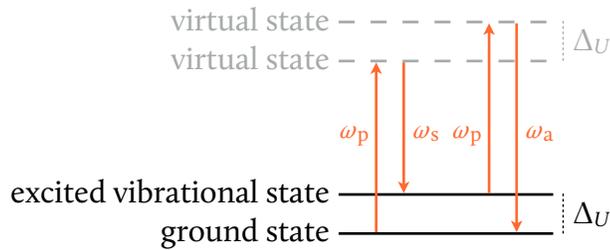


Figure 1.9: Diagram of coherent anti-Stokes Raman scattering. The anti-Stokes beam, with angular frequency $\omega_a = 2\omega_p - \omega_s$, is detected.

with $\omega_s < \omega_p$; in the context of light waves, \vec{E}_p is called the pump beam and \vec{E}_s is called the Stokes beam because it is Stokes-shifted with respect to \vec{E}_p (i.e. it has a lower frequency than \vec{E}_p). The total field is then $\vec{E} = \vec{E}_p + \vec{E}_s$. Because χ_3 is a tensor which acts on the electric field three times, P_3 is proportional to E^3 :

$$P_3 = \frac{\epsilon_0 \chi_3}{8} \left(E_{p0} \left(e^{i(\vec{k}_p \cdot \vec{r} - \omega_p t)} + e^{-i(\vec{k}_p \cdot \vec{r} - \omega_p t)} \right) + E_{s0} \left(e^{i(\vec{k}_s \cdot \vec{r} - \omega_s t)} + e^{-i(\vec{k}_s \cdot \vec{r} - \omega_s t)} \right) \right)^3.$$

P_3 therefore has, in addition to the original frequencies ω_p and ω_s , six combination frequencies: $3\omega_p$, $2\omega_p - \omega_s$, $2\omega_p + \omega_s$, $2\omega_s - \omega_p$, $2\omega_s + \omega_p$ and $3\omega_s$. In particular, $\omega_a = 2\omega_p - \omega_s$ is called the anti-Stokes frequency because it is shifted from ω_p in the anti-Stokes direction by the same amount as the Stokes beam (figure 1.9); the ω_a term of P_3 has amplitude $3\epsilon_0 \chi_3 E_{p0}^2 E_{s0}$.

This time, in contrast to spontaneous Raman scattering, the vibrations are driven by the two incident fields, which are coherent with each other. The vibrations are thus coherent with the incident fields, and so is the scattered light — hence the name “coherent anti-Stokes Raman scattering”, or CARS. If $\omega_p - \omega_s$ corresponds to the energy difference between two vibrational states of the material, the molecules vibrate resonantly and the amount of scattered light is greatly increased.

The CARS intensity generated by N identical molecules is

$$I \sim |\chi_3|^2 N^2 E_{p0}^4 E_{s0}^2 \sim N^2$$

because the scattering is coherent and thus all the φ_j are equal.

Stimulated Raman scattering

If, instead of detecting light at the anti-Stokes frequency, we detect light at the Stokes frequency ω_s or at the pump frequency ω_p , we note that the coherent Raman scattering process results in a net decrease in the number of photons of frequency ω_p and a net increase in the number of photons of frequency ω_s because the Stokes-side process is more efficient (occurs more often) than the anti-Stokes-side process: the pump beam intensity

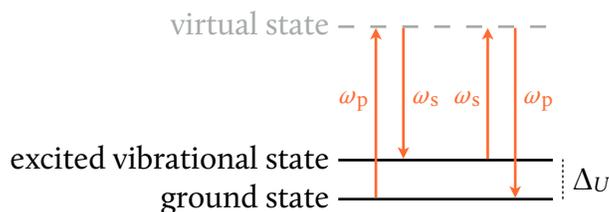


Figure 1.10: Diagram of stimulated Raman scattering. The higher efficiency of Stokes-side scattering compared to anti-Stokes-side scattering results in a net increase in the Stokes beam intensity and a corresponding decrease in the pump beam intensity.

decreases, while the Stokes beam intensity increases (figure 1.10). This process is called “stimulated Raman scattering”, or SRS, because the Stokes beam stimulates the conversion from ω_p to ω_s .

This is still a nonlinear process driven by two beams which can be coherent with each other, so it provides, like CARS, a signal orders of magnitude stronger than spontaneous Raman scattering does. At the same time, it only involves two frequencies rather than three, so, while the signal (which is proportional to $|\chi_3|N$ instead of to $|\chi_3|^2N^2$)³¹ is typically weaker than that of CARS, it does not suffer from a non-resonant background.^{32,33}

1.4. The cell membrane

1.4.1. The structure of the cell membrane

The fluid mosaic model

According to the fluid mosaic model proposed by Seymour Jonathan Singer and Garth Nicolson in 1972, the cell membrane comprises a lipid bilayer with embedded proteins.³⁴ The bilayer is typically 4–6 nm thick and, considering most cells have a radius between a few hundred nanometres and a few microns, has an area of the order of 10^6 – 10^8 nm².^{35,36} The area can thus be considered infinite compared to the thickness for most purposes (for example, because of its great area compared to its thickness and because membranes are closed surfaces, one may ignore edge effects when considering the electric properties of the membrane).

Lipids

Membrane lipids consist of a large hydrophilic head group and one or two hydrophobic hydrocarbon tails (figure 1.11). Therefore, they are amphiphilic and have cylindrical or conical shapes depending on the relative sizes of the head and tail groups and the relative orientation of the tail groups in the case in which there are two. Their amphiphilic nature causes them to form ordered structures, such as bilayers, when immersed in an aqueous medium³⁷ (see also section 3.1.1). Under physiological conditions, their tails are nearly fully extended.³⁶

The head group and tail(s) of a lipid molecule have different sizes, and the relationship between these sizes determines the overall shape of the molecule. Cell membranes are asymmetric in that their exterior and interior compositions are different, leading to differences in the shapes of the lipid molecules and thus their orientations relative to adjacent molecules. In general, this causes the membrane to be curved,³⁷ which in turn contributes to the fact that the voltage V_m across it is nonzero. The other contribution to the membrane voltage comes from differences in the concentrations of the ions which exist in the intracellular and extracellular media. Sections 1.4.2 and 1.4.3, respectively, will discuss these two contributions to V_m in further detail.

Individual lipid molecules have a large lateral diffusion rate (i.e. they are highly mobile within the membrane plane), making the membrane fluid. Lateral movement consists of

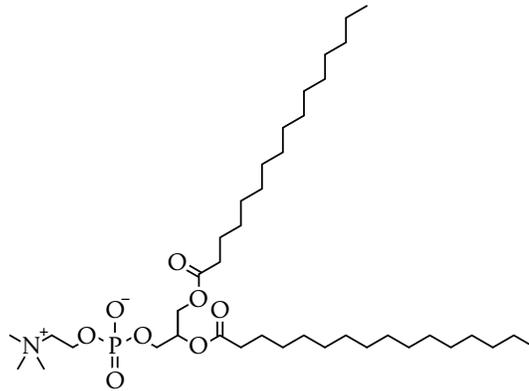


Figure 1.11: Chemical structure of dipalmitoylphosphatidylcholine (DPPC), an important lipid in some cell membranes.

alternating brownian motion and large jumps with a period of the order of 10 ms.³⁶ Thermodynamic variables such as temperature, pressure, pH and chemical structure affect the phase of the membrane;³⁷ cooling or an increase in pressure, for example, produce a phase transition from the physiological liquid state (called the liquid-disordered state) to a gel-like solid state (called the solid-ordered state) where the hydrocarbon tails are completely extended.³⁶ In this state, the tails are tilted with respect to the direction perpendicular to the membrane plane; the tilt increases as the concentration of water around the bilayer increases. In the solid-ordered state the lateral diffusion rate of the lipids decreases.³⁷ The melting temperature T_m of the lipid bilayer grows with tail length and saturation and is normally about 15 °C below body temperature,⁸ which, as mentioned earlier, means that under physiological conditions the membrane is in the liquid-disordered state; as their body temperature changes, some organisms are capable of modifying the composition of their cell membranes in such a way that fluidity is maintained.³⁶

In addition to lateral diffusion, lipid molecules are capable of undergoing rotations about their axes. In contrast, however, they very rarely move from one layer of the membrane to the other.³⁶ This means that asymmetry—and thus also V_m —is maintained.

Most membranes are made of at least two different types of lipid; a single eukaryotic cell can have more than a thousand different kinds of lipid.³⁸ These types have an inhomogeneous lateral distribution and are believed to sometimes organise themselves in homogeneous domains called lipid rafts.^{39,40}

Membrane proteins and ions

Membrane proteins are much larger than lipid molecules. Like lipids, they are amphiphilic; their hydrophobic part is embedded in the membrane, and their hydrophilic part(s) is/are exposed to one or the other surface of the membrane. Also like lipids, they can move laterally unless they are bound to cellular components, though their lateral diffusion rate drops as temperature decreases. Different proteins have different diffusion

rates, and some proteins occur only in certain membranes and only on one side of the membrane (if embedded in the membrane) or in one orientation with respect to the membrane (if spanning the entire membrane and exposed to both sides), suggesting they play an important part in very specific interactions;³⁶ however, the extent of their functional role is still debated (see sections 1.5 and 1.6).

Biological membranes are selectively permeable to different ions and molecules, and this selective permeability changes depending on various factors.³⁶ They are also flanked by several layers of ordered water.⁴¹

1.4.2. Flexoelectricity

Membrane curvature and transmembrane potential

Materials are made of molecules, and thus on a microscopic level matter is discrete and consists of electric charges. A lipid bilayer is made of lipids with inhomogeneous charge distributions and can thus be thought of as a collection of electric dipoles with dipole moment \vec{p}_j . If the bilayer is flat and symmetric with respect to the interlayer plane (i.e. one layer is the mirror image of the other), these electric dipoles cancel each other out and the net polarisation per unit length is zero. If, on the other hand, the membrane is asymmetric, the dipoles may not exactly cancel each other out and there will be a net electric potential between the layers; this will cause the dipoles to move, which will result in the membrane curving if there is no lipid exchange with the surrounding medium. Similarly, if the membrane is not flat, then the dipoles will be rearranged (figure 1.12), resulting in different surface charge densities on the two layers and thus a nonzero potential between the layers. The latter phenomenon is known as converse flexoelectricity or electrostriction; the former is known as direct flexoelectricity (we will simply refer to it as flexoelec-

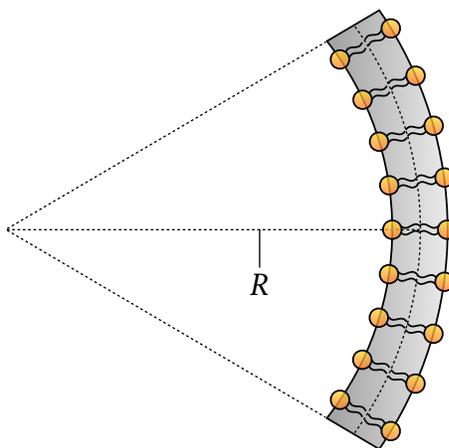


Figure 1.12: Membrane curvature leads to a change in the dipole moments of the membrane's constituent molecules.

tricity here).⁴²

The flexoelectric coefficient

For a membrane with radius of curvature R , the flexoelectric coefficient f is defined according to⁴³

$$P = \frac{f}{R},$$

where P is the magnitude of \vec{P} . Because P has units of charge/distance, f has units of charge. Measurements of f in artificial neutral lipid bilayers⁴⁴ have been of the order of 10^{-20} C; in charged lipid bilayers (and, in particular, in biological membranes),⁴⁵ f can be of the order of 10^{-19} C and sometimes even 10^{-18} C.

Flexoelectric voltage

The voltage across a membrane \mathfrak{M} which has curvature

$$\kappa = \frac{1}{R}$$

and is surrounded on both sides by conductive media is

$$V_m = \frac{1}{\epsilon A} \int_{\mathfrak{M}} P dA = -\frac{f}{\epsilon} \langle \kappa \rangle_{\mathfrak{M}}, \quad (1.4)$$

where ϵ is the membrane's permittivity, A is its area and

$$\langle \kappa \rangle_{\mathfrak{M}} = \frac{1}{A} \int_{\mathfrak{M}} \kappa dA$$

is the average membrane curvature.⁴³

For example, taking $f = 10^{-18}$ C, a cylindrical biological membrane with a radius of curvature of $50 \mu\text{m}$ has a curvature-induced transmembrane voltage of -1 mV.

1.4.3. The Nernst potential and the Goldman-Hodgkin-Katz equation

The chemical potential

The chemical potential μ (see section 1.1.2) of a system is the energy required to add one particle to the system. Given a molar concentration C of a solute in a solution at temperature T , the solute's chemical potential is³⁶

$$\mu = RT \ln(C),$$

where $R = k_B N_A$ is the ideal gas constant, k_B is the Boltzmann constant and N_A is Avogadro's number. In the case of a cell membrane inside of which an electrically neutral

solute has a concentration C_i and outside of which that same solute has a concentration C_o , the chemical potential difference across the membrane is

$$\Delta_\mu = RT \ln \left(\frac{C_i}{C_o} \right).$$

A negative chemical potential difference thus occurs when $C_i < C_o$ and causes some of the solute to travel through the membrane from the cell's exterior to its interior (assuming the membrane is permeable to the solute), whereas a positive chemical potential difference leads to the solute travelling in the opposite direction. In this case, the direction of flow of the solute depends solely on the relative concentrations on opposite sides of the membrane.

The Nernst potential

If the solute is charged, however, a concentration difference across the membrane will lead to a transmembrane potential V_m when the system formed by the membrane and the solutes is in equilibrium. In this case, the internal energy difference across the membrane is (recall equation 1.1)

$$\Delta_U = RT \ln \left(\frac{C_i}{C_o} \right) + qN_A V_m = RT \ln \left(\frac{C_i}{C_o} \right) + \frac{q}{q_e} F V_m,$$

where q is the solute's charge, q_e is the absolute value of the charge of an electron and $F = q_e N_A$ is Faraday's constant. Here, we have assumed that both sides of the membrane are at the same pressure.

From this expression it follows that, for positive univalent ions (i.e. $q = q_e$) and in thermodynamic equilibrium (i.e. when $\Delta_U = 0$),

$$V_m = -\frac{RT}{F} \ln \left(\frac{C_i}{C_o} \right).$$

This is called the Nernst equation after Walther Hermann Nernst's work in 1888.⁴⁶ In this case, V_m is called the Nernst potential of the ion in question and is equal to the electric potential necessary to counteract the chemical potential and thus prevent transport of the ion across the membrane.⁴⁷

The Goldman-Hodgkin-Katz equation

A biological membrane of thickness d has different permeabilities to different ions. A generalisation of the chemical potential difference to several solutes is³⁶

$$\Delta_\mu = RT \ln \left(\frac{\sum_{j,+} \mathcal{P}_j C_{i_j} + \sum_{j,-} \mathcal{P}_j C_{o_j}}{\sum_{j,+} \mathcal{P}_j C_{o_j} + \sum_{j,-} \mathcal{P}_j C_{i_j}} \right),$$

where the subscripts + and – in the sums indicate that the sums are over the positive and negative solutes (respectively), $\mathcal{P}_j = D_j/d$ is the membrane’s permeability coefficient for the j th solute and D_j is the j th solute’s diffusion coefficient (section 1.1.1). The Nernst equation thus becomes

$$V_m = -\frac{RT}{F} \ln \left(\frac{\sum_{j,+} \mathcal{P}_j C_{i_j} + \sum_{j,-} \mathcal{P}_j C_{o_j}}{\sum_{j,+} \mathcal{P}_j C_{o_j} + \sum_{j,-} \mathcal{P}_j C_{i_j}} \right).$$

This is known as the Goldman-Hodgkin-Katz equation; it was proposed by Alan Lloyd Hodgkin and Bernard Katz in 1949⁴⁸ following the work of David E Goldman in 1943.⁴⁹

As an example, the concentrations of potassium (K^+), sodium (Na^+) and chlorine (Cl^-) ions inside and outside of mammalian nerve cells and squid giant axons are presented on the table below (other ions are comparatively unimportant for the purposes of neural communication, as will be seen in section 1.5.2).⁵⁰ Mammalian nerve cells are about 2 times as permeable to Cl^- as to Na^+ and about 50 times as permeable to K^+ as to Cl^- ,³⁶ which leads to a transmembrane equilibrium potential of about -87 mV at $37^\circ C$ or about -82 mV at $20^\circ C$. The squid giant axon’s permeability to Cl^- , on the other hand, is about 0.45 times its permeability to K^+ , which is about 25 times its permeability to Na^+ ; this yields a transmembrane potential between -62 mV and -46 mV (depending on the intracellular Cl^- concentration) at $20^\circ C$. It is clear that flexoelectricity (equation 1.4) contributes very little to the total transmembrane potential in the case of the squid giant axon, which can have a radius of several hundred microns;⁵¹ in the case of mammalian neurons, which are much smaller, the flexoelectric contribution may be comparable to the chemical contribution.

Ion	Mammalian nerve		Squid giant axon	
	C_i (mM)	C_o (mM)	C_i (mM)	C_o (mM)
K^+	139	4	400	20
Na^+	12	145	50	440
Cl^-	4	116	40–150	560

1.5. The Hodgkin-Huxley model

The most well-known description of action potential generation and propagation was proposed in 1952 by Alan Lloyd Hodgkin and Andrew Fielding Huxley.⁵ It is purely electrical in nature; it essentially consists of modelling the axon membrane as a charged capacitor in parallel with a resistor and a voltage source and looking at how the voltage across it changes when an ionic current is present.

1.5.1. The membrane equation

The Hodgkin-Huxley representation of the cell membrane

As mentioned in section 1.4, there is an electric potential difference across the cell membrane. This is given by $V_m = V_i - V_e$, where V_i is the potential inside the cell and V_e is the potential outside it. At rest, V_m is negative for most cells⁴⁷ and takes values between -90 mV and -30 mV for axons.⁵² In addition, cell membranes have a specific capacitance of the order of $1 \mu\text{F}/\text{cm}$.⁵³

Hodgkin and Huxley's model begins by modelling the cell membrane as an electric circuit. A membrane segment is modelled as a capacitor with specific capacitance c_m , and an ion channel protein is modelled as a resistor with resistance R_{cp} (figure 1.13). A whole cell consists of an array of many such units in parallel (figure 1.14); if the cell's area is A , then the total membrane capacitance is $C_m = Ac_m$; if there are N ion channel proteins with resistances $\{R_j\}$, then the membrane resistance is given by

$$\frac{1}{R_m} = \sum_{j=1}^N \frac{1}{R_j}.$$

In the absence of net current (i.e. at rest), the membrane potential is V_r , as is the voltage across the capacitor. The membrane thus has a charge $q = C_m V_r$. This resting potential is a consequence of the curvature and of the differences in the concentrations of different ions on either side of the membrane, as explained in section 1.4.

The membrane equation

Suppose a current I_m flows through the membrane. Unlike Hodgkin and Huxley, we shall take the convention that a positive current flows outwards. This current will change the membrane voltage V_m to something different from V_r .

Let us simplify calculations by writing them in terms of the excess membrane potential (with respect to the resting potential), $\Delta_v = V_m - V_r$. The current through the capacitor

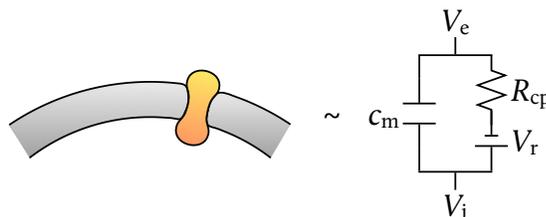


Figure 1.13: Equivalent circuit of a cell membrane segment (grey) with a protein channel (orange) embedded in it.

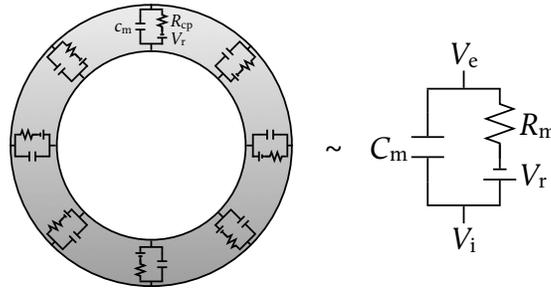


Figure 1.14: Equivalent circuit of the membrane of a whole cell; the equivalence was shown to be true in section 1.2.2.

is

$$I_C = C_m \frac{dV_m}{dt} = C_m \frac{d\Delta V}{dt}.$$

By Kirchhoff's law, $I_m = I_C + I_R$. Thus,⁵

$$C_m \frac{d\Delta V}{dt} + I_R - I_m = 0, \quad (1.5)$$

We will call equation 1.5 the membrane equation. Its solution depends on I_m .

1.5.2. The ionic current

Ion channel proteins

The most important ions for neural communication are potassium (K^+), sodium (Na^+) and chlorine (Cl^-); other ions, such as calcium (Ca^{++}) and Magnesium (Mg^{++}), are also present in neurons and their surroundings, but their concentrations are so low that they do not play important roles.⁵⁰ In its resting state, the cell membrane is permeable to potassium and chlorine but not to sodium.⁵⁴ As the membrane's permeability to potassium and sodium varies, the flow of these ions across the membrane changes, generating an ionic current which is ultimately responsible for the action potential, as we shall see later.

In Hodgkin and Huxley's model, the ion channel proteins are mutually independent, voltage-dependent resistances R_K and R_{Na} ; the membrane lacks chlorine channels, so its resistance to chlorine ion transport is merely the membrane resistance, R_m (figure 1.15). We may thus write $I_R = I_K + I_{Na} + I_{Cl}$. The individual ionic currents are given by

$$I_j = \frac{\Delta V - V_j}{R_j}$$

with $j \in \{K, Na, Cl\}$ and $V_j = V_{N_j} - V_r$, where the V_{N_j} are the potassium, sodium and chlorine ions' Nernst potentials (section 1.4.3).⁵ R_K and R_{Na} depend on both ΔV and t .

Gating particles

Hodgkin and Huxley employed the idea that “gating particles” control the degree to which the channels specific to each ion type are open (i.e. allow ions to pass through). No physical mechanism was given for these particles’ action on the proteins, nor was their physicochemical nature hinted at, but these particles are at the heart of the model.

We begin by defining $R_{j_v} = \min(R_j)$, where $\min(f)$ is the minimum value of a function f over the domain of its argument, which for the R_j is the membrane potential. We also associate probability functions $n, m, h \in [0, 1]$ with the ion channels, where n and m refer to the probability that a potassium and sodium (respectively) channel is open and h refers to the probability that a sodium inactivating particle is not inhibiting the transport of sodium through a channel.

To avoid modelling their data for the time evolution of R_K as a fourth-order polynomial, Hodgkin and Huxley postulated that conduction of potassium ions through a potassium channel protein requires four potassium gating particles to be bound to the channel simultaneously,⁵ so we have

$$I_K = \frac{\Delta_V - V_K}{R_{K_v}} n^4.$$

They further proposed first-order kinetics for the gating transition between the open and closed states of the potassium channel protein:

$$\frac{dn}{dt} = \alpha_n(1 - n) - \beta_n n,$$

where α_n and β_n are the (voltage-dependent) transfer rates of the potassium channel’s gating particles. The solution to the above differential equation is

$$n = \alpha_n \tau_n + \left((\alpha_n \tau_n)|_{\Delta_V=0} - \alpha_n \tau_n \right) e^{-\frac{t}{\tau_n}},$$

where $\tau_n = 1/(\alpha_n + \beta_n)$ is the time constant of the potassium gating particles and $t = 0$ corresponds to the beginning of the potassium-channel activation process. For the squid

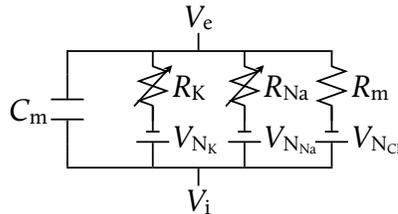


Figure 1.15: Equivalent circuit of an axon or dendrite in the Hodgkin-Huxley model. V_K and V_{Na} are the Nernst potentials for potassium and sodium ions, respectively.

giant axon, which they primarily worked with, Hodgkin and Huxley used⁵²

$$\alpha_n = \frac{V_1 - \Delta V}{V_1 t_1 \left(e^{\frac{V_1 - \Delta V}{V_1}} - 1 \right)},$$

$$\beta_n = \frac{e^{-\frac{\Delta V}{V_2}}}{t_2},$$

where $V_1 = 10$ mV, $V_2 = 80$ mV, $t_1 = 10$ ms and $t_2 = 8$ ms.

The case of the sodium channels is slightly more complicated. In this case, for a similar reason, three sodium gating particles and one sodium deactivating particle were proposed:⁵

$$I_{\text{Na}} = \frac{\Delta V - V_{\text{Na}}}{R_{\text{Na}_v}} m^3 h.$$

As in the case of n , we have

$$\frac{dm}{dt} = \alpha_m(1 - m) - \beta_m m,$$

$$\frac{dh}{dt} = \alpha_h(1 - h) - \beta_h h;$$

the solutions are

$$m = \alpha_m \tau_m + \left((\alpha_m \tau_m)|_{\Delta V=0} - \alpha_m \tau_m \right) e^{-\frac{t}{\tau_m}},$$

$$h = \alpha_h \tau_h + \left((\alpha_h \tau_h)|_{\Delta V=0} - \alpha_h \tau_h \right) e^{-\frac{t}{\tau_h}},$$

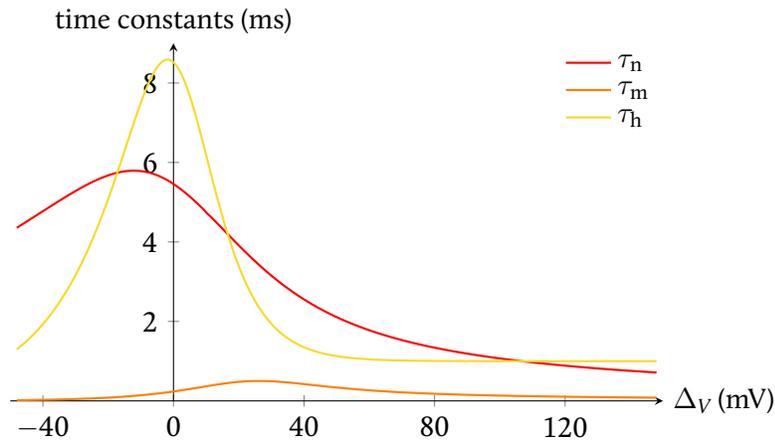


Figure 1.16: The time constants of the gating particles as functions of excess membrane voltage.

where

$$\begin{aligned}\tau_m &= \frac{1}{\alpha_m + \beta_m}, \\ \tau_h &= \frac{1}{\alpha_h + \beta_h}.\end{aligned}$$

Hodgkin and Huxley's data led them to write

$$\begin{aligned}\alpha_m &= \frac{V_3 - \Delta_V}{V_4 t_1 \left(e^{\frac{V_3 - \Delta_V}{V_1}} - 1 \right)}, \\ \beta_m &= \frac{1}{t_3} e^{-\frac{\Delta_V}{V_5}}\end{aligned}$$

for the activating particles and

$$\begin{aligned}\alpha_h &= \frac{1}{t_4} e^{-\frac{\Delta_V}{V_6}}, \\ \beta_h &= \frac{1}{t_5 \left(e^{\frac{V_7 - \Delta_V}{V_1}} + 1 \right)}\end{aligned}$$

for the inactivating particle, where $V_3 = 25$ mV, $V_4 = 1$ mV, $V_5 = 18$ mV, $V_6 = 20$ mV, $V_7 = 30$ mV, $t_3 = 0.25$ ms, $t_4 = 14.3$ ms and $t_5 = 1$ ms.

Comparison of the time constants (figure 1.16) reveals that the maximum of τ_h is nearly 20 times that of τ_m . This means that the sodium deactivating particle takes much longer to relax than the activating particle. On the other hand, the maximum of τ_h is of the same order of magnitude as that of τ_m : the sodium and potassium activating particles have similar relaxation times.

The action potential and the refractory period

With all these considerations, we may write the total ionic current through the membrane (figure 1.17) as

$$I_R = \frac{\Delta_V - V_K}{R_{K_V}} n^4 + \frac{\Delta_V - V_{Na}}{R_{Na_V}} m^3 h + \frac{\Delta_V - V_r}{R_m}.$$

At rest ($\Delta_V = 0$), the total ionic current is zero. If the membrane is hyperpolarised (i.e. the voltage is reduced, so $\Delta_V < 0$; recall that $V_r < 0$), an inward (negative) ionic current passes through the membrane and makes the cell interior more positive, which counteracts this hyperpolarisation (see figure 1.17). If the membrane is slightly depolarised ($0 < \Delta_V < 6.9$ mV), an outward ionic current again counteracts the voltage change.

If the depolarisation is strong enough, however, an inward current contributes to the depolarisation and initiates a positive feedback loop which gives rise to a voltage pulse which can travel along the axon — an action potential (see section 1.5.4). There is thus a threshold voltage below which the axon is not excited.

The loop is eventually interrupted by a delayed increase in h , which leads to an over-compensation (a slight hyperpolarisation with respect to the resting potential), since

$$\frac{dn}{dt} \ll \frac{dm}{dt}$$

(i.e. the potassium channels close much more slowly than the sodium channels and thus there is still some outward potassium current for a time after the sodium channels have closed). This means that the action potential is followed by a short period of time during which the axon is considerably more difficult to excite (the threshold voltage increases to more than 20 times its original value); this period of relative inexcitability is called the refractory period.⁵²

The above expressions for the transfer rates of the sodium and potassium channels were obtained by Hodgkin and Huxley assuming the axon was at a temperature of 6.3 °C. If we instead assume it is at 20 °C, the transfer rates all increase by a factor of about 4.5.⁵ This merely means that n , m and h change 4.5 times more quickly, which in turn means that the whole event lasts 4.5 times less. In other words, the only effect of a decrease in temperature on the action potential is to slow it down.

1.5.3. The membrane current

Equivalent circuit of a small length of axon

Let us now take into account a small length of axon, which includes the cytoplasm inside the membrane. If the cytoplasm and the extracellular medium are homogeneous and the latter is considered to be much more conductive than the former, the axon or dendrite has a relatively simple equivalent circuit (figure 1.18).⁵²

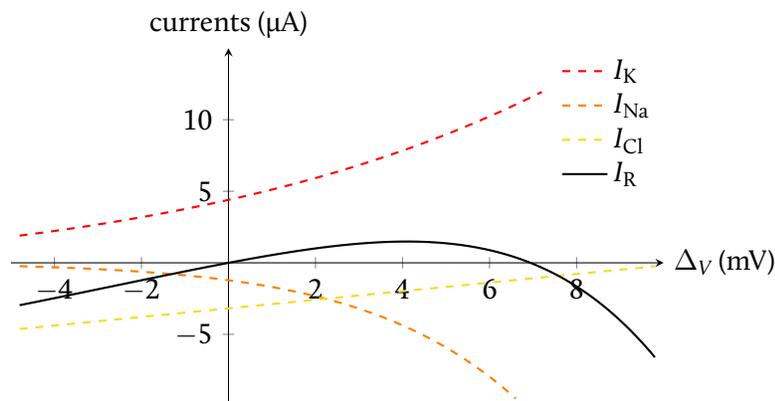


Figure 1.17: The total ionic current and its constituents at $t = 2.5$ ms.

Ohm's law for a small patch of membrane is

$$\frac{\partial V_m}{\partial z} = r_i I_i,$$

where r_i is the cytoplasm's resistance per unit length and I_i is the current flowing through the cytoplasm. The change in the current flowing through the cytoplasm is

$$i_i = \frac{\partial I_i}{\partial z} = \frac{1}{r_i} \frac{\partial^2 V_m}{\partial z^2} = \frac{1}{r_i} \frac{\partial^2 \Delta V}{\partial z^2}.$$

We may think of a small length of axon as a ring (of radius a) formed by many small patches of membrane (with their adjacent cytoplasm). The current density j_m through the membrane is then equal to i_i divided by the axon circumference:

$$j_m = \frac{i_i}{2\pi a}.$$

Similarly, the capacitance and resistance per unit length of the membrane are given by

$$c_m = 2\pi a C_m,$$

$$r_m = \frac{R_m}{2\pi a},$$

where C_m and R_m are the membrane's specific capacitance and resistance and have units of capacitance divided by area and resistance times area, respectively. On the other hand, the cytoplasmic part of a small length of axon is a solid disk. Thus, the resistance per unit length of the cytoplasm is given by

$$r_i = \frac{R_i}{\pi a^2},$$

where R_i is the cytoplasm's specific resistance and has units of resistance times length.

With these considerations, we may write the membrane current density as

$$j_m = \frac{a}{2R_i} \frac{\partial^2 \Delta V}{\partial z^2}.$$

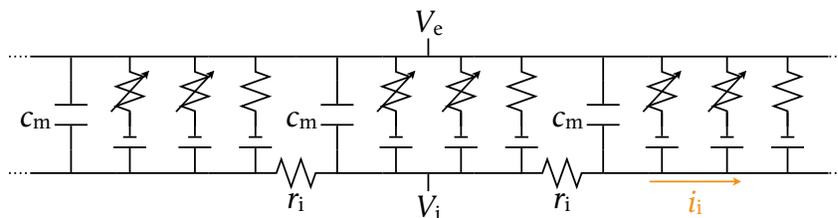


Figure 1.18: Equivalent circuit of a homogeneous membrane and the cytoplasm adjacent to it.

The membrane equation revisited

If we write the membrane equation in terms of the membrane current density and the membrane's specific capacitance and resistance (dividing the latter into the ionic specific resistances R_K , R_{Na} and R_m), we obtain

$$C_m \frac{\partial \Delta_V}{\partial t} + \frac{\Delta_V - V_K}{R_{Kv}} n^4 + \frac{\Delta_V - V_{Na}}{R_{Nav}} m^3 h + \frac{\Delta_V - V_r}{R_m} - j_m = 0.$$

Substituting the expression we have just found for j_m and rearranging terms, this becomes

$$C_m \frac{\partial \Delta_V}{\partial t} + \frac{\Delta_V - V_K}{R_{Kv}} n^4 + \frac{\Delta_V - V_{Na}}{R_{Nav}} m^3 h + \frac{\Delta_V - V_r}{R_m} = \frac{a}{2R_i} \frac{\partial^2 \Delta_V}{\partial z^2}. \quad (1.6)$$

This is Hodgkin and Huxley's membrane equation.⁵

An important consequence of equation 1.6 is that, like all differential equations which are of first order in the time derivative and of second order in the spatial derivative (see section 1.1.1), it involves dissipation. This means that some current "leaks" through the membrane as it travels along the axon or dendrite.⁵ Another consequence is that it does not have wavelike solutions with constant propagation speed.⁵²

1.5.4. The action potential in the Hodgkin-Huxley model

Despite equation 1.6 resembling a diffusion equation, Hodgkin and Huxley postulated that⁵

$$\frac{\partial^2 \Delta_V}{\partial z^2} = \frac{1}{u^2} \frac{\partial^2 \Delta_V}{\partial t^2};$$

in other words, they postulated a wavelike solution propagating at constant speed u . Substituting this into equation 1.6 to turn it into an ordinary differential equation in t and

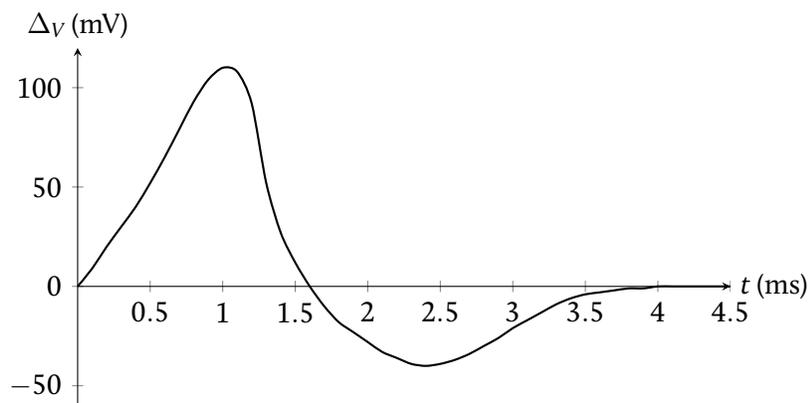


Figure 1.19: The typical shape of an action potential.

simplifying the ionic current, we obtain

$$\frac{a}{2R_i u^2} \frac{\partial^2 \Delta_V}{\partial t^2} - C_m \frac{\partial \Delta_V}{\partial t} - I_R = 0. \quad (1.7)$$

The only term in equation 1.7 which depends on the axon radius is the first one, and R_i is independent of the radius; thus, in order for the equation to hold, u^2 must be proportional to a : the propagation speed of an action potential scales as the square root of the axon radius and is thus an increasing function of a (this is true only for nonmyelinated axons, but a mathematical treatment of myelinated axons is beyond the scope of this work).

Hodgkin and Huxley's numerical solution of equation 1.7 (figure 1.19) accurately reproduces the shapes of the action potentials they observed in the squid giant axon. They obtained $u = 18.8$ m/s, and the experimentally measured propagation speed was $u = 21.2$ m/s, meaning their model yields a value of u which has an error of only 11%.⁵

1.6. The Heimburg-Jackson model

Since the Hodgkin-Huxley model was proposed, studies on individual axons and on entire nerve fibres have revealed many phenomena connected to action potentials: decreases in birefringence and light scattering, changes in axon temperature, and thickening followed by narrowing of the axon, to mention a few. Many of these findings disagree with a purely electrical model based on a dissipation equation (see chapter 2).

Already in 1966 Ling Yun Wei proposed a new model for action potential formation and propagation.⁵⁵ This model viewed the axon exterior, membrane and interior as a PNP junction and suggested that the molecular dipoles in the membrane changed orientation during action potential propagation.⁵⁶ Although it was still electrical in nature, it accommodated some of the optical and thermodynamic phenomena observed during action potential propagation. Two years later, J V Howarth, R D Keynes and J M Ritchie suggested that depolarisation leads to a reduction of entropy.⁵³ In 1971, Yonosuke Kobataki, Ichiji Tasaki and Akira Watanabe proposed a phase transition model for the action potential in another attempt to explain some of the thermodynamic and optical phenomena that had

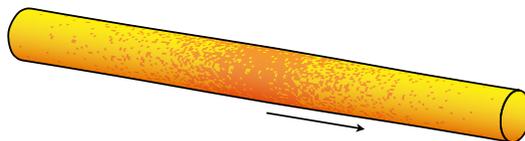


Figure 1.20: The action potential is generated by a propagating region where the membrane lipids are in a different thermodynamic phase (orange) than the lipids in the rest of the membrane (yellow).

been discovered, though their phase transition was one where only the chemical potential and electrostatic properties of the membrane changed.⁵⁷

Developing neurons have been found to mechanically sense the physical properties of their surroundings and grow in one direction or another depending on this information.⁵⁸ It is thus natural to assume neurons can also communicate mechanically. In 2005, Thomas Rainer Heimburg and Andrew Dumont Jackson developed a thermodynamic model for action potentials.⁶ Rather than seeing the action potential as a purely electrical phenomenon, they proposed that neurons communicate via localised phase transitions of the membrane lipids which travel along the membrane in a soliton-like fashion (i.e. without distortion or attenuation); in this model, the voltage change is merely one of a number of concomitant phenomena which are all associated with this phase transition (figure 1.20).

1.6.1. Lipid phase transitions

Lipid phases

Lipid bilayers at low temperature are not exactly solids; rather, they resemble gels. They possess a considerable degree of order, however, so this phase is often called the solid-ordered (SO) phase (figure 1.21, top). This phase is characterised by the lipid molecules having their hydrocarbon tails completely or almost completely expanded. When the bilayer melts, its order decreases, so its high-temperature phase is often called the liquid-disordered (LD) phase (figure 1.21, bottom), and “kinks” appear in the hydrocarbon tails.⁵⁹ Thus, a lipid membrane in the SO phase is thicker than one in the LD phase. Lateral density, however, decreases as temperature increases, as in most materials.

Lipid melting

Membrane lipids at body temperature (or growth temperature in the case of single-celled organisms) are typically about 15 °C above their melting temperature.⁸ As mentioned

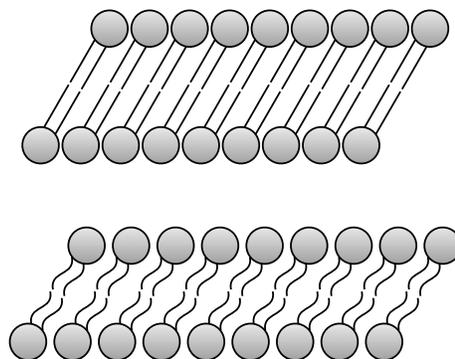


Figure 1.21: Solid-ordered (top) and liquid-disordered (bottom) phases of a lipid bilayer.

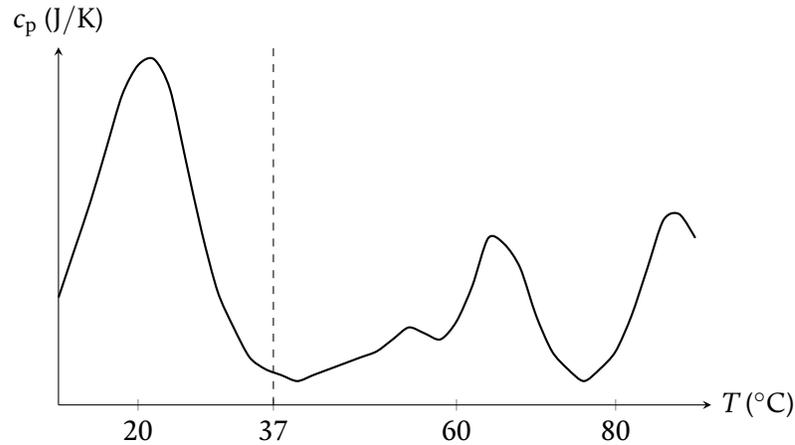


Figure 1.22: The specific heat capacity of escherichia coli membranes purified by centrifugation as measured in a calorimeter by Heimburg and Jackson.⁶ The peaks correspond to the phase transitions of the membrane components — those at temperatures above 37 °C correspond to protein unfolding, whereas the peak near 20 °C corresponds to lipid melting.

in section 1.1, phase transitions are accompanied by maxima in several thermodynamic quantities, including the specific heat capacity c_p (figure 1.22).⁶

As does any phase transition, the transition from the LD phase to the SO phase (as occurs during an action potential according to the Heimburg-Jackson model) has several effects on lipid membranes. Some of these are lateral compression, changes in elasticity and compressibility, curvature changes, decreases in forward- and high-angle-scattered light, changes in birefringence, and heat release. All of these changes have been found to accompany action potentials⁶ and will be discussed in detail in chapter 2.

1.6.2. Compression-wave propagation in membranes

Sound waves

A sound wave is a change of density which propagates in wavelike form through a medium. The equation for sound propagation in one dimension along a 2-dimensional material (such as a hollow cylinder) is⁶

$$\frac{\partial^2 \Delta_\rho}{\partial t^2} = \frac{\partial}{\partial z} \left(\frac{1}{\kappa_a \rho} \frac{\partial \Delta_\rho}{\partial z} \right),$$

where ρ is the (local) area density of the material the sound wave is propagating through, $\Delta_\rho = \rho - \rho_0$ is the propagating density change (with respect to the material's resting density ρ_0) and κ_a is the area compressibility of the material.

If the amplitude of the sound wave is small (i.e. $\Delta_\rho \ll \rho_0$) and κ_a is not a function of

position, this reduces to the wave equation

$$\frac{\partial^2 \Delta_\rho}{\partial t^2} \approx c_0^2 \frac{\partial^2 \Delta_\rho}{\partial z^2},$$

where

$$c(\rho) = \sqrt{\frac{1}{\kappa_a \rho}}$$

is the speed of the sound wave and $c_0 = c(\rho_0)$.

Nonlinearity

For linear materials, c is not a function of ρ . Real materials, however, are nonlinear and have $c = c(\rho) = c(\Delta_\rho)$. For a small-amplitude sound wave (and taking into account that κ_a is, in general, a function of ρ), we have

$$c^2 = \sum_{j=0}^{\infty} \frac{\Delta_\rho^j}{j!} \left. \frac{d^j(c^2)}{d\rho^j} \right|_{\rho=\rho_0} \approx c_0^2 \left(1 + \frac{a_1}{\rho_0} \Delta_\rho + \frac{a_2}{\rho_0^2} \Delta_\rho^2 \right)$$

with

$$a_1 = -1 - \frac{\rho_0}{\kappa_a} \left. \frac{\partial \kappa_a}{\partial \rho} \right|_{\rho=\rho_0},$$

$$a_2 = -a_1 + \frac{\rho_0^2}{\kappa_a^2} \left(\left. \frac{\partial \kappa_a}{\partial \rho} \right|_{\rho=\rho_0} \right)^2 - \frac{\rho_0^2}{2\kappa_a} \left. \frac{\partial^2 \kappa_a}{\partial \rho^2} \right|_{\rho=\rho_0}$$

and $a_1 < 0 < a_2$.⁶⁰

Dispersion

To take into account dispersion, Heimbürg and Jackson added a fourth-order term to the sound equation:⁶

$$\frac{\partial^2 \Delta_\rho}{\partial t^2} = \frac{\partial}{\partial z} \left(c^2 \frac{\partial \Delta_\rho}{\partial z} \right) - h \frac{\partial^4 \Delta_\rho}{\partial z^4}, \quad (1.8)$$

where c is as before and h quantifies the dispersion.

1.6.3. The action potential in the Heimbürg-Jackson model

The solitary-wave equation

We are interested in solutions to the sound equation with some constant propagation speed u , so we assume that Δ_ρ is only a function of $z - ut$. We define $\zeta = z - ut$ and, using the second-order Taylor expansion of c^2 calculated above, rewrite equation 1.8 as

$$u^2 \frac{d^2 \Delta_\rho}{d\zeta^2} = c_0^2 \frac{d}{d\zeta} \left(\left(1 + \frac{a_1}{\rho_0} \Delta_\rho + \frac{a_2}{\rho_0^2} \Delta_\rho^2 \right) \frac{d\Delta_\rho}{d\zeta} \right) - h \frac{d^4 \Delta_\rho}{d\zeta^4}.$$

This equation contains no terms with ζ , indicating that the solution must be symmetric with respect to $\zeta = 0$.⁶⁰ Integrating both sides of the equation with respect to ζ , we obtain

$$u^2 \frac{d\Delta_\rho}{d\zeta} = c_0^2 \left(1 + \frac{a_1}{\rho_0} \Delta_\rho + \frac{a_2}{\rho_0^2} \Delta_\rho^2 \right) \frac{d\Delta_\rho}{d\zeta} - h \frac{d^3\Delta_\rho}{d\zeta^3} + K_1,$$

where K_1 is an integration constant. Integrating again and rearranging terms, this becomes

$$h \frac{d^2\Delta_\rho}{d\zeta^2} = K_2 + (c_0^2 - u^2)\Delta_\rho + \frac{a_1 c_0^2}{2\rho_0} \Delta_\rho^2 + \frac{a_2 c_0^2}{3\rho_0^2} \Delta_\rho^3 + K_1 \zeta,$$

where K_2 is another integration constant. In order to preserve the aforementioned symmetry, K_1 must equal zero. If we now multiply this by the derivative of Δ_ρ with respect to ζ and integrate with respect to ζ once more, we finally arrive at

$$h \left(\frac{d\Delta_\rho}{d\zeta} \right)^2 = -K_3 + 2K_2\Delta_\rho + (c_0^2 - u^2) \Delta_\rho^2 + \frac{a_1 c_0^2}{3\rho_0} \Delta_\rho^3 + \frac{a_2 c_0^2}{6\rho_0^2} \Delta_\rho^4, \quad (1.9)$$

where K_3 is yet another integration constant and we have made use of the fact that

$$\frac{d^2\Delta_\rho}{d\zeta^2} \frac{d\Delta_\rho}{d\zeta} = \frac{1}{2} \frac{d}{d\zeta} \left(\left(\frac{d\Delta_\rho}{d\zeta} \right)^2 \right).$$

As was mentioned earlier, Δ_ρ is a function of ζ but not explicitly of t or z . This means that the shape of Δ_ρ is maintained as the pulse travels along the medium; in other words, it is soliton-like.⁸ The fact that this soliton-like compression wave must be localised (i.e. Δ_ρ approaches zero as ζ approaches infinity) means that $u \leq c_0$.⁶⁰

Equation 1.9 is mathematically identical to the one-dimensional mechanical equation of motion

$$\frac{m}{2} \left(\frac{dx}{dt} \right)^2 + \mathcal{V}(x) = \mathcal{U}$$

with Δ_ρ playing the role of x , ζ playing the role of t , $\mathcal{U} = -K_3$ playing the role of the total energy,

$$\mathcal{V} = -2K_2\Delta_\rho - (c_0^2 - u^2)\Delta_\rho^2 - \frac{a_1 c_0^2}{3\rho_0} \Delta_\rho^3 - \frac{a_2 c_0^2}{6\rho_0^2} \Delta_\rho^4$$

playing the role of the potential energy, and $h = m/2 > 0$. We may therefore apply some of our knowledge from classical mechanics.

Changing the value of K_3 is equivalent to adding a constant to \mathcal{V} , so K_3 is unimportant and is assumed to equal zero in the following. On the other hand, changing K_2 will directly affect the shape of \mathcal{V} , so its value must affect the solution to equation 1.9.

The ‘‘potential energy’’ \mathcal{V} is a fourth-order polynomial of Δ_ρ with a negative fourth-order term (recall that $a_2 > 0$). Therefore, it has two maxima and one minimum. As in classical mechanics, stable solutions for Δ_ρ must lie between the maxima but also require $\mathcal{V} < 0$ and $\Delta_\rho \geq 0$.⁶⁰

Propagation speed of the action potential

For $K_2 = 0$, the roots of \mathcal{V} take on the fairly simple forms

$$\Delta_\rho \in \left\{ 0, -\frac{a_1}{a_2} \rho_0 \mp \frac{\sqrt{a_1^2 c_0^4 - 6a_2 c_0^2 (c_0^2 - u^2)}}{a_2 c_0^2} \rho_0 \right\}$$

(recall that $a_1 < 0$). Obviously, then, $\mathcal{V} < 0$ for all positive Δ_ρ smaller than the smallest of the two nonzero roots of \mathcal{V} — that with the negative square root. This means that, when $K_2 = 0$, there exist stable solutions to equation 1.9 with $0 < \Delta_\rho < \Delta_{\rho\wedge}$ with

$$\Delta_{\rho\wedge} = -\frac{a_1}{a_2} \rho_0 - \frac{\sqrt{a_1^2 c_0^4 - 6a_2 c_0^2 (c_0^2 - u^2)}}{a_2 c_0^2} \rho_0.$$

Because $\Delta_{\rho\wedge}$, which is the amplitude of the sound wave, is the maximum possible value of Δ_ρ , it is the value of Δ_ρ for which the right side of equation 1.9 equals zero. It is worth noting that it is independent of h .

The two nonzero roots of \mathcal{V} become degenerate when

$$a_1^2 c_0^4 - 6a_2 c_0^2 (c_0^2 - u^2) = 0$$

or $u = u_\vee$ with

$$u_\vee = \sqrt{1 - \frac{a_1^2}{6a_2}} c_0.$$

When this happens,

$$\Delta_{\rho\wedge} = -\frac{a_1}{a_2} \rho_0.$$

Because u^2 has a positive sign in the argument of the square root in the expression for the roots of \mathcal{V} , setting $u > u_\vee$ will make the argument of the square root positive and setting $u < u_\vee$ will make it negative; the latter situation results in the two nonzero roots of \mathcal{V} being complex, which means that equation 1.9 has no non-trivial solution, so u_\vee is the lowest possible propagation speed of a solitary compression wave.

The negative sign of the square root in the expression for $\Delta_{\rho\wedge}$ means that as u grows $\Delta_{\rho\wedge}$ becomes smaller. Thus, stronger action potentials travel more slowly. The minimum possible amplitude occurs at $u = c_0$ and the maximum possible amplitude occurs at $u = u_\vee$.

Note that, in this model, the propagation speed does not depend on the axon radius, but rather on the amplitude of the action potential. This is in direct contradiction to the Hodgkin-Huxley model.

The refractory period

Negative values of K_2 or large positive values of K_2 result in no stable solutions to equation 1.9; a small positive value of K_2 leads to periodic (and therefore non-solitary) solutions with a separation determined by K_2 .⁶⁰ Two waves cannot be separated by a ζ smaller than this separation; this is the refractory period.

The action potential

Given a propagation speed u , the general solution to equation 1.9 is⁸ (figure 1.23)

$$\Delta_\rho = - \frac{a_1 \left(1 - \frac{u^2 - u_v^2}{c_0^2 - u_v^2} \right)}{a_2 + a_2 \left(1 + 2 \sqrt{\frac{u^2 - u_v^2}{c_0^2 - u_v^2}} \right) \cosh \left(\frac{c_0}{\sqrt{h}} \sqrt{1 - \frac{u^2}{c_0^2}} \zeta \right)} \rho_0. \quad (1.10)$$

In Heimburg and Jackson's view, the region near the maximum of Δ_ρ is a propagating region in which the membrane lipids have undergone a transition to the SO phase (see figure 1.20).⁶

The amount of dispersion (quantified by h) has no bearing on the soliton amplitude; instead, it controls the soliton width, with larger values of h leading to wider solitons (see equation 1.10). It is clear from this and from the mathematical analysis of section 1.6.2 that dispersion is necessary for the solitons to exist. The propagation speed u similarly controls the soliton width, with smaller values of u leading to wider solitons.

Lateral compression of an asymmetrically charged membrane will necessarily change the membrane potential V_m (see sections 1.1.2, 1.1.3, 1.4.2 and 1.4.3). Thus, the Heimburg-Jackson model is consistent with the observation of changes in the membrane potential associated with neural activity.

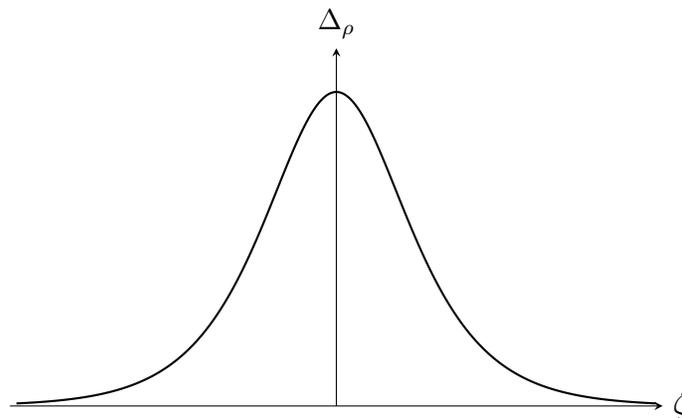


Figure 1.23: The general solution to equation 1.9.

I.7. References

1. Galvani L (1791): *De viribus electricitatis in motu musculari commentarium*, De Bononiensi Scientiarum et Artium Instituto atque Academia Commentarii 7, 363–418
2. Finkelstein G W (2013): *Emil du Bois-Reymond: neuroscience, self, and society in nineteenth-century Germany*, MIT Press, 179
3. Ramón y Cajal S (1894): *Les nouvelles idées sur la structure du système nerveux: chez l'homme et chez les vertébrés*, C Reinwald & Cie, 7–170
4. Bernstein J (1902): *Untersuchungen zur Thermodynamik der bioelektrischen Ströme*, Pflügers Archiv für die gesamte Physiologie 92, 521–562
5. Hodgkin A L & Huxley A F (1952): *A quantitative description of membrane current and its application to conduction and excitation in nerve*, Journal of Physiology 117, 500–544
6. Heimburg T R & Jackson A D (2005): *On soliton propagation in biomembranes and nerves*, Proceedings of the National Academy of Sciences of the United States of America 102, 9790–9795
7. Fick A (1855): *Ueber Diffusion*, Annalen der Physik 170, 59–86
8. Heimburg T R (2007): *Thermal biophysics of membranes*, Wiley-VCH, 9, 42–46, 78–79, 85–95, 312–317
9. Gould H & Tobochnik J (2010): *Statistical and thermal physics*, Princeton University Press, 60
10. Planck M (1945): *Treatise on thermodynamics*, third edition, Dover, 179
11. Linder B (2004): *Thermodynamics and introductory statistical mechanics*, Wiley-Interscience, 9
12. Fermi E (1956): *Thermodynamics*, Dover, 86
13. Klenin V J (1999): *Thermodynamics of systems containing flexible-chain polymers*, Elsevier, 50–52
14. Katzgraber H G (2007): *Phase transitions*, lecture notes, ETH Zürich, 22–38
15. Horowitz P & Hill W (1989): *The art of electronics*, second edition, Cambridge University Press, 3–21
16. Griffiths D J (1999): *Introduction to electrodynamics*, fifth edition, Prentice Hall, 330–331

17. Nahmad-Rohen A (2014): *The effective refractive index of blood as a tool for medical diagnosis*, BSc thesis, 2–6 (available upon request)
18. Born M & Wolf E (1980): *Principles of optics*, sixth edition, Pergamon Press, 665–667
19. Kaplan M W (1977): *Birefringence in biological materials*, Proceedings of SPIE 112, 112–119
20. Bass M (1995): *Handbook of optics, volume II*, second edition, McGraw-Hill, 17.47
21. Nahmad-Rohen A, Contreras-Tello H, Morales-Luna G & García-Valenzuela A (2016): *On the effective refractive index of blood*, Physica Scripta 91, 015503
22. Ponne C T & Bartels P V (1995): *Interaction of electromagnetic energy with biological material — relation to food processing*, Radiation Physics and Chemistry 45, 591–607
23. Kittel C (2005): *Introduction to solid state physics*, eighth edition, John Wiley & Sons, 444–445
24. Cheng J-X & Xie X S (2013): *Coherent Raman scattering microscopy*, CRC Press, 10–20
25. Raman C V (1928): *A new radiation*, Indian Journal of Physics 2, 387–398
26. Smekal A (1923): *Zur Quantentheorie der Dispersion*, Naturwissenschaften 11, 873–875
27. Stokes G G (1852): *On the change of refrangibility of light*, Philosophical Transactions of the Royal Society of London 142, 463–562
28. Nahmad-Rohen A (2015): *The dynamics of COSAN-cell interaction*, MSc thesis, 1–3 (available upon request)
29. Boltzmann L (1877): *Über die Beziehung zwischen dem zweiten Hauptsatze der mechanischen Wärmetheorie und der Wahrscheinlichkeitsrechnung respective den Sätzen über das Wärmegleichgewicht*, Wiener Berichte 76, 373–435
30. Boyd R W (2008): *Nonlinear optics*, third edition, Academic Press, 2–3, 44
31. Prince R C, Frontiera R R & Potma E O (2017): *Stimulated Raman scattering: from bulk to nano*, Chemical Reviews 117, 5070–5094
32. Popov K I, Pegoraro A F, Stolow A & Ramunno L (2012): *Image formation in CARS and SRS: effect of an inhomogeneous nonresonant background medium*, Optics Letters 37, 473
33. Langbein W & Borri P (2014): *Resonant nonlinear optical microscopy*, in *Proceedings of the International School of Physics* 181, 150–152,

34. Singer S J & Nicolson G L (1972): *The fluid mosaic model of the structure of cell membranes*, Science 175, 720–731
35. Franke T, Leirer C T, Wixforth A, Dan N & Schneider M F (2009): *Phase transition and dissipation driven budding in lipid vesicles*, ChemPhysChem 10, 2852–2857
36. Voet D & Voet J G (2011): *Biochemistry*, fourth edition, John Wiley & Sons, 395–411, 744–755, 771–775
37. Eeman M & Deleu M (2010): *From biological membranes to biomimetic model membranes*, Biotechnologie, Agronomie, Société et Environnement 14, 719–736
38. van Meer G, Voelker D R & Feigenson G W (2008): *Membrane lipids: where they are and how they behave*, Nature Reviews Molecular Cell Biology 9, 112–124
39. Simons K & Ikonen E (1997): *Functional rafts in cell membranes*, Nature 397, 569–572
40. Brown D A & London E (1997): *Structure of detergent-resistant membrane domains: does phase separation occur in biological membranes?*, Biochemical and Biophysical Research Communications 240, 1–7
41. Fukuma T, Higgins M J & Jarvis S P (2007): *Direct imaging of lipid-ion network formation under physiological conditions by frequency modulation atomic force microscopy*, Physical Review Letters 98, 106101
42. Vitkova V & Petrov A G (2013): *Lipid bilayers and membranes: material properties*, in *Advances in planar lipid bilayers and liposomes* 17, Elsevier, 89–138
43. Petrov A G (2001): *Flexoelectricity of model and living membranes*, Biochimica et Biophysica Acta 1561, 1–25
44. Petrov A G, Miller B A, Hristova K & Usherwood P N R (1993): *Flexoelectric effects in model and native membranes containing ion channels*, European Biophysics Journal 22, 289–300
45. Sachs F, Brownell W E & Petrov A G (2009): *Membrane electromechanics in biology, with a focus on hearing*, Materials Research Society Bulletin 34, 665
46. Nernst W H (1888): *Zur Kinetik der in Lösung befindlichen Körper: erste Abhandlung: Theorie der Diffusion*, Zeitschrift für Physikalische Chemie 2, 613–637, 634–637
47. Sherman-Gold A (2012): *The axon guide*, third edition, Molecular Devices, 17, 20–21, 111
48. Hodgkin A L & Katz B (1949): *The effect of sodium ions on the electrical activity of the giant axon of the squid*, Journal of Physiology 108, 37–77
49. Goldman D E (1943): *Potential, impedance, and rectification in membranes*, Journal of General Physiology 27, 37–60

50. Lodish H, Berk A, Matsudaira P, Kaiser C A, Krieger M, Scott M P, Zipursky L & Darnell J (2003): *Molecular cell biology*, fifth edition, W H Freeman, 253
51. Terakawa S (1983): *Changes in intracellular pressure in squid giant axons associated with production of action potentials*, Biochemical and Biophysical Research Communications 114, 1006–1010
52. Koch C (1999): *Biophysics of computation: information processing in single neurons*, Oxford University Press, 5–20, 25–32, 144–161
53. Howarth J V, Keynes R D & Ritchie J M (1968): *The origin of the initial heat associated with a single impulse in mammalian non-myelinated nerve fibres*, Journal of Physiology 194, 745–793
54. Hodgkin A F & Huxley A F (1947): *Potassium leakage from an active nerve fibre*, Journal of Physiology 106, 341–367
55. Wei L Y (1996): *A new theory of nerve conduction*, Spectrum 3, 123–127
56. Wei L Y (1971): *Possible origin of action potential and birefringence change in nerve axon*, Bulletin of Mathematical Biology 33, 521–537
57. Kobatake Y, Tasaki I & Watanabe A (1971): *Phase transition in membrane with reference to nerve excitation*, Advances in Biophysics 2, 1–31
58. Koser D E et al (2016): *Mechanosensing is critical for axon growth in the developing brain*, Nature Neuroscience 19, 1592–1598
59. Heimburg T R & Jackson A D (2007): *The thermodynamics of general anesthesia*, Biophysical Journal 92, 3159–3165
60. Villagrán-Vargas E, Lulu A, Hustert R, Gumrich P, Jackson A D & Heimburg T R (2011): *Periodic solutions and refractory periods in the soliton theory for nerves and the locust femoral nerve*, Biophysical Chemistry 153, 159–167

Chapter 2

Current understanding of action potentials and how to increase it

As mentioned at the beginning of chapter 1 and will be repeatedly seen throughout this chapter, neither model is fully consistent with the vast body of experimental results generated over the course of the twentieth century. In what follows, we will review these results, estimate the models' relevant predictions, and compare the experimental and predicted effects.

In order to gauge the full extent of applicability of each model and resolve the occasionally conflicting fragments of evidence, experiments more sophisticated than the ones performed thus far will be necessary. The last part of this chapter, section 2.6, will examine the requirements such experiments will need to satisfy if they are indeed to probe deeper into the mysteries of neural activity.

2.1. Mechanical aspects of action potentials

2.1.1. Axon radius, length and pressure changes

D K Hill: micrometric radius increase in cuttlefish giant axon

One of the earliest experiments in which changes in the radius of axons were recorded was performed in 1950 by D K Hill.¹ Hill extracted cuttlefish giant axons 90–125 μm in radius, placed each axon so its edge was visible as a division between bright (axon edge) and dark (exterior) regions in a narrow slit under a microscope, and measured the intensity of the light passing through the microscope eyepiece (figure 2.1). He excited the axon at 15–20 °C with square voltage waves at 100–200 Hz for 1 min and took light-intensity readings every 15 s before, during and after stimulation; an increase in intensity corresponded to the axon growing and occupying a larger fraction of the field of view. In this way, Hill observed an increase of 50–170 nm in the axon radius, which corresponds to an increase of 0.039%–0.139%. The swelling persisted for 15–75 s after stimulation ceased; this great variability in the duration of the swelling occurred even when multiple experiments were

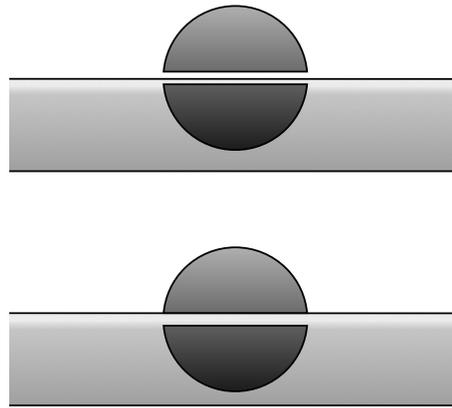


Figure 2.1: Diagram of D K Hill's 1950 experiment. The axon edge is partially visible in a slit under the microscope (top). The edge is bright, so expansion of the axon or movement in the direction in which more of the axon is visible (upwards in the figure) results in an increase in the intensity of the light passing through the microscope eyepiece, while contraction or movement in the opposite direction results in an intensity decrease. The axon edge is brighter than the axon interior, so changing the field of view to be completely within the axon (bottom) results in a decrease of light intensity if the axon expands or moves upwards and an increase of intensity if the opposite occurs.

carried out on the same axon. In some cases the axon radius returned to normal after a few minutes, but in other cases the axon remained swollen. The amount of swelling did not depend on axon size. Stimulation of an inexcitable axon by as much as 20 times the threshold voltage produced no visible radius change, indicating the swelling was related to the action potential even if it did not follow exactly the same time course. Moving the axon so the field of view included the bright axon edge and a darker region farther into the axon caused the light intensity to decrease, rather than increase, when the axon was stimulated, indicating that the change in intensity was due to movement of the axon edge relative to the field of view and not due to any optical changes (but see section 2.2 for further discussion on optical changes in the axon associated with the action potential).

Hill first attempted to explain his findings on the basis of water molecules crossing the membrane along with the ions exchanged during the action potential process. Sodium has a smaller atomic radius than potassium, which means that the bond between a sodium ion and its hydration shell (the water molecules forming a shell around the ion) is stronger than that between a potassium ion and its hydration shell. This in turn means that the hydration number of sodium (the number of water molecules forming a shell around the ion) is greater than that of potassium.² Assuming an ion's hydration shell is capable of following the ion as it crosses the membrane (which seems unlikely in the Hodgkin-Huxley model due to the high specificity of the ion channels), then, water molecules indeed enter the axon during the action potential. Hill assumed the exchange of one potassium ion for one sodium ion caused, on average, a net increase of 3.2 molecules of water in the axon interior. He also claimed that an action potential involves the exchange of about 10^{12}

ions/cm², which corresponds to about 10^{13} water molecules entering the axon. Assuming atmospheric pressure, this is about 2.6×10^{-10} ml/cm², which would cause a radius increase of about 2.6 pm in an axon 100 μ m in radius — more than 42,000 times less than observed on average. Even if one were to adopt Hill's view that the swelling was cumulative (which detailed analysis of his data shows no clear sign of) and that each square pulse generated an action potential, the 6,000–12,000 stimulations each axon received would only amount to 14%–28% of the average radius increase observed by him.

Flexoelectricity (equation 1.4) might appear to be a more likely explanation. However, a voltage change of 100 mV would cause a curvature change of 2×10^6 or 2×10^7 m⁻¹, depending on whether the membrane's flexoelectric coefficient is 10^{-18} C or 10^{-19} C, respectively. This would translate into a radius change of about 20–200 mm, which is 5–6 orders of magnitude too large (it is even larger than the axon itself). Thus, this cannot explain Hill's findings.

As was discussed in section 1.2.1, the charge in a capacitor's plates pulls the plates together with a pressure proportional to said charge (equation 1.3). The discharging of the capacitor relieves that pressure, which is mathematically equivalent to the appearance of an outwards pressure of the same magnitude. An experiment by Susumu Terakawa (which will be discussed shortly) found that the action potential leads to an outwards pressure of about 500 μ Pa; experiments by Kunihiko Iwasa and Ichiji Tasaki (also to be discussed shortly) found a radius increase of the order of 1 nm in the excited axon. If these numbers are to be believed, a pressure of 279 kPa (see section 1.2.1) would cause the axon to increase its radius by about 56 cm, which is again unrealistic. Of course, the 279 kPa calculated previously assumes that the insulator between the capacitor plates is immensely elastic and/or not attached to the plates and that nothing is holding the plates together; a real axon has an actin skeleton holding the membrane in place, which makes it difficult to relate pressure to swelling quantitatively.

The membrane molecules undergoing a phase transition could potentially explain the swelling. As was mentioned in section 1.6.1, the transition of membrane lipids from liquid-disordered to solid-ordered involves the complete or near-complete stretching of the lipids' hydrocarbon tails. These tails normally have 14–24 carbon atoms³ which are separated by a distance of about 0.15 nm and, in unstretched tails, form angles of about 55°. ⁴ This means that an unstretched tail is 0.92–1.63 nm long and a stretched one is 1.99–3.53 nm long. Because a membrane is a lipid bilayer, lipid tail stretching can account for a radius increase of 2.15–3.80 nm. This is two orders of magnitude smaller than the changes observed by Hill, but it is much closer than the other possible explanations. However, experiments with supported lipid bilayers⁵ have observed changes of about 1 nm during a phase transition (see also chapter 5), indicating that the lipid tails do not become fully stretched in the solid-ordered phase.

It is clear that neither theory is well equipped to explain Hill's findings. Flexoelectricity and capacitor pressure predict unrealistically large changes and are thus severely inappropriate, while hydration has a negligible effect and lipid tail stretching can account for only a small percentage of the observed radius change. Motion of the axon could po-

tentially explain these findings, although it would be odd for the axon to move always in the same direction when stimulated and back some time after stimulation ceased; the relative regularity of Hill's observations of brightness increase suggest mechanical changes, rather than motion, are responsible.

Iwasa and Tasaki: nanometric radius increase in squid giant axon

In 1980, in one of their multiple experiments over the course of several decades, Iwasa and Tasaki explored the swelling of the squid giant axon.⁶ They extracted the axon and placed it in a seawater-filled chamber at 23 °C. The axon was excited electrically at one end with an extracellular electrode pair, and its voltage response was recorded at the other end. Gold particles were placed on top of the axon, and the intensity of light reflected from the particles was detected by an optic fibre connected to an amplifier; the detected intensity was used to determine the amount of swelling the axon underwent during the action potential. They measured a radius increase of about 0.5 nm with a profile and time course identical to those of the action potential.

Two years later, they performed a similar experiment⁷ in which they excited a squid giant axon at 19–21 °C with variable series of electrical pulses at 200–333 Hz during several tens of milliseconds. They recorded the action potential both extracellularly (as before) and intracellularly with a 25- μm -radius silver wire introduced longitudinally in the axon. They recorded the mechanical changes in the axon using gold particles. This time, they measured an increase of 1.0–2.5 nm in the axon radius. The axon then became narrower by a similar amount, and this was followed by alternating swelling and narrowing phases with decreasing amplitude, as in a damped oscillation. This oscillation lasted for about 2 ms. No cumulative effects were observed. Chemical stimulation (by reducing the concentration of Ca^{++} ions outside the axon) resulted in a similar phenomenon with an initial amplitude of about 0.7 nm and a duration of about 4 ms. Although this did not follow the profile or time course of the action potential, it is possible that the change in the exterior calcium concentration could change the action potential properties in the Hodgkin-Huxley model or the lipid density in the Heimburg-Jackson model.

The possible explanations for these results are as previously. Tasaki and Iwasa did not give indications of the dimensions of their axons, but squid giant axons tend to be 250–500 μm in radius. The hydration explanation yields a radius increase of about 2.6 pm, which is now a larger percentage of the observed change but still orders of magnitude too small. The predictions of the capacitor plate pressure and flexoelectricity explanations are about an order of magnitude larger than for the somewhat smaller cuttlefish axon, so this explanation is even less viable than before. On the other hand, the lipid stretching explanation, as mentioned before, yields a radius increase of around 1 nm, in line with Iwasa and Tasaki's findings.

Tasaki and Iwasa: nanometric length decrease in crab claw nerve

In another experiment,⁸ Tasaki and Iwasa extracted and desheathed crab claw nerves 35 mm in length and measured the diameter of single axons to be 2–20 μm on an electron

microscope (in a separate experiment). They placed one nerve bundle at a time vertically in a seawater-filled chamber. The lower end of the nerve was held in place by a thread tied to the bottom of the chamber (nerve has a lower density than seawater), and the upper end hung by a thread from a reflective lever. A small weight on the opposite side of the lever caused the nerve to be stretched prior to excitation in order to have some initial tension. They excited the nerve electrically at one end and recorded the voltage change at the other end. The stretching or contraction of the nerve was recorded by measuring the intensity of light reflected from the lever. They recorded a shortening of the nerve by 5–10 nm for about 20 ms, which is longer than the action potential (Tasaki and Iwasa suggested that this delay was due to inertia of the lever).

Volume conservation can be used as an explanation if we assume that the axon density does not change. We must first recall that the action potential lasts about 1 ms and travels at a speed of a few tens of metres per second,⁹ which implies that the active region of the axon (or axons) is a few centimetres long (this is enormous by cellular standards and is a fact apparently unappreciated by many researchers studying action potentials). This, in turn, means that, at the peak of the shortening, the entire length of the nerve studied by Tasaki and Iwasa underwent the structural changes associated with the action potential. Thus, we may model the nerve as a homogeneous cylinder. Assuming a radius of 1–10 μm per axon and a radius increase of 0.5 nm per axon (for the sake of consistency with previous experiments by Tasaki and Iwasa), the required peak shortening for the volume to be conserved is 3.5–35 μm , which is between 350 and 7,000 times the observed shortening.

Another possible explanation is the change in membrane density which occurs in the activated region according to the Heimburg-Jackson model. The observed shortening requires an increase of 0.000014%–0.000029% in the average membrane density. According to the Heimburg-Jackson model, however, the transition from liquid to solid involves a maximum density increase of 20%–25%,¹⁰ which is about a million times larger (although other sources¹¹ give density increases as small as 4%); such a large change in density would have caused the nerve to shorten by about 5 mm.

It is possible that the lever and weight exhibited a large resistance to movement, and Tasaki and Iwasa assumed that the threads they used to fasten the nerve to the bottom of the chamber and to the lever were inelastic, which might not be true, but it is difficult to imagine that such experimental issues could account for such a large discrepancy between the order of magnitude of the observed changes and that of the theoretically predicted changes. On the other hand, the physiological relevance of these findings is questionable because nerves cannot realistically contract more than the size of a molecule without disrupting the neural network they are part of.

Terakawa: biphasic pressure change in squid giant axon

Terakawa's experiment in 1983 to measure pressure changes¹² involved extracting squid giant axons and placing them horizontally on top of electrodes within a seawater-filled chamber. The axon radius was 250–325 μm , and the experiments were performed at a temperature of $(21.0 \pm 0.5)^\circ\text{C}$ and a pH of 8.1. He inserted a buffer-filled glass pipette lon-

gitudinally into the axon at one end such that part of the pipette remained out of the axon; the buffer formed a convex surface at the end of the pipette which was outside the axon and held the intracellular pressure at 40–100 Pa. The axon was stimulated extracellularly by 50- μ s voltage pulses at 10 Hz, and the action potential was recorded extracellularly. The shape of the buffer surface was measured by reflection of light off it; the intensity increased as the surface became flatter. Movement of the surface towards or away from the detector (which was in the direction of the pipette and axon) was not taken into account even though a change in the radius or length of the axon would have caused such a movement of the buffer surface if the cytoplasm and buffer are assumed to be incompressible. Terakawa's results showed an increase in intracellular pressure followed by a decrease; the time course was roughly the same as that of the action potential, although the pressure peak lagged slightly behind the action potential peak (Terakawa attributed this to the fact that the conduction speed of the pressure wave in the pipette was finite; this was tested by using a longer pipette, which indeed resulted in a longer delay). The amplitude of the pressure wave varied between axons but was about 500 μ Pa for most of them; the amplitude of the decrease was at least the same as that of the increase and sometimes reached 7 times that, although there was great variability even in the same axon. Terakawa used sub-threshold voltages to stimulate axons and also tried stimulating anaesthetised axons with super-threshold voltages, both of which produced no pressure changes; this indicates that these changes were related to the action potential. The amount by which the pressure changed was otherwise independent of stimulus strength, further indicating that it is closely linked to the action potential.

We may recall equation 1.2, which relates the change in volume of a liquid which undergoes changes in pressure and temperature to the liquid's dilation coefficient β and isothermal compressibility κ . For water,¹³ which cytoplasm is similar to, $\beta = 2.07 \times 10^{-4} \text{ K}^{-1}$ and $\kappa = 4.59 \times 10^{-10} \text{ Pa}^{-1}$ at 20 °C. In order to produce an increase in intracellular pressure of 500 μ Pa, the action potential would have had to cause the axon volume to decrease by 0.31% assuming the temperature decreases by 15 °C during the action potential (consistently with the Heimburg-Jackson model); neglecting any changes in the axon length, this corresponds to a cytoplasm radius decrease of 0.155% or 400–500 nm. This is only slightly larger than the early observations of Hill but considerably larger than the results of the much more refined experiments by Iwasa and Tasaki.

If we instead assume the cytoplasm remains at constant volume, we may attempt to explain the increase in pressure using the hydration hypothesis of Hill. Because the squid giant axon is larger than the cuttlefish giant axon, we must scale the number of ion channels (and thus the number of ion exchanges during a single action potential). An axon with a radius of 250–325 μ m has 2.50–3.25 times the surface area of a 100- μ m axon; assuming the ion channel density is the same for both axons and Hill's remark about the number of ions exchanged is correct, this amounts to between 6.5×10^{-10} and $8.4 \times 10^{-10} \text{ ml/cm}^2$ of water. Now, an axon with a radius of 250–325 μ m and a surface area of 1 cm^2 has a volume of 12.5–16.3 μ l. Adding the amount of water calculated previously while keeping the volume constant requires increasing the density of water by about 0.000005%, which

translates to an increase in pressure of 2–5 μPa , two orders of magnitude smaller than observed.

Neither of these attempts at an explanation addresses Terakawa's observation that the reduction in intracellular pressure often exceeded the initial pressure increase.

2.1.2. Electron-microscope evidence of membrane thickness changes

Peracchia and Robertson: nanometric membrane thickness increase in crayfish axon

In 1971, Camillo Peracchia and J David Robertson observed axon membranes under an electron microscope during action potential propagation.¹⁴ They exposed crayfish circumesophageal connectives (the circumesophageal connective is the portion of the ventral nerve cord, where the giant axon is located, which lies between the brain and the first ganglion; see section 3.1.3 for more on lobster anatomy, which is essentially identical to crayfish anatomy) and excited them externally with platinum electrodes for 30–1,800 s at room temperature. The excitation pulses were square waves with an amplitude of 1–5 V and a duration of 100–500 μs at 2–60 Hz. They fixed the crayfish by injecting a fixative which had a pH of 7.4 (slightly more acidic than seawater normally is¹⁵) into the sternal artery. They measured the electrical activity of the nerves with external electrodes before, during and after fixation; activity ceased after a certain amount of fixative had entered the crayfish sternal artery. They then extracted the nerves, placed them in the same fixative,

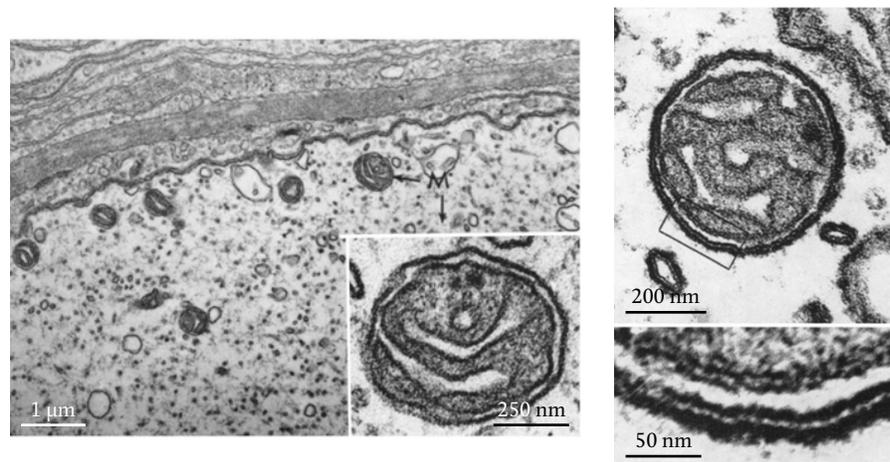


Figure 2.2: Control axon (left) and stimulated axon (right) in Camillo Peracchia and J David Robertson's 1971 experiment. The inset in the left image and the large circular structure in the right image are mitochondria. The mitochondrion in the stimulated axon appears thicker and darker than the one in the control axon. The same is true of the axon membranes near the mitochondria (above the mitochondrion, marked with an arrow, in the left image; above and to the right of the mitochondrion in the right image). Modified from figures 3 (left) and 8 (right) in the original¹⁴ to add scale bars based on the original figure captions.

washed them with saltwater (also at a pH of 7.4) and refixed them with a different fixative (also with the same pH). They finally cut 50-nm-thick sections of nerve and observed them under an electron microscope. They saw an increase in thickness to 12–15 nm (indicating that the thickness increased by 8–11 nm) in not only the axon membrane, but also the endoplasmic reticulum and mitochondrion membranes (figure 2.2). The membrane thickness increase only occurred in axons which were fixed during activity; axons fixed while at rest showed no such changes.

That several different membranes in the axon underwent the same changes during action potential propagation indicates that the physical phenomenon underlying the action potential is one that affects lipid bilayers universally rather than affecting different membrane types in different ways; so does the fact that the changes in giant axons were of the same magnitude as those in smaller axons.

While recent experiments would seem to indicate that mitochondria and endoplasmic reticula have voltage-gated potassium channels in their membranes,^{16,17} they are thought to regulate intracellular voltage rather than playing an active role in action potential generation and propagation;¹⁸ in addition, said experiments have been electrical in nature, so their results can be explained by the Heimburg-Jackson model without invoking ion channels, just as Hodgkin and Huxley's results can be explained without ion channels. Even if we were to assume the ion channels in the mitochondria and endoplasmic reticula play an active role in the action potential, the exchange of ions through their membranes would compensate the cell-membrane exchange at least partially, reducing the membrane voltage change and counteracting the action potential.

Conversely, the Heimburg-Jackson model—whose results are more general and thus applicable, in principle, to any lipid membrane—is better equipped to explain Peracchia and Robertson's findings. The model's prediction of a membrane thickness increase of the order of 1 nm does fall somewhat short of these findings. It is possible that more than one physical mechanism is at work during the action potential, although our previous calculations indicate that flexoelectricity and capacitor pressure cannot explain the observed membrane swelling and that hydration produces a negligible increase in membrane thickness; the fixation agent could have somehow contributed to the thickness change if it can somehow enter the membrane only during activity, but a quantitative treatment of this possibility is beyond the scope of this work.

2.1.3. Action potential generation by mechanical stimuli

Julian and Goldman: action potential generation by pressure increase in lobster giant axon and frog sciatic nerve

It is known that mechanical stimuli can generate action potentials. One experiment showing this was performed by Fred J Julian and David E Goldman in 1962.¹⁹ In it, segments of lobster giant axons with radii of about 50 μm and lengths of 3–4 cm were extracted and placed in a seawater-filled lucite chamber at room temperature. The axons were excited at one end either by rectangular voltage pulses smaller than 2 mV at 200 Hz

or by pressing a stylus tip 500 μm in radius against the central section of the axon. In the case of the mechanical excitation, the stylus was placed such that it made contact with the resting axon but did not press it; the stylus was then lowered by less than 15 μm at speeds of up to 50 $\mu\text{m}/\text{ms}$ and raised again at most 10 ms later at the same speed. Frog sciatic nerves, each with a radius of about 5 μm and containing 5–6 axons, received the same treatment, but the stylus displacement was at most 5 μm . The electrical response was recorded by electrodes placed on the central portion of the axon. Julian and Goldman observed an action-potential-like voltage response caused by the mechanical stimulation. The stylus displacement required to produce the stimulation was 2–5 μm for frog nerve and 10–15 μm for lobster giant axon, indicating that there was a threshold below which the axons were not excited, as in the electrical case. The speed of the depolarisation phase of the response was proportional to the stylus speed, and extremely slow stylus movement did not produce a response. In addition, sub-threshold electrical stimulation initiated at the peak of sub-threshold mechanical stimulation did elicit a response. In the case of mechanical excitation of the lobster giant axon, repolarisation and axon-shape recovery occurred several seconds after the depolarisation phase, so the time course of the changes in the axon was different for each type of stimulation. In addition, there was no post-action-potential hyperpolarisation.

Julian and Goldman concluded that the membrane resistance decreased due to the mechanical stimuli, a conclusion supported by their observation that removing all the sodium in the seawater in which the axon was bathed greatly reduced the amplitude of the depolarisation caused by the mechanical stimuli. However, they did not offer a possible mechanism for this decrease in resistance.

A more formal case can be made for a pressure-induced phase transition. A circular surface 500 μm in radius pressed 10–15 μm down against a lobster axon 50 μm in radius and 3–4 cm in length reduces the axon's volume by 0.13%–0.31%. Using equation 1.2 and a cytoplasm isothermal compressibility of $4.59 \times 10^{-10} \text{ Pa}^{-1}$, this corresponds to a pressure increase of 2.84–6.84 MPa. If the same pressure is felt by the membrane, this is about one order of magnitude higher than the calculation made in section 1.2.1. In the case of a frog sciatic nerve 5 μm in radius, the volume decreases by 0.36%–1.67% (assuming a length of 3–4 cm), which corresponds to a pressure increase of 7.77–36.62 MPa. This is between one and two orders of magnitude higher than the calculation of section 1.2.1. In both cases, the pressure is 10–11 orders of magnitude larger than the action-potential-produced pressure changes observed by Terakawa, but this assumes that the membrane and the cytoplasm feel the same pressure changes. It is also unclear why pressures 10 times smaller did not produce action potentials in Julian and Goldman's experiment.

Flexoelectricity might be initially thought of as another possible explanation, but a closer inspection of equation 1.4 reveals that it actually acts *against* depolarisation. When the axon is compressed by the stylus, the average curvature in the direction of the axon's axis remains zero because there are both regions of positive curvature and regions of negative curvature; conversely, the curvature in the lateral direction increases, which decreases the membrane voltage (i.e. it hyperpolarises the membrane). A quick estimate of the order

of magnitude of this hyperpolarisation can be obtained by assuming the axon becomes elliptical. In the case of the lobster giant axon, the ellipse's major semiaxis is 50 μm and its minor semiaxis is 42.5–45 μm , which results in an average curvature of 21.1–21.8 mm^{-1} . The curvature of the resting (circular) axon is 20.0 mm^{-1} , meaning the curvature change is 1.1–1.8 mm^{-1} . This corresponds to a decrease in membrane potential of 0.24–2.46 mV. In the case of the frog nerve, the resting curvature is 200 mm^{-1} and the compressed nerve's average curvature is 226–308 mm^{-1} . This translates into a curvature change of 26–108 mm^{-1} and a voltage decrease of 2.55–30.48 mV, resulting in hyperpolarisation of the membrane.

2.2. Optical aspects of action potentials

2.2.1. Membrane opacity changes

Solomon and Tobias: opacity increase in squid giant axon

In 1960, S Solomon and Julian M Tobias published their results on axon opacity changes concomitant with action potentials.²⁰ They extracted a giant axon from a squid, placed it on a slot in a slide and covered it with glass, creating what they called a “moist chamber”. They then inserted one end of a quartz needle with a radius of 40–50 μm into the axon and shone light through the other end of the needle, which consisted of a quartz rod 5 mm in diameter with the exterior painted black to prevent additional ambient light from entering the axon. They detected the light transmitted through the cytoplasm and membrane of the axon with a photomultiplier tube. They then electrically stimulated the axon, and recorded the induced action potentials, extracellularly using electrodes upstream from the point where the needle entered the axon; the stimulating signal was a train of 200- μs voltage pulses at 175 Hz lasting for 80–150 s; the stimulating voltage was 250–740 mV. They observed a decrease of 2%–5% in the intensity of the detected light in 72% of cases, a small increase of light in 1% of cases and no change in the remaining cases when an action potential was detected. There was a delay of 2–20 s from the onset of stimulation to the change in the optical signal, and the optical changes outlasted the stimulation by an unspecified amount of time which appears to have been several seconds. In 94% of the cases in which the intensity of the light passing through the axon decreased after stimulation began, the intensity eventually recovered after stimulation ceased. Sub-threshold stimulation produced no changes in the axon's optical transmission, indicating that the observed changes were associated with activity. The same conclusions can be drawn from their measurements on non-conductive axons, which also failed to measure a change in the transmitted light.

Several possible explanations were offered by Solomon and Tobias: a photochemical reaction causing the axon to become more opaque, an increase in the transmittivity of the cytoplasm resulting in less light being scattered sideways and reaching (and subsequently passing through) the membrane, an increase in the extinction of light by the membrane, a change in the refractive index of the cytoplasm leading to a change in the directional

distribution of refracted light (resulting in a similar effect as in the case of the change in cytoplasm transmittivity) and a decrease in the eccentricity of the axon leading to light being refracted more obliquely and less light reaching the photomultiplier.

A change in the shape of the axon or the properties of the membrane seems unlikely, as it would have had a time course similar to that of the action potentials and there would have been little or no delay; there certainly would not have been a delay of several seconds, as this would mean that the first few hundred or thousand stimuli did not cause an action potential while the remainder did, which seems unlikely. While the rest of the proposed explanations seems reasonable and any of those effects could have occurred, whichever effect (or combination thereof) occurred does not appear to be related to the action potential, as the time course it followed was very different from that of the action potentials and there was a considerable delay in the appearance of this effect. Unfortunately, Solomon and Tobias did not provide information related to the delay between the onset of the increased opacity and the start of illumination, so, while we know that there was great variability in the delay between the start of electrical stimulation and the opacity changes, we cannot know the delay between a possible photochemical reaction and the observed opacity changes.

It is clear that the observed changes cannot be the result of an increase in extinction by the axon membrane due to the action potentials making the membrane, and thus the distance the light must travel before exiting the axon, thicker; if this were the case, they would have had a time profile identical to that of the action potentials, and all excitable axons would have displayed them.

It is difficult, then, to imagine what could be causing these changes. One possibility, given that non-conductive axons did not exhibit changes in opacity, is that structural changes in the axon are indeed to blame but they are not a direct result of an action potential; rather, the axon could have become exhausted after extended stimulation, leading to said changes. This would explain the difference in time courses of the optical effects and the action potentials, the variability in the delay of the onset of the optical effects, and the delays themselves of the beginning and end of the optical effects.

2.2.2. Membrane birefringence changes

The cell membrane exhibits a small amount of birefringence²⁰⁻²² due to the orientation of its lipid molecules. Several researchers have observed changes in the birefringence of the axon during neural activity. Here we discuss their experiments and results. The Hodgkin-Huxley model cannot account for any of the birefringence-related observations, while the Heimburg-Jackson model, as will be seen, can explain them qualitatively; however, without knowing the exact contributions to birefringence made by the cytoplasm and both phases of the membrane, which would require knowing their exact composition and making measurements which are, unfortunately, unavailable, it is impossible to give a quantitative prediction of the birefringence change using the model.

Shaw and Tobias: birefringence decrease in squid giant axon

In 1951, Samuel N Shaw and Julian M Tobias extracted squid giant axons and mounted them in an oil-filled plastic chamber with seawater-soaked electrodes on the bottom.²³ They applied 15 V extracellularly through these electrodes; they did not specify whether this was a continuous voltage (i.e. DC) or not, so we assume it was. They recorded the action potentials extracellularly, and they also measured changes in the birefringence of the nerve by using a compensator crystal (figure 2.3). They found that the birefringence Δ_n (given by the difference between the extraordinary and ordinary refractive indices, as explained in section 1.3.1) decreased from its resting value of 1.64×10^{-4} ; the magnitude of the decrease was 4.6×10^{-6} . They did not continue measuring after the stimulation ceased, so it is unknown whether or not the birefringence change they observed was reversible.

The authors did not offer an explanation for their findings. However, in 1973, J M Ritchie²² proposed a possible explanation for a later experiment by Lawrence B Cohen and his colleagues (which will be discussed below). Ritchie suggested that the cellular activity caused changes in the orientation of the membrane lipids, which in turn caused the observed birefringence changes. This explanation is consistent with the Heimburg-Jackson model given what we now know about lipid bilayers (see section 3.1.1): a phase transition to the solid-ordered phase would have straightened the lipid tails and decreased the angular deviation of the lipids from a radial configuration (lipids in a bilayer are not perpendicular to the bilayer plane, but rather have a small slant), changing the birefringence. The change in birefringence which Shaw and Tobias detected is only about 3% of the resting birefringence; however, most of this birefringence probably comes from the cytoplasm, which, although it contributes much less birefringence per unit length than the membrane, is much more abundant than the membrane (for a squid giant axon, light passing through the centre of the axon travels through approximately 100,000 more cyto-

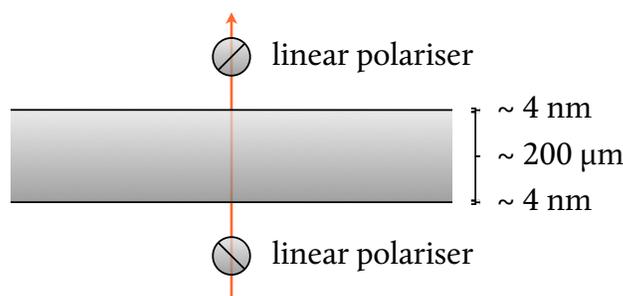


Figure 2.3: Diagram showing the geometry of Samuel N Shaw and Julian M Tobias's 1951 experiment. The horizontal grey rectangle is the axon, and the orange arrow indicates the direction of propagation of the light. The membrane constitutes an extremely small fraction of the distance travelled by the light within the axon.

plasm than membrane, as shown in figure 2.3), so the change in the birefringence of the membrane itself must have been much greater than 3%.

Tasaki et al: birefringence decrease in crab and lobster leg nerve and squid fin nerve

Ichiji Tasaki's experiments were not limited to mechanical findings. In 1968, he and his collaborators studied the optical effects of the action potential.²⁴ They extracted crab and lobster leg nerves and squid fin nerves. The nerves were placed in an acrylic chamber filled with seawater at 19 °C and stimulated extracellularly with a platinum electrode. A second platinum electrode recorded the action potentials, also extracellularly. The researchers used different filters to make the light incident on the nerves "monochromatic" (they did not specify a bandwidth or even the wavelengths they used) and measured changes in the nerves' birefringence by placing the chamber between crossed polarisers in such a way that the nerve was at a 45-° angle with respect to the polarisers' axes. They saw a decrease in the intensity of light by a factor of about 10^{-3} in crab and lobster nerve and 10^{-6} in squid nerve when the nerve was stimulated. Because the birefringence and its change are very small, we may approximate the amplitude of the detected electromagnetic field as a linear function of the birefringence. A change in the detected intensity by a factor of 10^{-3} translates into a change in the amplitude by a factor of about 0.03, or 3%, which implies a change in birefringence by about the same amount. This is the same change observed by Shaw and Tobias in squid giant axon. Strangely, however, the birefringence decrease was 0.1% in the squid nerve.

Tasaki's group reported that there was no significant dependence of the birefringence change on the wavelength of the light and that the time course of the decrease was very similar to that of the action potential. Both of these findings are consistent with what we would expect from a change in the membrane lipids' orientation and straightness and thus with the Heimburg-Jackson model of the action potential.

Cohen et al: birefringence decrease in crab walking nerve and squid giant axon

In the same year, Lawrence B Cohen, R D Keynes and their group initiated a long series of experiments which tested the birefringence of and scattering by axons and nerves.

In one such experiment, they extracted crab walking nerves and squid giant axons and placed them in seawater at 4.5 °C.²⁵ They illuminated a length of 1–3 mm of axon or nerve with white light and measured the birefringence in the same way as Tasaki and his group. With the crab nerve, they observed a decrease in intensity by a factor of $(1-9) \times 10^{-4}$, which implies a decrease in birefringence of 1%–3% (calculated as in the case of Tasaki et al's experiment), consistent with previous findings; however, this was over the course of several hundred milliseconds. With the squid giant axon; the birefringence decrease was again smaller, merely 0.05%–0.15%, although the time course was the same as the action potential's. These figures are all consistent with Tasaki et al's findings. Further tests revealed that the birefringence has a radial optic axis.

Two years later, Cohen and his group conducted a similar experiment²⁶ (carried out at several different temperatures between 5 °C and 22 °C, and with the added difference

that the excitation was intracellular) where they saw a decrease in birefringence of about 0.15% in both squid giant axon and crab nerve which followed the same time course as the action potential, including a slight overshoot during the return phase. There was no significant dependence on temperature or wavelength.

It is difficult to explain the former study's findings in crab nerve, given the great difference between the time courses of the birefringence change and the action potential. Perhaps Cohen et al saw the compound effect of the many axons in the nerve, and the effect of any changes caused by the action potential on the nerve sheath are unknown, so we cannot offer any insight into whether it contributed part of the observed change.

Conversely, their observations with squid giant axon are the same, and thus can be explained in the same way, as Tasaki et al's findings. That the membrane is solely responsible for the changes is made more apparent by their finding that the birefringence effect remained essentially unchanged when the cytoplasm was replaced with buffer solution and by the radial direction of the axon's optic axis; that the changes were caused by the action potential is shown by their binary nature (increasing the stimulation strength above the excitation threshold did not increase the change in birefringence) and by their cessation when nerve activity was blocked with tetrodotoxin (a powerful anaesthetic) or by changing the ions present in the seawater.

In addition to conformational changes, the authors suggested the changes might be caused by attraction of the two leaflets of the membrane (as between the plates of a capacitor). However, this would require a compression of the membrane, which is in direct opposition to the observations reported in section 2.1.

Tasaki and Iwasa: birefringence decrease in squid giant axon

In their 1982 study,⁷ Tasaki and Iwasa found a decrease in the intensity of the light transmitted by the squid giant axon and detected using the same method as before (placing the axon at 45 ° between crossed polarisers) by a factor of 4×10^{-5} , which constitutes a birefringence change of 0.63%. This change again had the same time course as the action potential. The explanation is as for the other reported experiments.

2.2.3. Membrane scattering changes

Several of the experiments which examined the birefringence of an axon or a nerve also looked at changes in the light scattered by it. As before, these results imply important conformational changes in the axon, which suggests that a phase transition might be responsible. Also as in the case of the changes in birefringence, it is difficult to give a quantitative explanation of the observations without knowing the scattering-amplitude matrix of an axon, but attempts at explaining the results qualitatively have been made.

Tasaki et al: side-scattering increase in crab and lobster leg nerve and squid fin nerve

Tasaki et al's 1968 experiment²⁴ also saw changes in the light scattered at 90 ° (in the direction perpendicular to both the axon axis and the direction of propagation of the light)

for different wavelengths in the near-ultraviolet and visible spectra. The intensity of the scattered light increased by a factor of $(0.8-5.7) \times 10^{-5}$ (crab), $(1.0-7.6) \times 10^{-5}$ (lobster) and $(0.9-5.8) \times 10^{-5}$ (squid). These changes did not significantly depend on the wavelength or polarisation of the light.

As with the birefringence, changes in the straightness and angle of the membrane lipids seems to be the most likely reason for these changes.

Cohen et al: biphasic angle-dependent scattering change in crab walking nerve and squid giant axon

Cohen et al's study of the same year²⁵ also looked at scattering changes. Like Tasaki and his colleagues, they saw an increase by a factor of $(2-9) \times 10^{-5}$ in the case of crab nerve. However, this change occurred over several hundred milliseconds, like the birefringence change. Therefore, we may attempt to explain this in the same way; it is possible that what Cohen et al saw was the compound effect of many action potentials occurring out of phase. They also detected a decrease in the scattering at 30° and 45° , but they failed to specify the amount by which the scattering at these angles decreased. With squid giant axon, on the other hand, they saw a decrease in scattering by a factor of about 3.8×10^{-6} which followed the time course of the action potential (with a delay of 0.25 ms) but overshoot and did not return to the baseline over extended periods. This change occurred only at 45° ; there was no observable change at 90° , in contrast to Tasaki's results. In all cases, changing the ions in the seawater or adding tetrodotoxin to block the action potentials made these scattering changes disappear, indicating that they were related to cell activity.

A later study by Cohen et al focussed on crab nerves, which were placed in chambers divided into three sections.²⁷ The central section was filled with seawater, and the other two sections were filled with oil. The nerves were excited extracellularly with electrodes at 10 Hz, and action potential detection was also extracellular; excitation and recording took place in the oil-filled sections. Light scattered by the central section was detected in a 50° arc centred at 35° or 90° . Cohen and his colleagues found that the light scattered at 90° decreased by a factor of $(3.0 \pm 0.1) \times 10^{-5}$, contrary to their previous results and to Tasaki et al's experiment, but also that this decrease was occasionally preceded by a small increase (by a factor of less than 2×10^{-6}). This was the result of integrating the scattered light over a relatively long period (about 300 ms); when they looked at the effect with increased temporal resolution, they found that the scattering increased by a factor of $(4.9 \pm 2.1) \times 10^{-5}$ before decreasing and overshooting to produce a net decrease of about the same magnitude (to within a factor of 2) as detected earlier. The authors concluded that in the integrated measurements they were seeing only the longer-term effect of this binary change. At 35° , the sign of the changes was reversed: a short-lived decrease in scattering was followed by a net increase. However, the authors noted that, due to the density of molecules in the axon (and particularly in the membrane), it is likely that multiple scattering took place, which casts uncertainty on the angles. As before, there was no wavelength dependence.

The authors suggested that an increase in the size of the nerve could be the cause of

the decrease in scattering, basing their guess on their observation that the scattering of light by erythrocytes decreased as the erythrocytes' size increased. While the size of the axons does increase slightly during action potentials, it is more likely that the scattering changes are due to conformational changes in the membrane lipids. For objects larger than the wavelength of light, such as axons in the case of visible light, the amplitude of the scattered field is proportional to the volume;²⁸ this means that an increase in radius by a factor between 10^{-5} (for giant axons some 100 μm in radius) and 10^{-3} (for axons only 1 μm in radius) would produce an increase in the scattered amplitude by a factor of $0.02 - 0.10$, which translates into an increase in the scattered intensity by a factor of $4.6 \times 10^{-4} - 1.0 \times 10^{-2}$. Even the lowest of these values is an order of magnitude higher than the increase observed by Cohen et al.

A final study by Cohen's group looked at squid giant axons only.²⁹ The axons were perfused with buffer solution and kept at $(12 \pm 1)^\circ\text{C}$, and a 4-mm length of axon was illuminated. The light scattered at several different angles between 10° and 120° was detected. The scattered intensity saw a decrease at angles smaller than 45° by a factor of 7.5×10^{-6} , no change at 45° and an increase at angles larger than 45° by a similar factor (1.94×10^{-5}). These changes were also biphasic, following the action potential's time course closely; interestingly, however, they returned to the baseline about 0.25 ms before the voltage did (and then overshoot, as with the crab nerves in the group's previous experiment). The overshoot at small angles had the same magnitude as the initial change, while the overshoot at large angles was small and eventually returned to the baseline over a long period. Increasing the temperature of the axons decreased the magnitude of the scattering change at small angles but not at high angles.

As with the crab nerves, the biphasic nature of these changes strongly suggests that they are ultimately caused by the action potential. The slow return to the baseline after overshooting seen at large angles is more difficult to explain; it may be the result of nerve fatigue due to repeated stimulation or of some long-term change in the cytoplasm.

2.3. Thermodynamic aspects of action potentials

2.3.1. Heat release

A V Hill: heat release in frog sciatic nerve

As early as 1916, Archibald Vivian Hill observed heat release during activity.³⁰ He extracted about 80 mg of frog sciatic nerve. Stimulation of the nerve at 140 Hz and recording of the response were extracellular. At the same time, the heat released by the nerve was measured by a thermopile, though it is unclear exactly where the thermopile was placed. "Immediately" after stimulation began, Hill detected a release of 24 $\mu\text{cal/g}$, which he calculated to be 7–10 ncal/g per action potential. Considering the nerve mass he used, this is 560–800 pcal per action potential. Unfortunately, Hill did not mention whether the heat release was reversible or not. He did assume it was cumulative, which suggests he believed it was irreversible. He believed the heat was necessarily linked to the action potentials

because nerve warming by the stimulating current would have been delayed rather than immediate. Indeed, he did detect such a delayed warming in addition to the immediate heat release. Unfortunately, it is unknown whether his measurements had the temporal sensitivity required for him to observe whether the time course of the thermal changes followed that of the action potentials.

Since the Hodgkin-Huxley model is purely electrical, this heat release must be explained in terms of dissipation by the circuit resistance. The heat dissipated by a resistor of resistance R is $Q = RI^2 = VI$, where I is the current flowing through the resistor and V is the voltage across it. We know that $V \approx -70$ mV; we must thus determine I .

The frog sciatic nerve is myelinated and, like all myelinated nerves, has small sections (only 1–2 μm in length) not covered by myelin, called nodes of Ranvier; the density of ion channel proteins in the nodes is much higher than in the internodal regions, but the sheer size difference between the two region types means that the total number of internodal channels is greater than the total number of nodal channels; for frog sciatic nerve the number of internodal channels has been found to be approximately 40 times the number of nodal channels.³¹ The internodal distance in frog sciatic nerve is about 2.2 mm on average, and the total current at a node has been measured as -2.15 nA.³² Since the nodes are tiny compared to the internodal distance, we may ignore their size and make the approximation that there are 41 times the number of nodal channels in every 2.2-mm stretch of nerve. The current through a single ion channel³³ is between -12 pA and -10 pA, so we calculate the total number of channel proteins per 2.2 mm of nerve as 7,346–8,815.

Now, a frog sciatic nerve has 1,000–100,000 axons which are 1–6 μm in radius.³⁴ The density of a cell membrane^{35,36} is 1.04–1.21 g/ml, while the cytoplasm density is approximately the density of water because cytoplasm is mostly water. Finally, the ratio between the inner and outer radii of the myelin sheath is 0.8–0.9 for most myelinated nerves.³⁷ Using these values, approximating the axons and nerve as tubes and assuming optimal packing of the axons in the space surrounded by the myelin (which means that the fraction of space occupied by the axons is $\pi/2\sqrt{3} \approx 0.9069$), we can calculate the volume fractions of all the components and, considering myelin is also made of stacked lipid bilayers and assuming the interaxonal material has the same density as lipid, estimate the radius and average density of the nerve. With this, we can calculate the total length of the nerves studied by Hill.

We can use these calculations to estimate that the total number of ion channel proteins in a nerve was between 1.4×10^4 and 8.6×10^7 . Using this and the current passing through a single channel, we finally obtain a figure for the total heat dissipated by the nerve: between 2.80×10^{-18} cal and 1.44×10^{-14} cal. It is clear that this cannot be the explanation.

In the Heimburg-Jackson model, most of the heat comes from a phase transition (and lowering of the axon temperature) and is fully reversible; a small amount is dissipated by the ion channels. The heat released in the cooling of the axon membrane is $Q_c = mC_p\Delta_T$, where m is the mass of the axons, C_p is the specific heat of the axon membrane and $\Delta_T \approx -15$ °C is the temperature change. To calculate the latent heat Q_l of the liquid-

gel phase transition, we use the Clausius-Clapeyron relation,

$$\frac{\partial P}{\partial T} = \frac{\Delta_S}{\Delta_V},$$

where Δ_S is the entropy change in the phase transition and Δ_V is the volume change of the membrane, and the equation relating entropy and heat,

$$Q_1 = T\Delta_S,$$

to obtain³⁸

$$Q_1 = \frac{\partial P}{\partial T} T\Delta_V.$$

The value of $\partial P/\partial T$ for a lipid membrane has been found to be 46 kbar/K, and its specific heat is approximately 25 cal/kg K close to the gel-liquid phase transition.³⁸ Hill did not specify the temperature at which he conducted his experiments, so we may assume they were conducted at room temperature (293.15 K); at any rate, any reasonable deviation from this value (e.g. conducting the experiments at 0 °C, or 273.15 K) will result in an error of less than 10% in our calculation of the latent heat. Making the same assumptions as before and again taking cytoplasm to be equivalent to water (meaning it has a specific heat equal to that of water, 1,000 cal/kg K), we can thus calculate the volume change of the membrane (recall that its thickness increases by about 1 nm) and thus obtain a value of 0.03 pcal for Q_c and a range of 0.6–3.5 pcal for Q_1 . The total heat released by the nerve, then, was about 100 times the amount of heat the Heimburg-Jackson model can account for.

Howarth et al: partially reversible heat release in rabbit vagus nerve

Half a century later, in 1968, J V Howarth, R D Keynes and J M Ritchie saw partially reversible heat release.³⁹ They extracted rabbit vagus nerves 70–80 mm in length, removed the myelin sheath and mounted them on a device consisting of a thermopile and eleven electrodes. The nerve being studied was submerged in saline solution at a temperature close to 5 °C. It was stimulated by a single 2-ms pulse at a time, but the voltage was not specified by the authors; they merely claimed it was “sufficient to excite all the fibres in the nerve”. Recording and stimulation were both extracellular. The thermopile detected a slow increase of 4.2×10^{-6} °C over 206 ms; the authors calculated that this was equivalent to 7.20 µcal/g. This was followed by a decrease in temperature at the thermopile of 2.9×10^{-6} over 380 ms, a reabsorption of 4.90 µcal/g. Oddly, when the nerves were made inexcitable, there was only an increase, which followed exactly the same time course as the partially reversible increase seen with excitable nerves but then continued until it reached a total of 5.0×10^{-6} °C (8.14 µcal/g).

It is difficult to explain these findings using either model. The rabbit vagus nerve contains approximately 40,000 axons, each about 1.5 µm in radius on average.⁴⁰ Howarth et al reported that the area of membrane in 1 g of rabbit vagus nerve is 6,000 cm², which we

can use, along with the same assumptions as before, to calculate the necessary parameters for the Hodgkin-Huxley and Heimbürg-Jackson calculations.

If we assume that the channel density and internodal distance (even though the nerves were desheathed by the researchers, they were originally myelinated, so the ion channel distribution was obviously that of a myelinated nerve) are the same for rabbit vagus nerve as for frog sciatic nerve, we can calculate the number of channels in the nerve and obtain a value of about 2.36×10^{-18} cal/g using the Hodgkin-Huxley model. Conversely, the authors claimed the heat release predicted by the Hodgkin-Huxley model is 750 ncal/g, but they assumed a resting membrane voltage of -34 mV; using -70 mV, this would be $3.18 \mu\text{cal/g}$. Their calculation would require the rabbit vagus nerve to have about 10^{12} times more ion channels than the frog sciatic nerve, which is unrealistic; it is unknown how they reached this value. At any rate, the model cannot explain the heat reabsorption which followed the release.

Using the same values as before for the thermodynamic variables, which seems reasonable because these refer to the lipid membrane itself and not to species-specific details of nerve anatomy, the Heimbürg-Jackson model can account for a total release of 0.82 pcal (which amounts to 37 pcal/g). This is about 10^5 times smaller than the experimental observation. This release should be almost fully reversible, however. In any case, the time course of the heat release is around 100 times longer than an action potential's duration.

Neither model can explain why the inexcitable nerves also produced heat. It seems likely, given the slow time course of the heat changes and the results with excitable and inexcitable nerves, that what was detected by Howarth et al was heating of the nerve tissue due to the current applied and that the reversibility seen with excitable nerves was (relatively slow) regulation of the axon temperature by the axons themselves simply as a result of them being alive.

2.3.2. Action potential termination by heating

Spyropoulos: activity termination by heating in toad thigh nerve

In 1961, Constantine S Spyropoulos extracted the nerves connected to the sartorius and semitendinosus muscles of the toad and carried out some heating and cooling experiments on them.⁴¹ In each experiment, he stimulated them extracellularly with 5 pulses at a time with a frequency of 0.1 Hz and recorded the electrical response at the nodes of Ranvier also extracellularly. Either Ni^{++} , NiCl_2 or NaCl was added to the solution surrounding the nerve, which initially at a temperature of $2-5$ °C; the addition of these ions or compounds to the solution and the low temperature of the system served to prolong the cells' electrical activity to several seconds. In one of his experiments, the nerve was surrounded by a NiCr or wolfram $0.1-\Omega$ coil 0.3 mm in diameter through which a 3-A current pulse lasting 2–5 ms travelled; this current heated the wire, which heated the solution and the nerve by conduction. The temperature of the wire was measured with a thermocouple or a thermistor in contact with the coil, and it was assumed that the duration of the experiment was enough for the nerve, solution and wire to reach thermal

equilibrium. In another experiment, he heated the nerve at the node with a stream of hot solution shot from a syringe. In this case, the temperature change was calculated by Spyropoulos using the heat dissipated by the current (as mentioned previously, this is equal to RI^2). In both experiments, he observed that electrical activity was interrupted (with an occasional delay) by a rise of 2–5 °C. He did not quantify the delay, so it is impossible to analyse it rigorously, but we may speculate that it was due to the finite time required by heat dissipation, as indeed suggested by Spyropoulos himself. However, he did state that the temperature rise required depended on factors such as the initial temperature of the nerve, the rate of temperature change, the delay between the onset of activity and the current pulse and the duration of the electrical activity itself.

These findings can be easily explained by invoking the phase transition model; as stated before, the phase transition of the membrane occurs about 15 °C below the membrane's normal temperature, and activity probably does not involve a much greater cooling than required, so heating the cells during activity brings them above the transition temperature and back into the liquid-disordered phase.

Explaining the observed effect of mid-activity heat under the Hodgkin-Huxley model requires extrapolation. Spyropoulos remarked that another experiment had seen an increase of the resting membrane voltage of the squid giant axon with cooling. By assuming the opposite temperature change causes the opposite voltage effect, we may speculate that a sudden rise in temperature hyperpolarises the membrane, terminating the activity.

A change in initial temperature will affect how far the transition temperature is and thus how easy it is to reach it (and also how easy it is to return above the transition temperature, assuming the change is 15 °C irrespective of how much change would be actually needed); this can serve as an explanation for the initial-temperature dependence of the temperature change required to terminate activity, but only if we assume the membrane composition does not change to adapt to the reduction in resting temperature. On the other hand, recalling the effect of reducing the resting temperature on the membrane potential and assuming the action potential has the same amplitude regardless of the resting membrane potential, we may see how the voltage (and thus temperature) change required to terminate activity increases or decreases depending on the cells' initial temperature.

There is no reason in the Heimburg-Jackson model for there to be a significant change in how much heating is required depending on what fraction of the action potential has occurred. The transition is binary, meaning that once a given region of the membrane has undergone the phase transition it remains in the solid-ordered phase until the action potential moves beyond that region and it returns to the liquid-disordered phase; it shouldn't matter how much time has elapsed since that particular region underwent the transition. In the Hodgkin-Huxley model, however, the membrane potential is directly related to the elapsed time, so the voltage (and thus temperature) change required to bring the voltage back to its resting value changes with time.

Conversely, the rate of temperature change, as long as it remains quicker than the action potential, should have no bearing on the required magnitude of the change in an electric model, but it makes sense to think that it does have an effect in a thermodynamic

model — if this rate is sufficiently high, the membrane will not be in thermodynamic equilibrium while the temperature changes, meaning there might be a concomitant change in other variables (such as pressure) which acts together with or against the heating, modifying the required magnitude.

2.3.3. Action potential generation by cooling

Spyropoulos: activity initiation by cooling in toad thigh nerve

Spyropoulos's 1961 experiments⁴¹ also looked at the effect of cooling on the cells' activity. In an inversion of the experiment where he used solution blasts from a syringe, he used cold blasts and observed that sufficient cooling could initiate activity. It is rather unfortunate that he does not mention how much cooling was necessary for this effect to be seen. In another experiment, he greatly reduced the amount of solution around a node of Ranvier and cooled it with an air flow during activity; here he saw that cooling resulted in a lengthening of the activity by a factor of more than 3.

The initiation of activity by cooling is consistent with the Heimburg-Jackson model if and only if the magnitude of the cooling was enough to bring the axons below the transition temperature. It is consistent with the Hodgkin-Huxley model if we once again invoke the finding in squid giant axon that lowering the temperature depolarises the axon, although one would have to assume that this effect continues far enough to bring the membrane potential above zero, which is not known.

The slowing of activity by decreasing temperature is an effect already seen and exploited by Hodgkin and Huxley and other researchers. For example, in a 1982 experiment similar to the one they made in 1980 where they measured a shortening of the axon during the action potential, Tasaki and Iwasa observed⁷ that action potentials in squid giant axons lasted about twice as long at 6 °C as at 20–22 °C. Unless cooling changes the time constants of the gating particles, however, it cannot be easily explained with the Hodgkin-Huxley model. The Heimburg-Jackson model can partially explain it by recalling the sound equation from section 1.6.2: a temperature decrease results in a density increase, which in turn results in a decrease in the speed of sound. However, a lengthening by a factor of at least 3 (as seen by Spyropoulos) and up to 5 (as seen by Hodgkin and Huxley) would require at least a 9-fold increase, and up to a 25-fold increase, in density, which is unrealistic. It could be that other properties of the cell are altered when the temperature is lowered, changing some other aspect of the nature of the action potential and not merely the speed of sound.

2.4. Chemical aspects of action potentials

2.4.1. Action of anaesthetics and poisons

It was discovered by Hans Horst Meyer in 1899⁴² and independently by Charles Ernest Overton in 1901⁴³ that the median effective dose of an anaesthetic (the concentration of

an anaesthetic that must be present in the body to produce a narcotic effect in half of a population, denoted by ED_{50}) is inversely proportional to the anaesthetic's solubility in vegetable oil. This is true for a wide variety of anaesthetics, including noble gases such as xenon, inorganic molecules such as N_2O , small organic molecules such as chloroform ($CHCl_3$) and diethylether, and larger organic molecules such as octanol and procaine, all of which have very different chemical properties. Further, it seems to hold independently of the animal in question.⁴⁴ This discovery, now known as the Meyer-Overton rule, has been widely assumed by the scientific community to imply that an anaesthetic's ED_{50} is inversely proportional to its solubility in cell membranes as well.

In 2007, Heimburg and Jackson measured the transition temperature of pure dipalmitoylphosphatidylcholine (DPPC), an important and common lipid present in cellular membranes, as well as that of various mixtures of DPPC and the anaesthetic octanol.⁴⁴ They found that the melting temperature decreased by an amount proportional to the octanol concentration in the lipid with a coefficient of determination (r^2) of 0.9989.

The Meyer-Overton rule strongly suggests a thermodynamic origin for the action of anaesthetics. Indeed, if all anaesthetics act the way octanol does (i.e. if they all lower the transition temperature), their action is fully consistent with the Heimburg-Jackson model. In fact, Heimburg and Jackson went as far as reformulating the Meyer-Overton rule to say that all anaesthetics lower the melting temperature of the membrane by an equal amount at ED_{50} .⁴⁴

Defenders of the Hodgkin-Huxley model have long held that anaesthetics act by blocking ion channels. While the activity of some proteins in their purified form, such as luciferase, is indeed inhibited by 50% by approximately ED_{50} doses of various different kinds of anaesthetics,⁴⁵ for this to be a universal anaesthetic mechanism these proteins would all need to have binding sites for an incredibly wide variety of chemical substances, including noble gases, given the seemingly universal applicability of the Meyer-Overton rule. This unrealistic requirement has been recognised by the scientific community, and it has been proposed⁴⁶ that anaesthetics act on proteins not by specific chemical binding, but by increasing the lateral pressure of the membrane at the hydrophilic edges, which in turn alters the structure of the ion channels in such a way that they cannot act normally. This is a Hodgkin-Huxley-consistent explanation for the correlation between anaesthetic potency and anaesthetic solubility in lipid membranes. However, it has been observed that a drastic increase in hydrostatic pressure both increases the melting temperature of lipids and counters the effect of anaesthetics on lipid membranes (though, notably, not on luciferase).^{44,48} This is consistent with the Heimburg-Jackson explanation of anaesthetic action and counter to the lateral-density-increase explanation.

While some poisons, such as tetrodotoxin, do not seem to affect the cell membrane⁴⁷ and do seem to act by inhibiting proteins according to computer simulations,⁴⁹ which is consistent with the Hodgkin-Huxley picture of neural activity, there is no reason to believe poisons should act in the same way as anaesthetics. While the Heimburg-Jackson model does not require proteins to explain action potential propagation along the axon, it says nothing about how action potentials are generated in the first place; if ion channels

are involved in the *initial production* of the action potential, any molecule that inhibits them will act as a poison regardless of what the mechanism of action potential *propagation* is. This is also true for anaesthetics that target luciferase or other proteins; the existence of narcotics and poisons which act by inhibiting activity-related proteins is not mutually exclusive with the existence of narcotics which act by lowering the melting temperature of the membrane.

2.4.2. Anaesthetic potency reduction by pH

Ueno et al: anaesthetic potency reduction by inflammation and pH

In 2008, Takahiro Ueno et al studied the effects of pH reduction and peroxy-nitrite (NO_3^-) on the interaction between artificial lipid bilayers and certain anaesthetics.⁵⁰ They prepared supported lipid bilayers (commonly used as a model for cell membranes; see chapter 3) from palmitoyl-oleoylphosphatidylcholine (POPC), palmitoyl-oleoylphosphatidylethanolamine (POPE), palmitoyl-oleoylphosphatidylserine (POPS), sphingomyelin and cholesterol, all major lipids in the membranes of peripheral neurons. The bilayers were labelled with diphenylhexatriene (DPH), a fluorescent molecule which emits light which is more strongly polarised the less fluid its surrounding medium is. The bilayers were treated with either lignocaine, prilocaine or bupivacaine, all of which are charged anaesthetics,⁵¹⁻⁵³ at clinical concentrations. Since drugs interacting with a lipid membrane increase its fluidity, the degree to which the anaesthetics interacted with the bilayers was determined by measuring the fluorescence polarisation of DPH; a decrease of polarisation indicated that the membrane fluidity had increased, which in turn was an indicator of anaesthetic interaction. The experiments were carried out at a pH of 7.4, similar to physiological conditions. Ueno et al observed very similar polarisation effects with the three anaesthetics: compared to the fluorescence polarisation in control measurements (bilayers without an anaesthetic), the polarisation in the presence of an anaesthetic was reduced by roughly the same amount by all three drugs. This reduction was decreased (i.e. the polarisation was less reduced compared to the control) when the pH was reduced to 6.4, suggesting the interaction between the bilayer and the anaesthetic is weaker at lower pH; the reduction at a pH of 6.4 was about 70% of the reduction at a pH of 7.4. The authors then added peroxy-nitrite to the bilayers. At both pH values, the effect of this was to decrease the polarisation reduction by about 50% (compared to the polarisation reduction at the corresponding pH without peroxy-nitrite).

As a possible explanation, the authors examined A Punnia-Moorthy's 1987 proposal⁵⁴ that inflammation reduces pH. Punnia-Moorthy found this by injecting rats with various inflammatory agents (carrageenan, dextran and the bacterium *staphylococcus aureus*) and measuring the pH of the rats' tissue at the site of injection and of the injected solution, as well as the concentration of leukocytes in the solution, at different time intervals from 4 hr to 48 hr after injection; there was a correlation between inflammation (gauged by the leukocyte count) and pH decrease in dextran and *staphylococcus* (though not in carrageenan), leading Punnia-Moorthy to conclude that inflammation caused a decrease in

the pH of the inflamed tissue by 0.15–0.49. While a decrease in pH did result in a weaker interaction between the membrane and the anaesthetics in Ueno et al's experiment, the authors rejected pH decrease as the mechanism for the reduction in anaesthetic potency. One of their reasons to reject it was that the same experiment with liposomes resulted in the polarisation reduction at a pH of 6.4 being only about 30% of the reduction at a pH of 7.4 when the liposomes were made of DPPC, POPC and cholesterol, POPE and cholesterol, or sphingomyelin and cholesterol (notably, liposomes made of POPS and cholesterol exhibited practically no pH dependence of the polarisation reduction, which the authors attributed to the fact that POPS is a charged lipid and charged drugs have a much higher affinity for charged lipids), which means that the decrease in pH had an effect more than twice as strong for these liposomes than for the model neuron membrane. Their other reason was their suggestion that in live organisms there exist physiological mechanisms related to inflammation which can result in the removal of anaesthetics from the inflamed site or compensate the reduction of pH. Instead, they proposed that peroxynitrite itself binds to the anaesthetics and reduces their efficacy; this conclusion is supported by the fact that peroxynitrite is more reactive at lower pH.

Two important remarks must be made about this experiment in order to properly estimate its significance in the matter of the mechanism behind neural activity.

One is that their experiment did not involve any proteins. Their sample was a lipid membrane representative of the membranes of certain neurons. This is once again strongly suggestive of a thermodynamic mechanism for neural activity where the lipids are the responsible constituents. It is known that, for pH values between 5 and 9, decreasing the pH by 1 results in an increase of the melting temperature of the *Escherichia coli* membrane by almost 2 °C.⁴⁴ If the effect is linear, this means that a decrease by 0.5, as found by Punmia-Moorthy with carrageenan, raises the melting temperature of the membrane by about 1 °C, directly counteracting the effect of anaesthetics; Heimburg and Jackson have suggested that this is sufficient to compensate the effect of anaesthetics, which lower the transition temperature by 0.6 °C.⁴⁴

The other remark is more cautionary. While Ueno et al's results are certainly suggestive, both they and Punmia-Moorthy's findings used by Heimburg and Jackson to estimate the effect of pH on lipid melting temperature must be taken with reserve. First, Punmia-Moorthy only found a correlation between inflammation and pH change for two of the three anaesthetics studied, and it was precisely carrageenan, where the largest pH decrease was observed, that did not display this correlation. Second, the pH reduction induced by dextran and *Staphylococcus* was markedly smaller, only 0.15–0.27. Third, the lower pH value studied by Ueno et al, 6.4, is outside the pH range Punmia-Moorthy's experiments took place; the effect of inflammation might be considerably smaller than the effect seen by Ueno et al, since it produces a pH decrease smaller than the one studied by them. Fourth, Ueno et al's experiment is a rather indirect study with considerable extrapolation to the living case: it was performed on artificial bilayers without proteins, meaning no control for the possible interaction of anaesthetics with proteins was included (providing such a control was, of course, not the authors' intention, as they were merely studying

the relationship between inflammation and pH reduction, but it does make conclusions concerning the mechanism of anaesthetics as related to lipids and proteins less definitive), and it involved a fluorescent label, which may have modified the behaviour of the bilayer.⁵ Fifth, while peroxynitrite is highly toxic and is produced by inflammatory cells,^{55,56} this does not mean that it *causes* inflammation; Ueno et al's use of this ion to simulate inflammation may have been misguided and, in fact, not representative of the effects of inflammation's effect on lipid membranes and their interaction with anaesthetics.

2.4.3. Action potential generation by chemical stimuli

Tasaki and Iwasa: action potential generation by external calcium depletion

In their 1982 experiment, Tasaki and Iwasa also observed generation of action potentials by external calcium depletion.⁷ They immersed a squid giant axon in medium containing only 2-mM Ca^{++} and observed generation of action potentials without the need for electrical or mechanical stimulation. These action potentials occurred with a frequency of 100–150 Hz. They did not, in fact, attempt to record any electrical activity during this experiment, but they observed mechanical and optical responses essentially identical to the ones they had seen with electrical stimulation, so they concluded what they were seeing was a train of action potentials.

The interior and exterior Ca^{++} concentrations are normally 300 nM and 10 mM, respectively, for the squid giant axon.⁵⁷ Both of these concentrations are tiny compared to those of the monovalent ions (K^+ , Na^+ and Cl^-), with the exception of the exterior K^+ concentration, which is only twice the external Ca^{++} concentration, as mentioned in section 1.4.3. Furthermore, the permeability of the squid giant axon membrane to Ca^{++} is only about 0.009 times its permeability to K^+ .^{58,59} This is countered to some extent by the fact that Ca^{++} contributes twice as much charge per ion as the monovalent ions, but this effect is not enough to compensate the low concentrations of and permeability to the divalent ion. Generalising the Goldman-Hodgkin-Katz equation to include divalent ions^{58,60} and taking the values from section 1.4.3 for the monovalent ions, one can see that Ca^{++} contributes less than 1 mV to the membrane potential, so a depletion of 80% of the exterior Ca^{++} concentration cannot produce an effect strong enough to depolarise the membrane beyond the threshold voltage.

The mechanism behind the generation of action potentials by calcium activity, then, must lie elsewhere. Changing the ion concentration on either side of the membrane will alter the pH, which, as we saw earlier in this section, changes the melting temperature of the membrane, but the effect of this is unlikely to be of sufficient magnitude given the low calcium concentration — and, even if the concentration were high enough for depletion to produce a significant pH change, this alone would be insufficient to generate action potentials, since the melting temperature would have to be increased by 10–15 °C. In a reversal of the causal relationship neurophysiologists usually take for granted, Tasaki and Iwasa suggested a swelling of the membrane caused by the calcium depletion might be responsible for the generation of action potentials. While it is unclear why the partial

depletion of an ion which occurs in such a low concentration compared to the univalent ions would cause swelling, this remains the most likely explanation.

It is possible that the response Tasaki and Iwasa saw was not due to action potentials, though this is unlikely given the time course and shape of the mechanical and optical responses, which correspond to the time course and shape of the mechanical and optical effects they observed when they electrically recorded action potentials generated by electrical stimulation.

2.5. The current state of our understanding of action potentials

There is a vast body of experimental results which at first glance have shed some light on the mechanism behind neural activity. However, on second inspection it is evident that there remains much confusion, in large part because neither the Hodgkin-Huxley model nor the Heimburg-Jackson model offers an entirely adequate explanation of all the observations and in part because these observations occasionally contradict each other. A large fraction of the reason these issues have not yet been resolved is that no experiments have yet been designed specifically for the purpose of directly detecting the presence or absence of a phase transition in the axon membrane associated with activity; all of the experiments we have reviewed here either were undertaken with other goals in mind or obtained results which can be explained both with and without invoking phase transitions.

Neither of the models is very well equipped to explain the mechanical aspects of the action potential. The Heimburg-Jackson model's prediction of membrane swelling by lipid tail stretching is closer to the experimental findings in all cases but sometimes falls short by an order of magnitude; this model is also the only one which can acceptably explain Peracchia's findings of mechanical changes in multiple membranes. Conversely, all of the Hodgkin-Huxley predictions are several orders of magnitude above or below the experimental observations. Furthermore, flexoelectricity offers a rather poor explanation of the observed mechanical phenomena, in some cases predicting unphysical changes in the axon and in other cases predicting changes with the wrong sign. Treating the cytoplasm as an ideal liquid and assuming the membrane experiences the same pressure changes as it is also inappropriate.

The optical aspects of the action potential seem, in general, particularly consistent with the Heimburg-Jackson model; every fragment of evidence suggests the cause of all of the observed changes is a reversible conformational change in the membrane lipids, which is precisely what a phase transition in the membrane would cause. That said, it is beyond the scope of this work, and perhaps impossible with current knowledge, to offer a quantitative explanation of the observed changes in scattering, birefringence and opacity, as mentioned earlier. It is also difficult to explain some of the observations on nerves without assuming the behaviour of the axons is blurred by superposition of all the axons' poorly coordinated activity, but this is merely an explanation of why the results of the experiments on nerves are strange and sometimes mutually inconsistent rather than an adequate explanation of which changes are occurring in the nerve.

As in the mechanical and optical cases, there is evidently a large gap in our understanding of the thermodynamic aspects of the action potential and a need for more-sensitive experiments to fill this gap. Neither model can adequately explain the observed heat release. Given the enormous magnitude of this release compared to the models' predictions, it is likely that some other phenomenon, such as heating of the nerves by the current supplied to them, is responsible for the experimental observations and completely masks the effects predicted by the models. If we accept this, then it is possible that both models make correct predictions but these have been thus far too small to detect. The effects of temperature on cell activity are variously consistent with one model or the other, but rarely with both, casting confusion on the matter of which model, if any, is more complete. Activity initiation by cooling and termination by heating is precisely what the Heimburg-Jackson model predicts, while the Hodgkin-Huxley model must resort to invoking a poorly understood and characterised depolarisation caused by cooling and assuming the inverse of this effect is also true in order to explain this. Likewise, thermodynamics may explain why the rate, as well as the magnitude, of the temperature change is important, while the Hodgkin-Huxley model has no in-built explanation for this. Conversely, observations like the slowing of the action potential with cooling and the temporal variation of the threshold temperature change required for action potential termination are consistent with the electric model but difficult to explain with the Heimburg-Jackson model.

While the Heimburg-Jackson theory offers a particularly elegant way to describe the mechanism of anaesthetics and explain the generality of the Meyer-Overton rule, both theories are largely consistent with the observations related to anaesthetics, and elegance alone does not constitute a valid reason to choose one theory over another. Of the chemical findings discussed in this chapter, the pressure and pH reversal of anaesthetic action is alone in being consistent with only the Heimburg-Jackson model, although the pH case must be taken with caution due to the somewhat dubious nature of the experiments leading to this observation. Excitation by calcium depletion remains mysterious under either theory.

It is frequently seen that the results obtained with one species are different from, and sometimes opposite to, those obtained with another. Whether this is due to differences in experimental configurations or to actual differences in the animals' nerve cells is presently unknown. This, together with the fact that several of these studies were on artificially induced, long-lived voltage changes rather than induced or spontaneous action potentials, increases the difficulty of understanding exactly what happens to the cell during activity and developing a theoretical model to explain it. Any explanation that invokes membrane proteins must necessarily consider the fact that these proteins vary somewhat between species, while thus far it has been the norm to assume they all function in much the same way. Conversely, it would be difficult to imagine such a general phenomenon as a phase transition is species-specific; if it occurs, a thermodynamic model would be truly universal without requiring every individual species to be studied separately. Put simply, while biochemistry is different in each species, thermodynamics is not. That said, a phase transition has never been directly observed and thermodynamic models of neural activity have

received comparatively little attention.

2.6. Requirements for increasing our understanding of action potentials

It is now clear that much work remains to be done in this field. While the Heimburg-Jackson model offers a multidisciplinary perspective on action potentials and enjoys a somewhat greater degree of compatibility with experimental findings than the Hodgkin-Huxley model, it remains incapable of explaining some of these findings; furthermore, while the electrical behaviour of the axon membrane is well characterised (if not necessarily well understood), there is conflicting evidence of some of the non-electrical aspects of neural activity. Despite the evidence agreeing better with a phase-transition model than with cable theory, whether the membrane undergoes a phase transition or not is still unknown.

Practically all experimental work on action potentials has so far been electrophysiological. Voltage recording has become the standard way to observe neural activity, with other techniques being subject to electrophysiological confirmation that whatever signal is detected is concomitant with an action potential. However, electrophysiology is necessarily an extremely invasive technique which often involves rather artisanal methods and invariably results in disturbance to the cellular environment and, in the case of internal recording (and/or stimulation when it is required), the death of the cell.

Unfortunately, electrophysiology remains the gold standard even today, nearly a century after its earliest use in the field.⁶¹ This means that, in order to one day reach the ideal of noninvasive recording of action potentials, we must first transition through a phase of confirmation of some noninvasive technique using electrophysiology until such a time as action potential recording with said technique is well established.

This establishes the first two characteristics our hypothetical better technique: it must be noninvasive, and it must be possible to use it simultaneously with electrophysiology. Of course, it must also be capable of distinguishing signals coming from phase transitions from other types of signals.

There are countless options when it comes to sample type. Giant axons are visible by eye (see chapter 3) and relatively easy to handle, but they must be stimulated, duplicating the difficulty of the electrophysiological component of the experiment; conversely, cultured neurons often display spontaneous activity, but they are microscopic and fragile. The experimental setup must be capable of solving one of these two problems. In the case of cultured neurons, mammalian ones are the most relevant to us as humans. In all cases, the axon radius is well above the typical spatial resolution of standard microscopy using visible light, which is a few hundred nanometres; the problem lies not in resolving the axon, but in manipulating it.

The action potential lasts around 1 ms at body temperature and can be stretched to a few milliseconds by working at lower temperatures. This means that the Nyquist limit

for acquisition frequency is between about 500 Hz and 2 kHz; desirable acquisition frequencies are 10 kHz or more.

The requirement for noninvasiveness all but restricts the hypothetical technique to optical sensing. Mechanical, electrical and thermal probes (such as atomic force microscopy, electrophysiology and temperature sensing via a thermocouple, respectively) all require the sensor to be in contact with the sample and often with the cell itself. Electron microscopy provides a strong contrast and unprecedented spatial resolution, but it is highly destructive and requires fixing the sample somehow, making it incompatible with in-vivo measurements. In contrast, optical techniques can be noninvasive, do not always require a contrast agent, and involve low enough power not to cause photodamage or thermal damage to the sample. These requirements themselves rule out optical techniques such as fluorescence microscopy and staining, which require the addition of a contrast agent to the sample; x-ray scattering microscopy, which requires averaging the signal of multiple acquisitions and can thus not provide real-time information; and transmission microscopy techniques such as phase microscopy and differential interference contrast, which are not quantitative and do not have enough axial resolution to observe 1-nm thickness changes (it will be seen in chapter 3 that differential interference contrast can be made quantitative and sensitive enough to resolve these changes, but that requires taking two successive images, which renders it incapable of providing a signal in real time).

Interferometric reflectometry and coherent Raman scattering microscopy, on the other hand, satisfy these constraints. They are both noninvasive and, as will be seen in chapters 4 and 5, capable of discriminating between different phases of a material with high temporal resolution.

It will be seen in chapter 4 that the reflectometry signal corresponding to a 1-nm increase in lipid bilayer thickness is approximately 0.1% of the bilayer signal in amplitude (a change from 99.8% of the glass-water reflection coefficient for 4 nm to 99.7% for 5 nm assuming the refractive index does not change appreciably, as would be the case if there is no phase transition; if it does (which would be the case if there is a transition),^{11,62} it might go from about 1.440 to about 1.487, in which case the signal change would be reduced to about 0.025% but up to 30% of it in phase (a change from about 0.8 rad to about 1.1 rad). We will see whether the sensitivity of interferometric reflectometry is high enough to detect such signal differences; if it can resolve at least the higher-signal non-transition case, it will allow us to determine whether such a transition takes place during cell activity.

The difference in coherent-Raman signal depends heavily on the wave number but can easily reach 50% in some cases. This technique has been used in the past to distinguish between lipid types and phases in bulk lipid (see chapter 5); we will see whether this can also be done with single lipid bilayers.

2.7. References

1. Hill D K (1950): *The volume change resulting from stimulation of a giant nerve fibre*, Journal of Physiology 111, 304-327

2. Damask A C & Swenberg C E (1984): *Medical physics, volume III: synapse, neuron, brain*, Academic Press, 219
3. Alberts B, Johnson A, Lewis J, Raff M, Roberts K & Walter P (2008): *Molecular biology of the cell*, fifth edition, Garland Science, 618
4. Small D M (1984): *Lateral chain packing in lipids and membranes*, *Journal of Lipid Research* 25, 1490–1500
5. Regan D, Williams J B, Borri P & Langbein W (2019): *Lipid bilayer thickness measured by quantitative DIC reveals phase transitions and effects of substrate hydrophilicity*, *Langmuir* 35, 13805–13814
6. Iwasa K & Tasaki I (1980): *Mechanical changes in squid giant axons associated with production of action potentials*, *Biochemical and Biophysical Research Communications* 95, 1328–1331
7. Tasaki I & Iwasa K (1982): *Further studies of rapid mechanical changes in squid giant axon associated with action potential production*, *Japanese Journal of Physiology* 32, 505–518
8. Tasaki I & Iwasa K (1980): *Shortening of nerve fibers associated with propagated nerve impulse*, *Biochemical and Biophysical Research Communications* 94, 716–720
9. Bullock T H, Horridge G A, Bern H A, Hagadorn I R & Smith J E (1965): *Structure and function in the nervous systems of invertebrates, volume I*, W H Freeman, 136–137
10. González-Pérez A, Budvytyte R, Mosgaard L D, Nissen S & Heimburg T R (2014): *Penetration of action potentials during collision in the median and lateral giant axons of invertebrates*, *Physical Review X* 4, 031047
11. Huang W-T & Levitt D G (1977): *Theoretical calculation of the dielectric constant of a bilayer membrane*, *Biophysical Journal* 17, 111–128
12. Terakawa S (1983): *Changes in intracellular pressure in squid giant axons associated with production of action potentials*, *Biochemical and Biophysical Research Communications* 114, 1006–1010
13. Otero L, Molina-García A D & Sanz P D (2002): *Some interrelated thermophysical properties of liquid water and ice: a user-friendly modeling review for high-pressure processing*, *Critical Reviews in Food Science and Nutrition* 42, 3399–352
14. Peracchia C & Robertson D (1971): *Increase in osmiophilia of axonal membranes of crayfish as a result of electrical stimulation, asphyxia, or treatment with reducing agents*, *Journal of Cell Biology* 51, 223–239
15. Chester R (1990): *Marine geochemistry*, Unwin Hyman, 251

16. Bednarczyk P (2009): *Potassium channels in brain mitochondria*, Acta Biochimica Polonica 56, 385–392
17. Sepehri H, Eliassi A, Sauvé R, Ashrafpour M & Saghiri R (2007): *Evidence for a large conductance voltage gated cationic channel in rough endoplasmic reticulum of rat hepatocytes*, Archives of Biochemistry and Biophysics 457, 35–40
18. Khodae N, Ghasemi M, Saghiri R & Eliassi A (2014): *Endoplasmic reticulum membrane potassium channel dysfunction in high fat diet induced stress in rat hepatocytes*, Experimental and Clinical Sciences Journal 13, 1075–1087
19. Julian F J & Goldman D E (1962): *The effects of mechanical stimulation on some electrical properties of axons*, Journal of General Physiology 46, 297–313
20. Solomon S & Tobias J M (1960): *Thixotropy of axoplasm and effect of activity on light emerging from an internally lighted giant axon*, Journal of Cellular and Comparative Physiology 55, 159–166
21. Bear R S, Schmitt F O & Young J Z (1937): *Ultrastructure of nerve axoplasm*, Proceedings of the Royal Society of London B 123, 505–519
22. Ritchie J M (1973): *Energetic aspects of nerve conduction: the relationships between heat production, electrical activity and metabolism*, Progress in Biophysics and Molecular Biology 26, 147–187
23. Shaw S N & Tobias J M (1951): *Preliminary measurements of birefringence changes due to electrical polarization of isolated squid giant axones*, Biological Bulletin 101, 229–230
24. Tasaki I, Watanabe A, Sandlin R & Carnay L (1968): *Changes in fluorescence, turbidity, and birefringence associated with nerve excitation*, Proceedings of the National Academy of Sciences of the United States of America 61, 883–888
25. Cohen L B, Keynes R D & Hille B (1968): *Light scattering and birefringence changes during nerve activity*, Nature 218, 438–441
26. Cohen L B, Hille B & Keynes R D (1970): *Changes in axon birefringence during the action potential*, Journal of Physiology 211, 495–515
27. Cohen L B & Keynes R D (1971): *Changes in light scattering associated with the action potential in crab nerves*, Journal of Physiology 212, 259–275
28. Bohren C F & Huffman D R (1983): *Absorption and scattering of light by small particles*, Wiley, 158–159
29. Cohen L B, Keynes R D & Landowne D (1972): *Changes in light scattering that accompany the action potential in squid giant axons: potential-dependent components*, Journal of Physiology 224, 701–725

30. Hill A V (1926): *The heat production of nerve*, Journal of Pharmacology and Experimental Therapeutics 29, 161–165
31. Shrager P (1988): *Ionic channels and signal conduction in single remyelinating frog nerve fibres*, Journal of Physiology 404, 695–712
32. Shrager P (1987): *The distribution of sodium and potassium channels in single demyelinated axons of the frog*, Journal of Physiology 392, 587–602
33. Ribaiy H N (2017): *A short guide to electrophysiology and ion channels*, Journal of Pharmacy and Pharmaceutical Science 20, 48–67
34. Tasaki I (1952): *Properties of myelinated fibers in frog sciatic nerve and in spinal cord as examined with micro-electrodes*, Japanese Journal of Physiology 3, 73–94
35. Peters T J & Cairns S R (1985): *Analysis and subcellular localization of lipid in alcoholic liver disease*, Alcohol 2, 447–451
36. Parks J E, Arion J W & Foote R H (1987): *Lipid of plasma membrane and outer acrosomal membrane from bovine spermatozoa*, Biology of Reproduction 37, 1249–1258
37. Adam A & Friede R L (1988): *The number of frog axons increases continually during body growth*, Anatomy and Embryology 178, 537–541
38. Mitaku S, Jippo T & Kataoka R (1983): *Thermodynamic properties of the lipid bilayer transition*, Biophysical Journal 42, 137–144
39. Howarth J V, Keynes R D & Ritchie J M (1968): *The origin of the initial heat associated with a single impulse in mammalian non-myelinated nerve fibres*, Journal of Physiology 194, 745–793
40. de Neef K J, Jansen J R C & Versprille A (1982): *Developmental morphometry and physiology of the rabbit vagus nerve*, Developmental Brain Research 4, 265–274
41. Spyropoulos C S (1961): *Initiation and abolition of electric response of nerve fiber by thermal and chemical means*, American Journal of Physiology 200, 203–208
42. Meyer H H (1899): *Zue Theorie der Alkoholnarkose. Erste Mittheilung. Welche Eigenschaft der Anästhetica bedingt ihre narkotische Wirkung?*, Archiv für experimentelle Pathologie und Pharmakologie 42, 109–118
43. Overton C E (1901): *Studies of narcosis*, Chapman and Hall
44. Heimburg T R & Jackson A D (2007): *The thermodynamics of general anesthesia*, Biophysical Journal 92, 3159–3165
45. Franks N P & Lieb W R (1984): *Do general anaesthetics act by competitive binding to specific receptors?*, Nature 310, 599–601

46. Cantor R S (1997): *The lateral pressure profile in membranes: a physical mechanism of general anesthesia*, *Biochemistry* 36, 2339–2344
47. Appali R, van Rienen U & Heimburg T R (2012): *A comparison of the Hodgkin-Huxley model and the soliton theory for the action potential in nerves*, in *Advances in planar lipid bilayers and liposomes* 16, Elsevier, 275–299
48. Blicher A, Wodzinska K, Fidorra M, Winterhalter M & Heimburg T R (2009): *The temperature dependence of lipid membrane permeability, its quantized nature, and the influence of anesthetics*, *Biophysical Journal* 96, 4581–4591
49. Chen R & Chung S-H (2014): *Mechanism of tetrodotoxin block and resistance in sodium channels*, *Biochemical and Biophysical Research Communications* 446, 370–374
50. Ueno T, Tsuchiya H, Mizogami M & Takakura K (2008): *Local anesthetic failure associated with inflammation: verification of the acidosis mechanism and the hypothetic participation of inflammatory peroxynitrite*, *Journal of Inflammation Research* 1, 41–48
51. Binshtok A M et al (2009): *Co-application of lidocaine and the permanently charged sodium channel blocker QX-314 produces a long-lasting nociceptive blockade in rodents*, *Anesthesiology* 111, 127–137
52. Sweet W D, Wood W G & Schroeder F (1987): *Charged anesthetics selectively alter plasma membrane order*, *Biochemistry* 26, 2828–2835
53. Chahar P & Cummings III, K C (2012): *Liposomal bupivacaine: a review of a new bupivacaine formulation*, *Journal of Pain Research* 5, 257–264
54. Punnia-Moorthy A (1987): *Evaluation of pH changes in inflammation of the subcutaneous air pouch lining in the rat, induced by carrageenan, dextran and staphylococcus aureus*, *Journal of Oral Pathology & Medicine* 16, 36–44
55. McCafferty D M (2000): *Peroxyntirite and inflammatory bowel disease*, *Gut* 46, 436–439
56. Pacher P, Beckman J S & Liaudet L (2007): *Nitric oxide and peroxynitrite in health and disease*, *Physiological Reviews* 87, 315–424
57. Lodish H, Berk A, Matsudaira P, Kaiser C A, Krieger M, Scott M P, Zipursky L & Darnell J (2003): *Molecular cell biology*, fifth edition, W H Freeman, 253
58. Hodgkin A L & Keynes R D (1957): *Movements of labelled calcium in squid giant axon*, *Journal of Physiology* 138, 253–281
59. Caldwell P C & Keynes R D (1960): *The permeability of the squid giant axon to radioactive potassium and chloride ions*, *Journal of Physiology* 154, 177–189
60. Hodgkin A L & Katz B (1949): *The effect of sodium ions on the electrical activity of the giant axon of the squid*, *Journal of Physiology* 108, 37–77

61. Adrian E D (1954): *The basis of sensation*, British Medical Journal 1, 287–290
62. Howland M C, Szmodis A W, Sani B & Parikh A N (2007): *Characterization of physical properties of supported phospholipid membranes using imaging ellipsometry at optical wavelengths*, Biophysical Journal 92, 1306–1317

Chapter 3

Methods

Cells are extremely complicated objects. A cell houses densely packed organelles, proteins, aminoacids and small molecules in its interior; its membrane is comprised of hundreds, or sometimes over a thousand, different kinds of lipids, as mentioned in section 1.4.1, and contains many embedded proteins; and it is surrounded by ions, connective tissue and other cells. Furthermore, cells capable of showing electrical activity are alive and thus capable of reacting to experimental conditions. Finally, since no two cells are exactly identical in terms of size, shape, developmental stage, health, access to nutrients, nutrient storage and intercellular connections, often two cells of the same type will react somewhat differently to identical experimental conditions. All of this introduces a large number of factors which complicate experimentation on live cells.

While experiments seeking to record action potentials must necessarily be undertaken on cells which are alive and healthy, some simpler experiments, such as attempting to observe the optical signal caused by a phase transition in a lipid membrane or determine a membrane's composition, are best undertaken on simpler systems which replicate the aspects of the cell which are important to the experiment but dispose of the unwanted complexity.

When modelling parts of a complex system using a simpler system (called a model system), it is important to choose the simpler system's properties carefully. Since we wish to observe phase transitions in lipid bilayers (so we may later determine whether or not they take place in the context of neural activity), the model system must be a lipid bilayer formed either by lipids whose melting temperature can be crossed during the experiment or by a mixture of lipids which are in different phases at the experimental temperature. Since the phase transition we are interested in involves only the lipids, our model system can be protein-free. Because we wish to undertake noninvasive experiments, the bilayer must contain no fluorophores, nanoparticles or other labels. Thus, for experiments where we merely wish to determine whether our equipment is capable of detecting a phase transition or distinguishing between lipid phases but we do not need to perform electrophysiology, a supported lipid bilayer made of lipids with the aforementioned transition temperatures is an adequate model system.

This chapter will describe the procedures followed when preparing the model and live samples (section 3.1) which were used in the experiments whose results are discussed in chapters 4 and ??, as well as the experiments themselves (sections 3.2–3.5).

3.1. Sample preparation

None of the samples used were prepared by the author of this work, with the exception of the lobster samples (section 3.1.3).

3.1.1. Supported lipid bilayers

Coverslip etching

Lipid molecules are amphiphilic (i.e. they have a hydrophobic part and a hydrophilic part). Therefore, when submerged in water or a water-like substance they form structures with a hydrophilic exterior and a hydrophobic interior. One such structure is a lipid bilayer, a (relatively) flat double layer of lipid molecules with the hydrophobic tails of each layer pointing at the other layer and the hydrophilic heads pointing outwards¹ (figure 3.1). Thus, in order for a lipid bilayer to properly adhere to a substrate, the substrate must be hydrophilic. Untreated glass is not sufficiently hydrophilic, so it must be made so (for example, by acid etching²).

Glass coverslips to be used were first washed with acetone to remove inorganic contaminants and then submerged in 60 ml of sulphuric acid. This was then placed in a water bath at 95 °C. After several minutes (so the acid had had time to warm up to the water bath temperature), 20 ml of hydrogen peroxide were added to the sulphuric acid. The

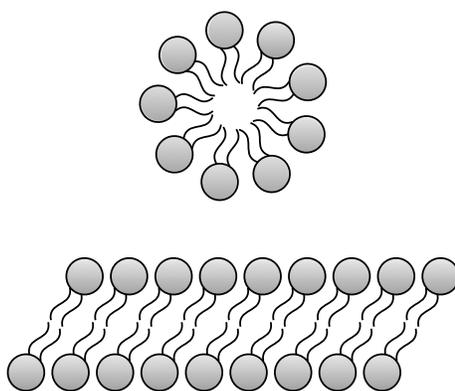


Figure 3.1: Diagram of some of the structures formed by lipids: a micelle (top) and a bilayer (bottom). These structures are formed when the lipid is submerged in a hydrophilic medium, such as water; both leave the hydrophilic heads pointing outwards, in contact with the surrounding medium, and the hydrophobic tails sequestered inside the structure.

coverslips were allowed to sit in the mixture for 1 hr, after which they were washed with distilled water and dried under a stream of nitrogen gas. The etching process also served to remove organic contaminants.

While waiting to be used, etched coverslips were stored in nitrogen gas at 4–6 °C so their hydrophilicity would be preserved, as exposure to air would have resulted in oxidation and thus a decrease in the hydrophilicity of the glass.

Lipid solution preparation

Lipid solutions were prepared by mixing dioleoylphosphatidylcholine (DOPC, Avanti Polar Lipids), chicken egg sphingomyelin (sm, Avanti Polar Lipids) and cholesterol (ch, Avanti Polar Lipids) in the following concentrations (molar ratio, DOPC:sm:ch) in glass vials: 1:0:0 (pure DOPC); 0:7:3 (sm+ch); and 3:5:2, 2:2:1 and 11:5:4 (ternary mixtures). A pure-dipentadecanoylphosphatidylcholine (DC₁₅PC, Avanti Polar Lipids) solution was also prepared. The vials were heated to 50 °C on a hot plate, and the lipids were agitated by a weak flow of nitrogen gas so they would mix properly. The vials were then placed inside a desiccator for 1 hr to remove any trace solvent present in the mixture. After desiccation, the lipid solutions were finally diluted in solvent to achieve a concentration of 0.8–1 mg/ml, which was previously found to be the optimal range of concentrations for mostly unilamellar (single-bilayer) lipid patches to form over part of the coverslip with the spin-coating procedure used; for reasons which will become clear in section 5.1, each sample needed to have only partial coverage of the coverslip. The solvent used for the DOPC sample was 95% chloroform and 5% acetonitrile (by volume), and that used for the other samples was isopropanol.

The lipid solutions were stored at –20 °C while not used.

Lipid bilayer formation

As mentioned in section 1.6.1, bilayers are partially ordered but can still be fluid; they can be in a gel-like phase called the solid-ordered (SO) phase, or they can be in a liquid state where the ordering of the molecules decreases considerably, called the liquid-disordered (LD) phase^{3,4} (figure 3.2, top and bottom). The melting temperatures of DOPC bilayers^{3,5} and sphingomyelin bilayers^{3,6,7} —that is, the temperatures below which the bilayers are in the SO state and above which they are in the LD state— are –16.5 °C and 37 °C, respectively. Cholesterol doesn't form bilayers on its own, but when mixed with other lipids that do it interacts with the hydrocarbon tails of those lipids, influencing their structure and giving rise to the liquid-ordered (LO) phase, which is still liquid but has a molecular order intermediate between those of the LD and SO phases^{3,4} (figure 3.2, centre). At room temperature, then, a DOPC bilayer is in the LD phase and a sphingomyelin-cholesterol mixture containing between 30% and 50% cholesterol is in the LO phase^{7,8} (see also figure 3.3). Cholesterol exhibits a higher affinity for saturated lipids, such as sphingomyelin, than for unsaturated ones, such as DOPC.^{4,9,10} Therefore, a ternary mixture of these three lipids segregates into LO domains made of sphingomyelin and cholesterol in a 17:10 ratio and LD domains made of DOPC mixed with the small remaining amount of cholesterol

in the mixture.^{7,9,11}

Planar lipid bilayers were formed on the etched coverslips by either rupturing of electroformed giant unilamellar vesicles (2:2:1 mixture) or spin-coating (all other samples). The procedures are described below. Which technique was used for which lipid solution was guided by observations of which method produced more desirable results in each case; the parameters evaluated were lipid coverage, regularity of the lipid patches and cleanliness of the regions without lipid.

In the giant unilamellar vesicle (GUV) rupturing method,^{12,13} 20 μl of the lipid solution are deposited on a pair of tantalum electrodes under a weak flow of nitrogen gas. The lipid is deposited only in the lower half of the electrodes, and only a few droplets at a time to allow the solution to be dried by the nitrogen flow. The electrodes are then placed in vacuum for 1 hr in order to remove any remaining solvent. Once this is done, 1.2 ml of distilled water are poured into a microcentrifuge tube, degassed for 5 min in vacuum and finally warmed to 70 $^{\circ}\text{C}$. The lipid-coated ends of the electrodes are submerged in the hot water, and the dry ends of the electrodes are connected to a function generator. The generator provides first a square voltage wave with a peak-to-peak amplitude of 1.2 V and a frequency of 10 Hz for 1 hr and then a 1.5-V sinusoidal wave with variable frequency:

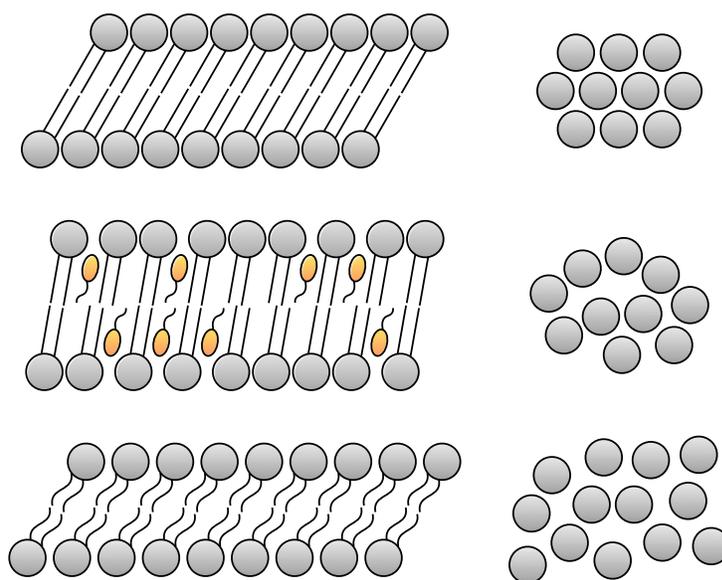


Figure 3.2: Lateral (left) and top-down (right) diagrams of three of the phases lipid bilayers can be in: the SO phase (top), the LO phase (centre) and the LD phase (bottom). In the SO phase, the lateral structure is crystalline and the lipid tails are stretched. In contrast, the LD phase is characterised by a disordered lateral packing of the lipid molecules and “kinks” in the tails which cause the bilayer to be thinner. The LO phase is formed by some lipids in the presence of cholesterol (shown in yellow), which nestles among the tails, causing them to be straight, as in the SO phase, but disrupting the lateral packing of the bilayer lipids.

5 Hz for 30 min, then 2 Hz for 15 min and finally 1 Hz for 15 min. The voltage results in the formation of GUVs. A 120- μm -thick, 13-mm-diameter adhesive imaging gasket is then adhered to an etched coverslip to create a shallow well into which 260 μl of the GUV solution are deposited; the solution is placed 65 μl at a time with a temporal spacing of 5–10 min so the GUVs can sink to the bottom of the well. 0.750–1.365 ml of phosphate-buffered saline (PBS) are degassed for 5 min to prevent air bubbles from forming in the sample. 65 μl of degassed PBS are added to the well. The sample is finally diluted to the desired concentration by repeatedly adding another 65 μl of PBS and pipetting out an equal volume of the GUV-water-PBS mixture in the well; depending on the desired concentration, this step is performed a total of 10–20 times. The pipetting takes approximately 1 min and thus does not allow enough time for the PBS to mix with air, which would counteract the degassing.

In bilayer spin-coating,¹⁴ 150 μl of the mixture are deposited on the centre of an etched coverslip, which is then spun at 3,000 rpm for 42 s on a spin-coater. The spin-coater provides a 6-s constant acceleration at the beginning of the 42-s period and a 6-s constant deceleration at the end of it, so the coverslip only rotates at constant speed for the central 30 s. The coverslip is then placed in a centrifuge tube with a small piece of wet tissue so the lipid can absorb some of the moisture and later, when PBS is added to the sample, the lipid is already hydrated and does not absorb a large amount of liquid in a short period of time, which could destroy the structure of the sample. The tube is filled with nitrogen gas to prevent lipid oxidation, sealed and incubated at 37 $^{\circ}\text{C}$ for 1 hr in an

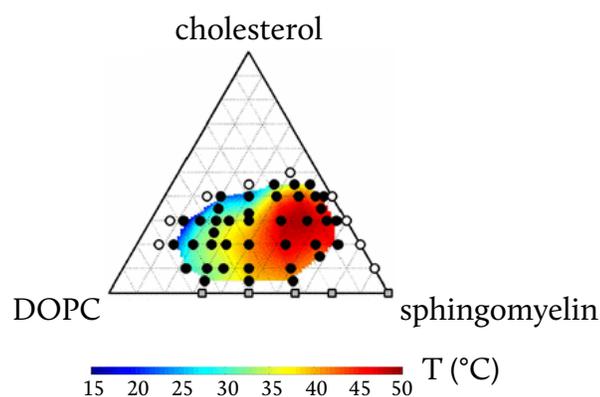


Figure 3.3: Phase diagram of a DOPC+sm+ch mixture. The miscibility temperature of a lipid mixture is the temperature above which separate lipid domains no longer coexist and the mixture melts into a single phase. Here, black circles denote mixtures where lipid segregation into liquid-ordered (LO) and liquid-disordered (LD) domains has been observed, the coloured surface is an interpolation of the miscibility temperature measured at the points represented by the black circles, white circles denote measurements where either no miscibility temperature is observed or it is lower than 10 $^{\circ}\text{C}$, and grey squares denote mixtures where lipid segregation occurs but there are solid-ordered (SO) domains instead of LO domains. Modified from the original.⁸

oven. A shallow well is created on the lipid-coated side of the coverslip by placing an adhesive imaging gasket with the same specifications as in the GUV rupturing procedure. The well is finally filled with degassed PBS.

Regardless of bilayer formation method, a glass microscope slide was washed with acetone and adhered to the top of the gasket to close the cylindrical well. In the case of the ternary mixture, slides were stored at 4 °C for at least a few hours prior to imaging so the LO and LD domains would have enough time to form.

Multilamellar (multiple-bilayer) versions of the sm+ch and DOPC samples were also prepared. The entire procedure was the same as for the unilamellar samples except for the fact that 20 times more lipid was present in the lipid solution.

3.1.2. Mouse neuron culture

Glass coverslips 13 mm in diameter (Agar Scientific, AGL46R13-1) were coated with poly-L-lysine (Sigma-Aldrich) in preparation for neuron culture.

A hippocampal neuron suspension was prepared by extracting and triturating the hippocampi (see figure 3.4¹⁵) of 18-day-old mouse foetuses and treating them with trypsin (Worthington Biochemical Corporation) and deoxyribonuclease (Roche Applied Science).¹⁶

The neurons were then placed on the coated coverslips at densities of 18,000 and 25,000 cells/coverslip (approximately 13,600 and 18,800 cells/cm², respectively). Each coverslip was placed in a 35-mm plastic well with a mixture of 98% neurobasal A medium (Invitrogen) and 2% B27 medium (Invitrogen)¹⁷ with 60 µg/ml of penicillin and 100 µg/ml of streptomycin (Gibco); this way, a single coverslip at a time could be manipulated with-

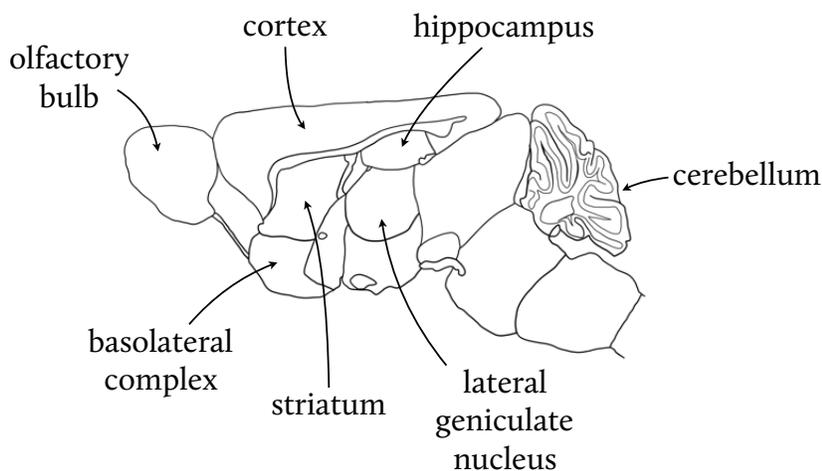


Figure 3.4: Diagram of a mouse brain showing several different sections, including the hippocampus (near top centre). Copied from the original.¹⁵

out disturbing the other coverslips. 0.5-mM GlutaMAX I (Invitrogen) was added to half of the coverslips.

The cultures were incubated at 37°C. The air inside the incubator had 5% CO₂. The culture medium was replaced twice a week.

The neurons were used for electrophysiology experiments from 8 days old since culturing (8 days in vitro, or DIV-8) to 16 days old since culturing (DIV-16), although each individual culture was only used once. Cells lasted about 2 hours under the microscope (see section 3.5). At DIV-16, the cells were heavily degraded; at DIV-17, they were dead.

3.1.3. Lobster giant axon and leg nerves

The work described in this section (3.1.3), as well as that described in sections 3.5.3 and 4.4.1, was performed during a visit to Thomas Rainer Heimbürg's laboratory at the University of Copenhagen, Denmark.

This section contains photographs of dissected lobsters which may not be suitable for every reader.

Saline solution

Saline solution was prepared by mixing 462-mM sodium chloride (NaCl), 16-mM potassium chloride (KCl), 26-mM calcium chloride (CaCl₂), 8-mM magnesium chloride (MgCl₂), 11-mM glucose and 10-mM tris(hydroxymethyl)aminomethane in water in a glass flask. This solution had a pH of approximately 10, which was adjusted to 7.6 with hydrogen chloride (HCl); approximately 0.08% HCl by volume was required for this.^{18,19}

Lobster dissection

Figure 3.5 shows a diagram of a lobster with some of the body parts labelled.²⁰

Each lobster was purchased at a food market and stored at approximately –20 °C for 30 min in order to anaesthetise it. Afterwards, it was held upside-down and its eyes were covered with tissue paper so it would relax; once relaxed, the lobster would extend its legs and uncurl its tail, which made the remainder of the procedure easier. Once this had occurred, the lobster was placed right-side-up on a dissecting table and the dissection began. The thoracic and abdominal cords (which together comprise the lobster's ventral cord, its central nervous system) and the leg nerve bundles were extracted in the following way.¹⁹

All walking legs, including the chelipeds, were severed as close to the body as possible using a large kitchen knife. The chelipeds were set apart for cooking, while the other walking legs were stored in an open container with saline solution in a refrigerator at approximately 4 °C until such a time as the nerves were to be extracted. All antennae were severed as close to the body as possible using large scissors.

The abdomen was then separated from the thorax just behind where the carapace ended, again using the large kitchen knife. The thorax and head were placed upside-down on a petri dish large enough to hold them, and the dish was filled with saline solution.

For extraction of the thoracic cord, the carapace was cut from below on both sides with large scissors and removed to expose the internal organs (figure 3.6, top). The heart, intestine and stomach were removed, and the thorax was cleaned. The thoracic muscles were cut off, and the internal skeleton was cut longitudinally down the centre using medium-sized scissors, maintaining the scissors as horizontal as possible at all times (figure 3.6, upper centre). The two halves of the skeleton were then pried apart slightly to expose the cord, and the cord was gently lifted with small tweezers to avoid damaging it (figure 3.6, lower centre and bottom). The connections between the cord ganglia and the muscles were cut with small shears, as were the parts of the subesophageal ganglion (the brain) that were attached to the lobster's head; care was taken to save as much brain tissue as possible.

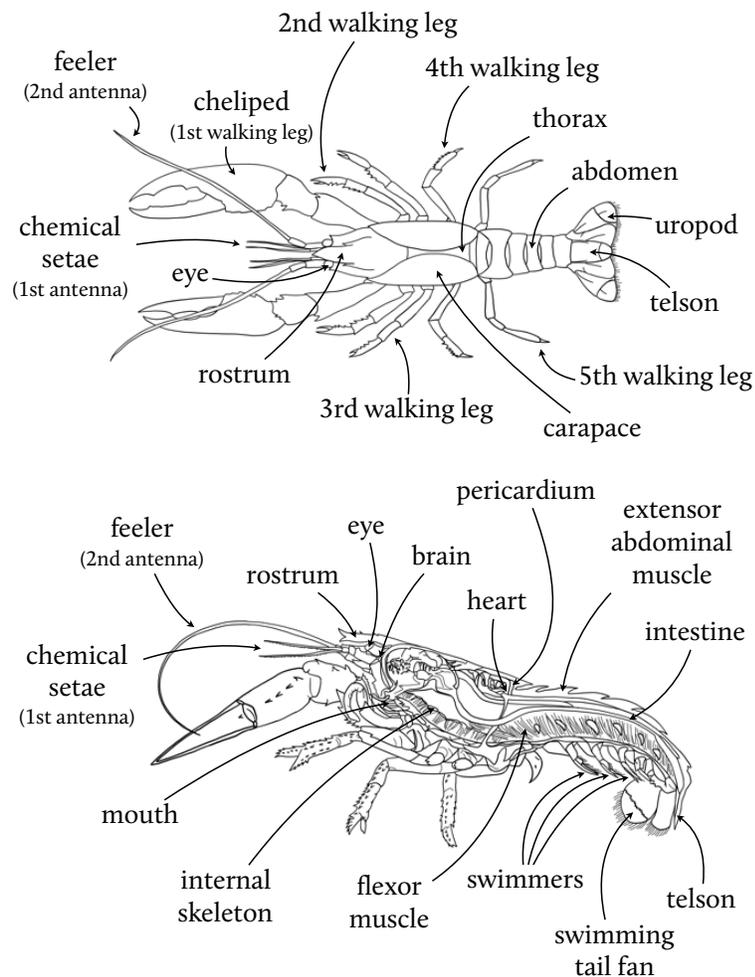


Figure 3.5: Diagrams of a lobster from the top (top) and from the side (bottom, copied from the original²⁰) showing different body parts, including those mentioned in this section.

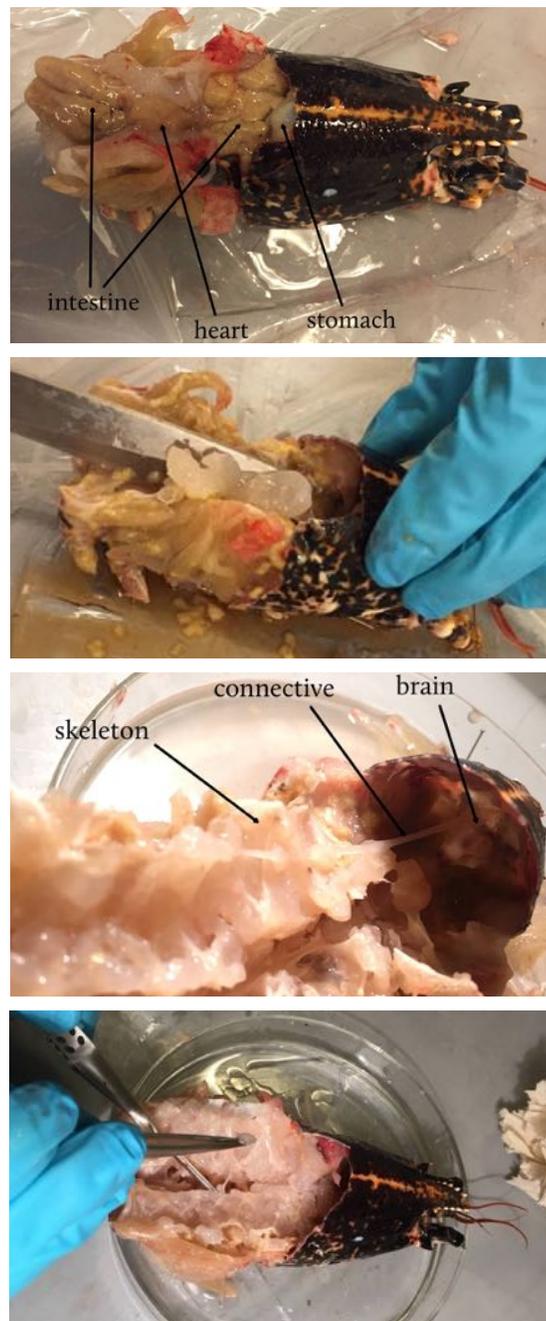


Figure 3.6: Internal organs of the thoracic cavity of the lobster (top) with some organs labelled. Longitudinal cut of the internal skeleton after cleaning of the thoracic cavity (upper centre). Skeleton after longitudinal cut and separation of the halves (lower centre). The ventral cord has been extracted from the skeletal cavity. The subesophageal ganglion is still attached to the inside of the head. Pulling the cord from the skeleton requires patience and care (bottom).



Figure 3.7: Longitudinal cuts down the sides of the abdomen (top) and peeling of the internal skeleton (bottom) to reveal the abdominal section of the ventral cord.

For extraction of the abdominal cord, the swimmers were cut from the abdomen. When this was done, care had to be taken to maintain the fingers away from the underside of the abdomen to prevent them from being trapped by the abdominal carapace in the event that the abdominal flexor muscle contracted and the tail curled up. The carapace was then cut on both sides longitudinally from the front of the abdomen to the uropod using large scissors (figure 3.7, top). The skeleton was then lifted slightly on the front side, and the muscle tissue was cut where it had been joined with the thorax in order to separate it from the skeleton. The skeleton was peeled off (figure 3.7, bottom); this was done very gently to avoid damaging the abdominal cord, which can be adhered to either the skeleton or the muscle. The rest of the abdomen, like the chelipeds, was set apart for cooking.

The extracted parts of the ventral cord were placed in petri dishes with saline solution and stored at 4 °C. They remained excitable for 2 days.

Giant axon extraction

The ventral cord is a collection of ganglia joined by a pair of parallel nerve bundles which are surrounded by a thin sheath. Four giant axons run down each cord from brain to tail: two lateral giant axons and two median giant axons. Figure 3.8 shows a diagram of the complete ventral cord, including a transverse cut where the giant axons can be seen. The lateral giant axons have a larger diameter and are thus easier to isolate and manipulate.

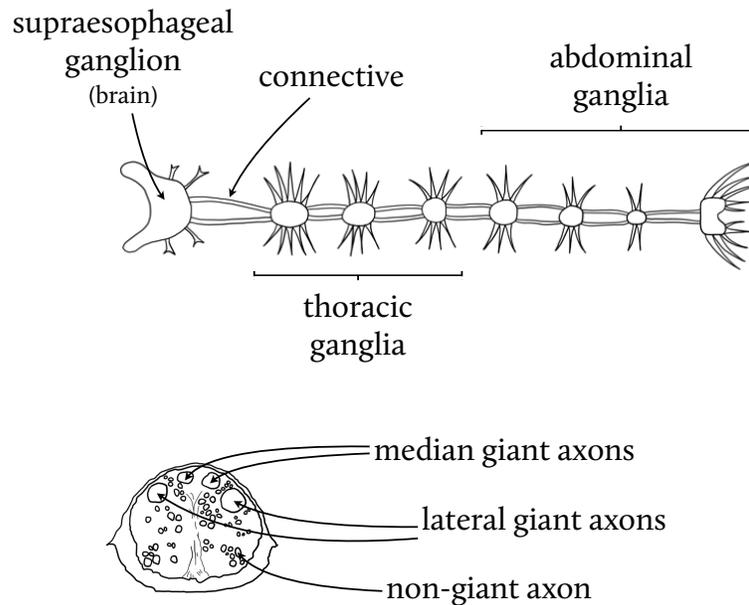


Figure 3.8: Diagram of the complete ventral cord of a lobster (top). The nerve terminals connecting the ganglia to other body parts are shown as protrusions from the ganglia, including the subesophageal ganglion. A transverse cut of the connective (bottom) reveals the internal structure. The exterior cover is a sheath that runs down the connective and protects the nerve bundle from the exterior. The nerve bundle consists of many regular axons and four giant axons. The lateral giant axons are larger than the median ones.

The thoracic ventral cord was cut transversally just below the first thoracic ganglion, leaving only said ganglion, the brain and the two connectives joining them (see figure 3.8). This was for ease of manipulation; this length of giant axon was sufficient for stimulation. The brain, connectives and first ganglion were placed on a petri dish with silicone elastomer, and saline solution was added to the dish to maintain the cells alive during extraction of the giant axon; they were placed upside-down to expose the nerve ends for better visibility. The brain and ganglion were cut in half longitudinally to separate the connectives, taking care not to cut the nerve ends, and one connective was set aside. The remaining brain hemisphere and half-ganglion were pinned right-side-up to the silicone elastomer with either small metallic pins or glass pipette tips (see section 3.5.1 for pipette tip preparation); since the giant axons are on the bottom of the connective, this configuration made it less likely that they would be damaged during removal of the sheath. The petri dish was then placed under a microscope with a low-magnification objective; in this case, a 20 \times objective was used.

Looking at the connective under the microscope, the sheath surrounding the nerve bundle was very gently cut longitudinally with small shears, taking care not to damage the bundle itself. Once a cut from brain to ganglion had been made, a transversal cut was made on the sheath just above the ganglion, any axons attached to the sheath were gently

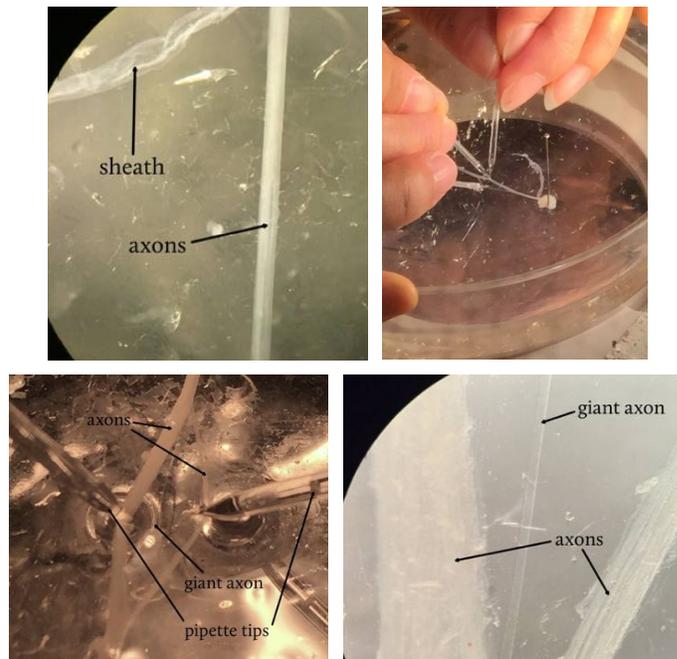


Figure 3.9: Connective as seen under a microscope (top left) and with the naked eye (top right) after partial removal of the sheath. The sheath can be seen as a wrinkled white mass extending away from the brain, which is pinned to the silicone elastomer substrate, in an arc in the top-right photograph. Isolated lateral giant axon as seen with the naked eye (bottom left) and under a microscope (bottom right). The axon is so thick that it can be seen with the naked eye despite being clear.

separated, and the sheath was slowly peeled in the direction of the brain (figure 3.9, top left).

With the sheath removed and the axons exposed, the brain and ganglion were unpinned, turned upside-down so the giant axons would be at the top, and pinned to the silicone elastomer once again. A glass pipette tip was used to gently separate the axons by running the tip longitudinally down the bundle. When one of the lateral giant axons was located, a length of it was separated and the rest of the nerve bundle was held apart from it with pins or pipette tips (figure 3.9, top right). Thus isolated, the giant axon—which, as mentioned in chapter 1, typically has a diameter of a few hundred microns—was visible with the naked eye (figure 3.9, bottom).

The abdominal cord has been used by others¹⁹ but was not used in our case.

Leg nerve extraction

Extraction of the leg nerves was considerably easier than extraction of the giant axons. First, any one of the refrigerated walking legs was placed on a petri dish with silicone elastomer and saline solution, and the top end of the leg was pinned to the silicone. The skin was lifted with small tweezers, and any muscles adhered to it were cut with small

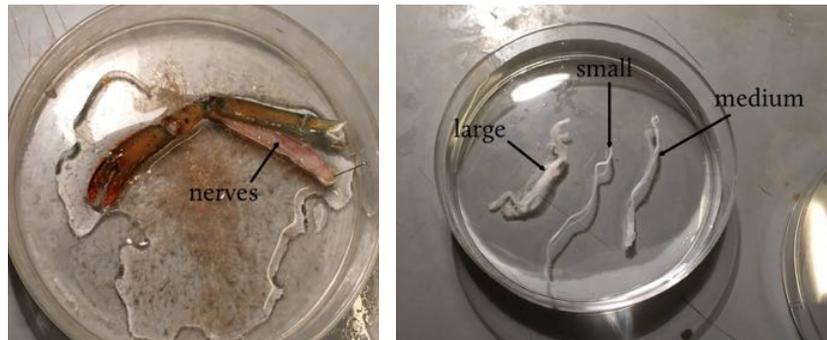


Figure 3.10: Opened leg and its nerves (left). The three nerves, once extracted (right), show pronounced differences in thickness. Each consists of multiple axons.

shears.

The leg has three nerve bundles: a large one, a medium one and a small one (figure 3.10). The nerves were separated from each other using a knife and cut at the joint. They were then placed in a separate petri dish filled with saline solution and stored at 4 °C.

Disposal of unwanted parts and cleaning of the workspace

Unused lobster parts were disposed of as food waste. Surfaces and tools were cleaned with ethanol.

3.2. Quantitative differential interference contrast

3.2.1. Differential interference contrast

Differential interference contrast (DIC) microscopy works by taking advantage of the phase shift a sample causes when light passes through it.

In the simplest form of the technique, light passes through a linear polariser followed by a Nomarski prism, which spatially separates its polarisation components. If we denote by \vec{H} (horizontal) and \vec{V} (vertical) the polarisations into which the prism separates the incident beam, then the polarisation of the light prior to the prism should be at an angle of 45 ° with respect to \vec{H} in order for the two beams separated by the prism to have the same amplitude. These beams, which have mutually orthogonal linear polarisations, are focussed by the condenser onto slightly different positions in the sample, separated by a distance r_0 ; this distance, called the shear, depends on the prism. After passing through the sample and being collimated by the objective, they are incident on another Nomarski prism and recombined. The light is finally transmitted through a linear polariser with its axis perpendicular to the original polarisation of the light, behind which is a detector.

If the beams are incident on the sample at positions $\vec{r} \pm \vec{r}_0/2$ and the sample introduces a position-dependent phase φ in the light that passes through it, then there will be a phase

shift

$$\Delta_\varphi = \varphi\left(\vec{r} + \frac{\vec{r}_0}{2}\right) - \varphi\left(\vec{r} - \frac{\vec{r}_0}{2}\right)$$

between the two orthogonal polarisations. If the sample's optical thickness (the product between its thickness and its refractive index) is the same at the two positions, then $\Delta_\varphi = 0$ and no light reaches the detector. On the other hand, if the optical thickness is different at the two positions, then $\Delta_\varphi \neq 0$, which converts the original linear polarisation of the light into elliptical polarisation, and some light reaches the detector (unless the difference is such that $\Delta_\varphi = 2N\pi$ for some integer N , so the technique is appropriate only for samples with relatively small optical thicknesses). This results in an image which only shows the contours of the structures of the sample, where the sample's optical thickness changes. The name of the technique comes from the fact that the image is produced by the interference of the two orthogonally polarised beams when they are recombined, as well as the fact that these beams come from different points in the sample, which means the contrast is created by differences in the optical thickness of the sample across a small distance; in this sense, the image can be thought of as an approximation of the derivative of the optical thickness in the direction of \vec{r}_0 . (This means that the image contains no information of optical thickness changes in the direction orthogonal to \vec{r}_0 .)

An alternative setup (shown in figure 3.11) is to place a quarter-wave plate with its fast axis at 45° immediately after the first linear polariser; this is the configuration often used

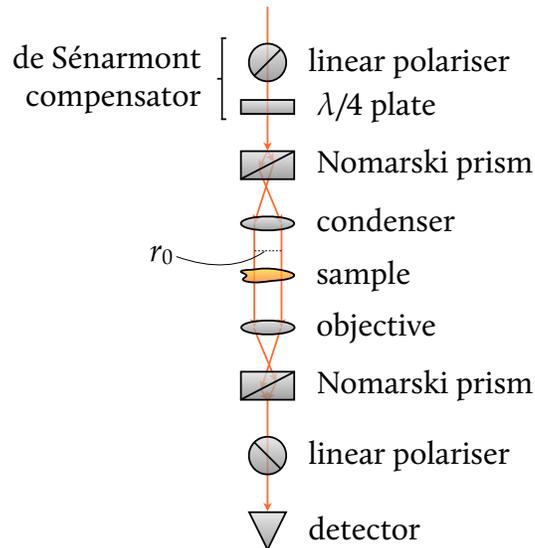


Figure 3.11: Experimental setup used for quantitative DIC imaging. The quarter-wave plate rotation was initially done manually but was eventually controlled by a motor which provided rotation angles between -40° and 40° .

in commercial DIC setups. This combination of linear polariser and quarter-wave plate is called a de Sénarmont compensator. The first polariser can be rotated about the optical axis. The wave plate introduces a phase shift between the polarisations equal to twice the polariser rotation angle ψ , as can be shown by some simple Jones calculus, which consists of writing the polarisation of the beam as a (normalised) 2-dimensional vector (whose components are the complex amplitudes of the polarisation components of the beam) and the effect of each optical component on this polarisation vector as a 2×2 matrix.

If we denote the Jones matrix of the quarter-wave plate by $M_{\lambda/4}$ and the Jones vector of the beam after it has passed through the linear polariser (which, as has been mentioned, is at an angle $\psi + 45^\circ$) by \vec{E}_i , we have²¹

$$M_{\lambda/4} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -i \\ -i & 1 \end{pmatrix},$$

$$\vec{E}_i = \begin{pmatrix} \cos(\psi + \pi/4) \\ \sin(\psi + \pi/4) \end{pmatrix}.$$

The Jones vector of the beam after it has passed through the wave plate is then

$$\vec{E} = M_{\lambda/4} \vec{E}_i = \frac{e^{-i\pi/4}}{\sqrt{2}} \begin{pmatrix} e^{-i\psi} \\ e^{i\psi} \end{pmatrix} \sim \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ e^{2i\psi} \end{pmatrix},$$

where the omission of the global phase factor $e^{-i(\pi/4 + \psi)}$ is due to the fact that global phases do not affect the polarisation.

If the \vec{H} and \vec{V} beams are incident on the sample at $\vec{r} + \vec{r}_0/2$ and $\vec{r} - \vec{r}_0/2$, then the Jones vector of the light after the beams have been recombined by the second Nomarski prism is

$$\vec{E} = \frac{1}{\sqrt{2}} \begin{pmatrix} e^{i\varphi(\vec{r} + \vec{r}_0/2)} \\ e^{2i\psi + i\varphi(\vec{r} - \vec{r}_0/2)} \end{pmatrix} \sim \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ e^{i(2\psi - \Delta\varphi)} \end{pmatrix}.$$

The Jones matrix of a polariser orthogonal to the first polariser (i.e. at -45°) is

$$M_{-45^\circ} = \frac{1}{2} \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}.$$

Therefore, the Jones vector of the light that reaches the detector is

$$\vec{E} = \frac{1}{\sqrt{2}} M_{-45^\circ} \begin{pmatrix} 1 \\ e^{i(2\psi - \Delta\varphi)} \end{pmatrix} = \frac{1 - e^{i(2\psi - \Delta\varphi)}}{2\sqrt{2}} \begin{pmatrix} 1 \\ -1 \end{pmatrix}$$

and the detected intensity is

$$I = I_s \left| \frac{1 - e^{i(2\psi - \Delta\varphi)}}{2\sqrt{2}} \begin{pmatrix} 1 \\ -1 \end{pmatrix} \right|^2 = \frac{I_s (1 - \cos(2\psi - \Delta\varphi))}{2}, \quad (3.1)$$

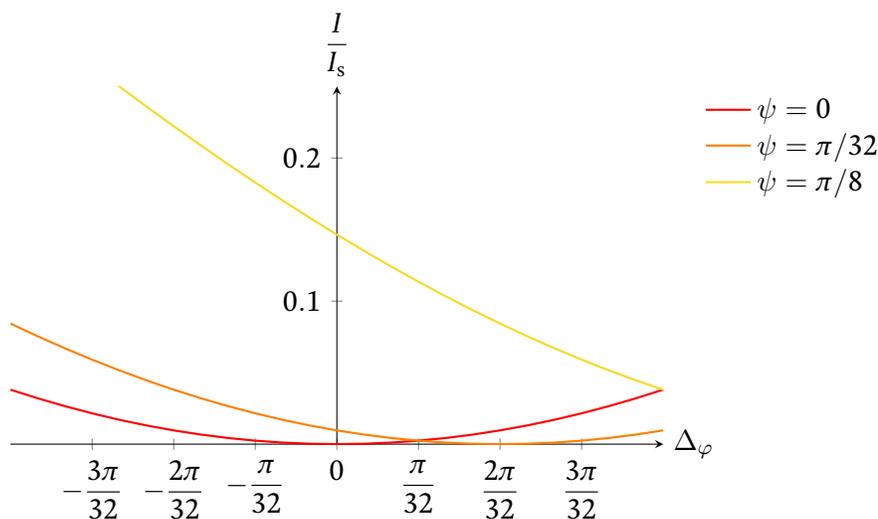


Figure 3.12: Normalised DIC intensity as a function of Δ_φ for different values of ψ . For $\psi = 0$ (red curve), the intensity carries no information about the sign of Δ_φ . If the maximum of $|\Delta_\varphi|$ is, for example, $\pi/8$, then $\psi < \pi/16$ (orange curve) results in the intensity being degenerate for some values of Δ_φ . If ψ is equal to or greater than this maximum phase (yellow curve), then the intensity takes a different value for each value of Δ_φ throughout the range of Δ_φ .

where I_s is the intensity of the beam after it leaves the sample.

The purpose of the wave plate now becomes clear. If $\psi = 0$, there is no retardation and the contrast depends solely on Δ_φ ; however, the intensity carries no information about the sign of Δ_φ . If, on the other hand, ψ is an angle between 0 and $\pi/4$, then I will be nonzero for $\Delta_\varphi = 0$, will increase for negative values of Δ_φ and will decrease for positive values of Δ_φ (figures 3.12 and 3.13). Of course, in order for there to be a one-to-one relationship between I and Δ_φ , we must set ψ to a value such that $0 \leq 2\psi - \Delta_\varphi \leq \pi/2$ for all Δ_φ in the field of view.

Our setup (figure 3.11) consisted of an inverted microscope (Nikon Ti-U) with a water-immersion $60\times$ objective with numerical aperture 1.27 and a $1.5\times$ tube lens. Images were taken by a PCO.edge 5.5 digital camera capable of taking images 2,560 px wide and 2,160 px tall; the pixel size was $6.5 \mu\text{m}$, which corresponds to 72.22 nm at the sample with the magnification. Exposure times were 50 ms, and images were averaged either 128 or 256 times to increase the signal-to-noise ratio. The condenser lamp emitted white light which passed through a green interference filter, which resulted in the mean wavelength being 530 nm and the full width at half maximum being 70 nm . The Nomarski prisms (Nikon N2) provided a shear of $(238 \pm 10) \text{ nm}$.²² The quarter-wave plate was initially rotated by hand, but it was eventually motorised, providing rotation angles between -40° and 40° ; when the rotation was manual, one image (either $I(\psi)$ or $I(-\psi)$, depending on the last angle used) was taken, the average number of electrons per pixel was recorded, and the quarter-wave plate was rotated to an angle with the opposite sign until the aver-

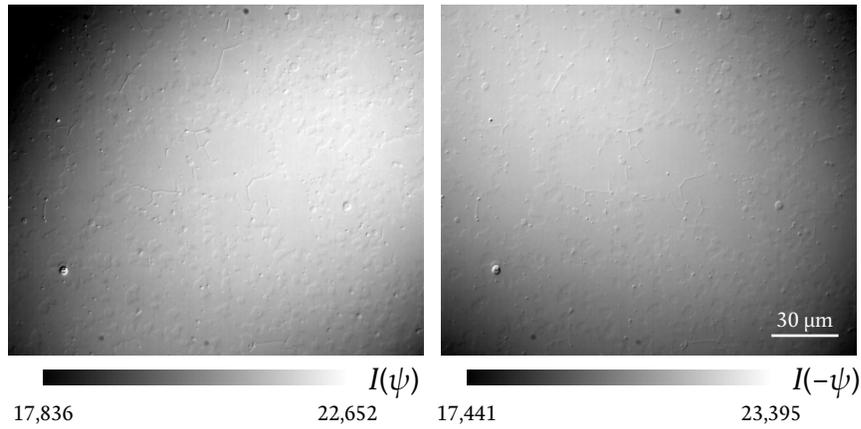


Figure 3.13: Differential interference contrast images of a supported DC₁₅PC bilayer. The quarter-wave plate rotation angle is $\psi = 12.9^\circ$. The shear direction is from top left to bottom right. The grey range (in electrons) has been reduced in both images to increase the contrast; the full ranges are 6,960–27,895 for $I(\psi)$ (left) and 4,149–30,528 for $I(-\psi)$ (right). Ideally, the range would be the same for the two images; the small deviation from this ideal situation was due to the fact that the average number of electrons per pixel, not the maximum signal, was used to determine when the negative rotation angle was equal in magnitude to the positive rotation angle (see section 3.2.2).

age number of electrons per pixel was equal to the recorded number to within about 30 electrons. Typical average electrons per pixel were around 20,000. Images taken before the motor was installed used $\psi = 12.9^\circ$, and images taken after it was installed used $\psi = 15^\circ$.

3.2.2. Making DIC quantitative

DIC microscopy is not quantitative because I_s is unknown; this is due to spatial inhomogeneities in the illumination intensity and detector efficiency, as well as to spatial artefacts caused by slight misalignment or imperfections of the optical components of the setup. To make it quantitative, we must remove the contribution of I_s to I .

This can be done by imaging the sample with a quarter-wave plate rotation angle $-\psi$ as well (figure 3.13). We may then calculate the contrast

$$I_c = \frac{I(\psi) - I(-\psi)}{I(\psi) + I(-\psi)}.$$

Using equation 3.1, this can be written in terms of ψ and Δ_φ :

$$I_c = \frac{\cos(2\psi + \Delta_\varphi) - \cos(2\psi - \Delta_\varphi)}{2 - \cos(2\psi + \Delta_\varphi) - \cos(2\psi - \Delta_\varphi)} = -\frac{\sin(2\psi) \sin(\Delta_\varphi)}{1 - \cos(2\psi) \cos(\Delta_\varphi)}.$$

Solving for Δ_φ , we obtain

$$\Delta_\varphi = -\arcsin\left(I_c \sin(2\psi) \frac{1 \pm \sqrt{1 - \sin^2(2\psi) - I_c^2 \cos^2(2\psi)}}{\sin^2(2\psi) + I_c^2 \cos^2(2\psi)}\right). \quad (3.2)$$

3.2.3. Obtaining the optical thickness of a sample

Recall that DIC gives an approximation of the derivative of the sample's optical thickness. Therefore, it is natural to think of integrating the intensity to retrieve the optical thickness itself at each point on an image.

To recover the phase φ at each point of the image, we first note that Δ_φ can be written as the convolution of $\varphi(\vec{r})$ with two Dirac delta "functions" δ , one centred at $\vec{r} + \vec{r}_0/2$ and the other centred at $\vec{r} - \vec{r}_0/2$:

$$\begin{aligned} \Delta_\varphi &= \varphi\left(\vec{r} + \frac{\vec{r}_0}{2}\right) - \varphi\left(\vec{r} - \frac{\vec{r}_0}{2}\right) \\ &= \int \varphi(\vec{\rho}) \delta\left(\vec{\rho} - \vec{r} - \frac{\vec{r}_0}{2}\right) d\vec{\rho} - \int \varphi(\vec{\rho}) \delta\left(\vec{\rho} - \vec{r} + \frac{\vec{r}_0}{2}\right) d\vec{\rho} \\ &= \int \varphi(\vec{\rho}) \delta\left(\vec{r} + \frac{\vec{r}_0}{2} - \vec{\rho}\right) d\vec{\rho} - \int \varphi(\vec{\rho}) \delta\left(\vec{r} - \frac{\vec{r}_0}{2} - \vec{\rho}\right) d\vec{\rho} \\ &= \varphi(\vec{r}) * \left(\delta\left(\vec{r} + \frac{\vec{r}_0}{2}\right) - \delta\left(\vec{r} - \frac{\vec{r}_0}{2}\right)\right), \end{aligned}$$

where $*$ denotes the convolution of two functions and we have used the evenness of δ for the third equality. The Fourier transform of this is

$$\mathcal{F}[\Delta_\varphi] = \left(e^{i\pi\vec{\xi}\cdot\vec{r}_0} - e^{-i\pi\vec{\xi}\cdot\vec{r}_0}\right) \mathcal{F}[\varphi] = 2i\sin\left(\pi\vec{\xi}\cdot\vec{r}_0\right) \mathcal{F}[\varphi],$$

where $\vec{\xi}$ is the spatial frequency of the image.

For (discrete and finite-sized) experimental data, we must take the discrete Fourier transform or the fast Fourier transform. One drawback of this is that the transform assumes the data is periodic, which is not the case. This results in artefacts, in part due to discontinuities at the edges of the image. To address this problem and reduce the edge artefacts, the images are subjected to two procedures.

In the first one, a second-order polynomial is fitted to each image and subtracted from it. This sets the signal where there is no contrast to zero.

In the second one,²³ each image is padded with arrays of the same size of the image in both x and y to triple its dimensions, after which the enlarged image is apodised as follows. The value of the pixel at each corner of the original image is extended outwards in the x and y directions all the way to the edge corresponding edges of the enlarged image. The remainder of the edge pixels are extended outwards in the direction perpendicular to the corresponding edge, and this extended data is convolved with a normalised

gaussian function. The gaussian's standard deviation σ is equal to $ws(8\ln(2))^{-1/2}$, where w is the distance from the edge of the original image and s is a smoothing factor. This blurs the top, bottom, left and right sections of the enlarged image increasingly with distance from the original image. For each corner section, the value at each (x, y) is given by $\cos^2(\theta)X + \sin^2(\theta)Y$, where X is the value of the edge pixel of the vertically adjacent section whose horizontal coordinate is x , Y is the value of the edge pixels of the horizontally adjacent section whose vertical coordinate is y , and θ is the angular coordinate of the point (x, y) taking the origin as the point the corner section shares with the original image and with θ increasing from the vertical edge ($\theta = 0$) to the horizontal edge ($\theta = \pi/2$); see figure 3.14 for a diagram of these quantities. Finally, the array is multiplied by a broad \cos^2 function of the distance from the centre, resulting in the central image being unaffected and the outer edges of the extended image becoming zero, as well as their first derivative. This effectively serves to push the edges much farther out, confining any edge artefacts (which should be small anyway because the intensity at the edges of the enlarged image is zero) to a region far from the original image.

Dividing by

$$K = 2i \sin(\pi \vec{\xi} \cdot \vec{r}_0)$$

and taking the inverse Fourier transform of the result would yield φ , which is the quantity

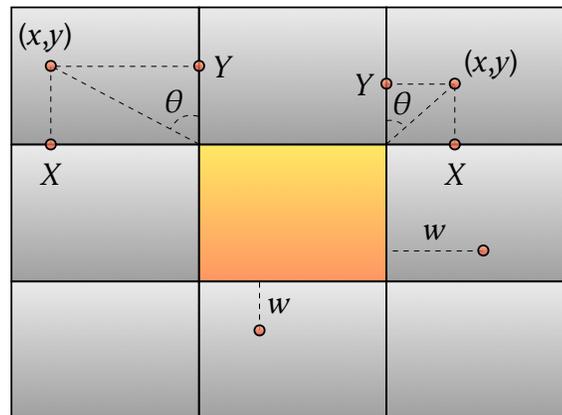


Figure 3.14: Diagram illustrating the meaning of the various quantities used in the apodisation process. The yellow section is the original image. Together with the original image, the grey sections (which have the same dimensions as the original image) form the extended image. For the top, bottom, left and right sections, w is the distance from the nearest edge of the original image (indicated here for points in the bottom and right sections). For the corner sections, given a position (x, y) , X and Y are the values of the edge pixels of the adjacent sections which share a coordinate with that position, and θ is the position's angular coordinate measured from the vertical edge of the corresponding adjacent section (indicated here for points in the top-left and top-right sections).

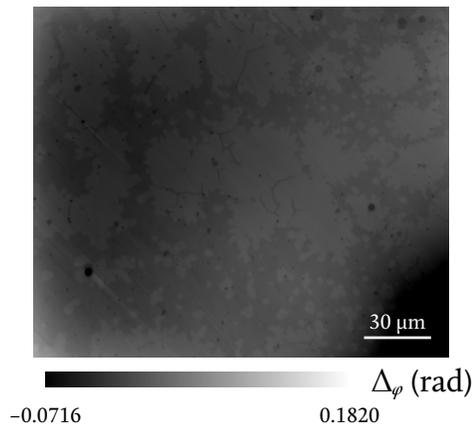


Figure 3.15: Integrated phase (Δ_φ) image after Wiener deconvolution. The bilayer appears as a dark, fractal-like structure, while lighter regions are empty (these regions will henceforth be referred to as “no-layer regions”). There are artefacts in the form of lines parallel to the shear; these are the result of the lack of information about the sample in the direction perpendicular to the shear.

we seek, but $K = 0$ when $\vec{\xi} \cdot \vec{r}_0 \in \mathbb{Z}$. Instead, we multiply $\mathcal{F}[\Delta_\varphi]$ by the Wiener filter function²³

$$W = \frac{1}{K + \frac{1}{\kappa K^*}},$$

where κ is the estimated signal-to-noise ratio of the image and serves to limit the amplification of any noise in the resulting image.²³ Only after multiplying $\mathcal{F}[\Delta_\varphi]$ by W do we take the inverse Fourier transform to retrieve φ (figure 3.15). This process is called Wiener deconvolution.²⁴

After Wiener deconvolution, the padding is removed, leaving only the original image.

The sample’s optical thickness is the product of its thickness d and its refractive index n_s . If the light’s wavelength is λ , then the sample’s optical thickness at any position \vec{r} is

$$dn_s = \frac{\lambda\varphi(\vec{r})}{2\pi}. \quad (3.3)$$

3.3. Interferometric reflectometry

3.3.1. Microscope and optics

The experimental setup used for interferometric reflectometry is shown in figure 3.16 and described below. The theory is described in chapter 4.

A 100-fs pulsed 820-nm titanium-sapphire laser (Spectra-Physics Mai Tai, not shown in the figure) excites an optical parametric oscillator (OPO, Inspire Radiantis), which emits

550-nm 150-fs pulses with an approximately gaussian temporal profile at a rate of 80 MHz. The polarisation is initially horizontal (\vec{H}). The beam enters an acousto-optic modulator (AOM); only the zeroth- and first-order diffracted beams are of interest here.

The first-order beam passes through a quarter-wave plate and a half-wave plate before entering the inverted microscope (Nikon Ti-U) and being reflected towards the objective and the sample by total internal reflection in a prism. Since total internal reflection in the prism changes the relative phase between the polarisation components of the beam and the last mirror before the microscope is actually a dichroic beam splitter (the transmitted light is used for alignment), which also changes the relative phase, this combination of wave plates is used to ensure that the polarisation is left circular (\vec{L}) at the sample. The reflected beam, which (assuming the sample is not birefringent) has right circular (\vec{R}) polarisation, travels back through the same components, which turn its polarisation into vertical (\vec{V}). Because this beam interacts with the sample, we call it the probe beam and denote its field by \vec{E}_p . The power at the sample is typically 10–100 μW . The tube lens provided a $1.5\times$ magnification, and the objective (Nikon CFI Plan Apo) was a $60\times$ water-immersion objective with a numerical aperture of 1.27.

The zeroth-order beam travels down a long path (in order to compensate for the ex-

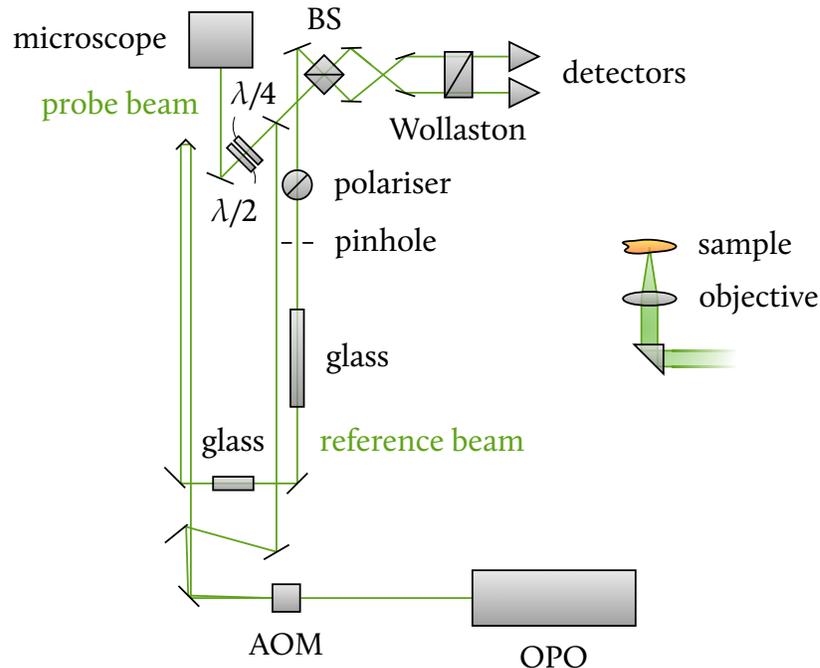


Figure 3.16: Simplified schematic of the experimental setup used for interferometric reflectometry (top view). Not shown are the exciting laser, filters and beam expanders used. The inset on the right (side view) shows the beam being reflected upwards by a prism and focussed onto the sample by the microscope objective. Not shown is the tube lens of the microscope.

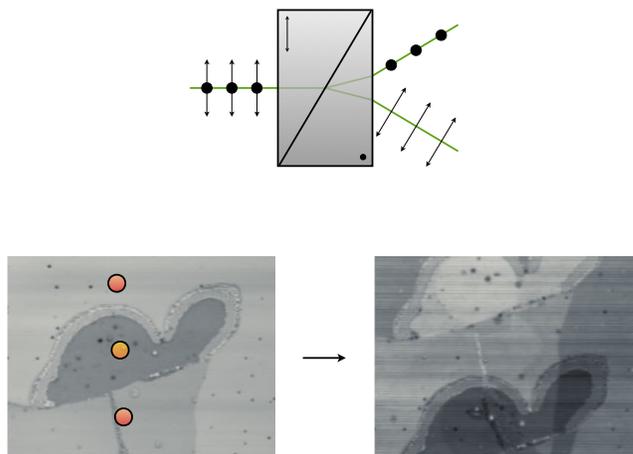


Figure 3.17: Schematic of a Wollaston prism (top) and its effect on imaging when inserted close to the objective's back focal plane (bottom). A Wollaston prism consists of two birefringent triangular prisms adhered to each other so they form a rectangular prism. The optic axes of the two triangular halves are perpendicular with respect to each other. When used in such a way that the direction of propagation of a light beam is perpendicular to both halves' optic axes (and so the axes are parallel to the two polarisation components of the beam), refraction at the boundary between the prisms results in a spatial separation of the two polarisation components. This, in turn, results in two orthogonally polarised beams being incident on the sample at positions (represented by red circles) symmetric about the position (represented by a yellow circle) at which the beam would be incident on the sample if the Wollaston prism were not there, creating a double image. The image is of a ternary lipid bilayer (DOPC:sm:ch = 11:5:4).

tended path followed by the probe beam). The glass blocks in the beam path match the chirp of this beam to that of the probe beam, which is caused by the glass components of the microscope (see section 3.4.1). The beam passes through a linear polariser, which turns its polarisation into linear at 45° , before recombining with the reflected probe beam at a non-polarising beam splitter. Because this beam does not interact with the sample, we call it the reference beam and denote its field by \vec{E}_r .

Before reaching the detectors, the combined beams pass through a Wollaston prism. The effect of this is to spatially separate the horizontal and vertical polarisation components of the interfering fields. We finally have four detectors, two for each linear polarisation, arranged in a square configuration; the detectors are balanced photodiodes connected to a lock-in amplifier (Zurich Instruments HF2). The beam splitter introduces a phase shift of $\pi/2$ in the reflected components of the beams, which for one of the detectors corresponding to each polarisation is the probe beam and for the other detector is the reference beam, resulting in a total difference of π between the components in the two detectors. This allows the detector pair corresponding to each polarisation (\vec{H} and \vec{V}) to use the beam interference to obtain the amplitude and phase, rather than merely the intensity, of the reflected field.

Where an additional Wollaston prism was placed in the probe beam's path (see fig-

ure 3.17) for referencing (see section 4.2.2), it was placed close to the back focal plane of the objective. Three prisms with different shear angles (0.25° , 0.5° and 1°) were available. Ultimately, only the 0.5° prism was used. The shear distance of this prism with the $60\times$ objective was measured to be approximately $35\ \mu\text{m}$. The shear of the 0.25° prism was measured as $12\ \mu\text{m}$, which was too close to the size of many of the imaged features and was thus too small to be useful, and the shear of the 1° prism was calculated to be about $70\ \mu\text{m}$, which was approximately the image size and was thus too large to be useful. The prisms were set into gears attached to rotating motors controllable through a computer, and the gears and motors were mounted on custom-built sliders compatible with the microscope's objective slider slots normally used for DIC imaging.

3.3.2. Detection

In addition to modulating the amplitude of the probe and reference beams, the AOM increases the optical frequency of the probe beam by 82 MHz. The lock-in amplifier is set to 2 MHz, the difference between the laser pulse rate and the AOM upshift, so it amplifies the interference of the probe and reference beams and filters out any detected signal at other frequencies, which is merely noise; in this way, most noise is discarded.^{25,26}

Due to the aforementioned relative phase shift of π introduced by the beam splitter, the signals S_1 and S_2 at the two detectors corresponding to each polarisation are

$$S_1 \sim \left| \frac{E_r + E_p}{\sqrt{2}} \right|^2$$

and

$$S_2 \sim \left| \frac{E_r - E_p}{\sqrt{2}} \right|^2.$$

The balanced detector combines these signals into $S_1 - S_2 = E_p E_r^* + E_p^* E_r = 2\text{Re}(E_p E_r^*)$ and provides this signal to the lock-in amplifier, from which the amplitude and phase of the probe beam can be obtained given that the amplitude and phase of the reference beam are known. This way, we have both amplitude and phase for each polarisation.

3.3.3. Image acquisition

At any given time, only the light reflected by the region of the sample which is being illuminated by the laser is detected; this is not a wide-field imaging technique. The confocal nature of the setup means that either the sample or the laser must be scanned in order to produce an image. In our case, the sample stage, which was motorised, was scanned.

Each pixel corresponds to roughly 108 nm in every image we took; this pixel size was selected in order to have two points per resolution, which was $\lambda/2\text{NA} \approx 217\ \text{nm}$. The acquisition time was 200 μs per pixel. Images were square and between $50\ \mu\text{m}$ and $80\ \mu\text{m}$ on a side. The acquisition time was thus 43–110 s per image.

Immediately before each reflectometry image was taken, the region of interest was imaged using quantitative DIC. This allowed comparisons between the reflectometry images and an already established technique and made it easier to determine what each region in an image corresponded to (e.g. different domains in a lipid bilayer).

3.3.4. Regularisation

Because the position of the scanning stage is not perfectly accurate, the reflectometry data was improved via a process called regularisation. For each pixel of an image, the actual position \vec{r}' of the stage was recorded, along with the time t since the start of the scan; the nominal position of the stage, \vec{r} , was taken as the centre of the corresponding pixel. Then, for every point \vec{r} , the signal $S(\vec{r})$ was calculated as

$$S(\vec{r}) = \sum_{t'} S'(t') W(|\vec{r} - \vec{r}'|),$$

where $S'(t')$ is the detected signal at time t' , W is a weighting function which could be a gaussian function or a step function, and the sum was taken over all the time points t' such that $|\vec{r}(t) - \vec{r}'(t')|$ was smaller than twice the size of one pixel (in the case of the step function) or twice the standard deviation (in the case of the gaussian function; see figure 3.18). In our case, W was a gaussian function:

$$W(|\vec{r} - \vec{r}'|) = e^{-\frac{|\vec{r} - \vec{r}'|^2}{\delta^2}},$$

where δ is the spatial pixel size.

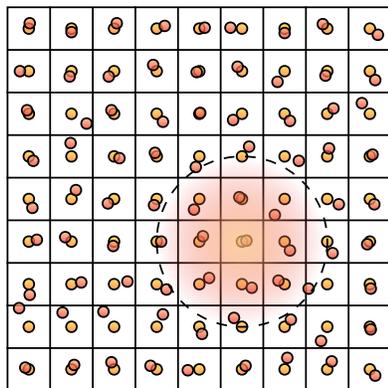


Figure 3.18: Schematic illustrating the principle of regularisation in the case in which the weighting function W is a gaussian. Each square is a pixel. The yellow circles, which lie at the centre of their pixels, indicate the nominal coordinates $\vec{r} = (x, y)$ of the corresponding pixel. The red circles are the actual coordinates \vec{r}' of the points of the measurement. The measured signal at the red points within the dashed circumference is used in the calculation of the signal value at the pixel the circumference is centred on; the shading in the circle represents the weight.

3.4. Coherent Raman scattering microscopy

3.4.1. Microscope and optics

As in the case of reflectometry, the experimental setup used for coherent Raman scattering¹ (figure 3.19) employed a pulsed titanium-sapphire laser emitting 820-nm, 150-fs pulses at 80 MHz. In this case, this was used both to pump the OPO and as the pump beam for SRS and CARS; the OPO idler beam was used as the Stokes beam. The Stokes beam, which was tunable, was set to 1,070 nm, and a spectral range of 2,700–3,100 cm^{-1} was obtainable through spectral focussing,^{28–30} which was achieved by chirping the beams with glass blocks so the pulse duration increased to 1–2 ps; the spectral resolution was 30 cm^{-1} . The amplitude of the Stokes beam was modulated by an acousto-optic modulator with a 2.5-MHz square wave.

Chirping is created by dispersion. A pulse with a finite duration is not monochromatic,

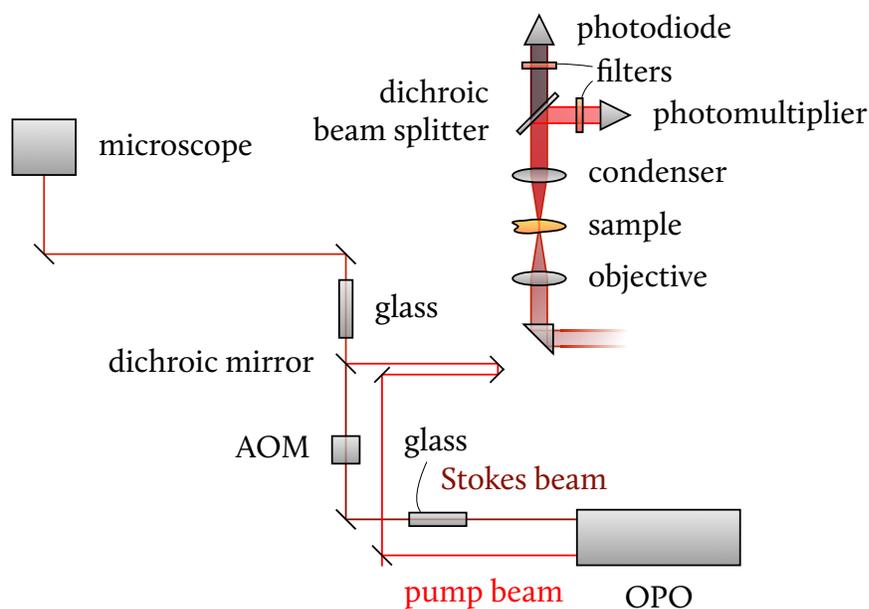


Figure 3.19: Simplified schematic of the experimental setup used for coherent Raman scattering microscopy (top view). Not shown are the exciting laser and beam expanders used, as well as wave plates to fine-tune the pump beam polarisation. The inset on the right (side view) shows the beam being reflected upwards by a prism and focussed onto the sample by the microscope objective. The coherent Raman signals are collected in the forward direction; they are separated by a dichroic beam splitter, since the CARS signal (detected by a photomultiplier tube) is on the anti-Stokes side and the SRS signal (detected by a photodiode) is detected as pump loss. Not shown is the tube lens of the microscope and the lenses required to image the back focal plane of the condenser onto the detectors.

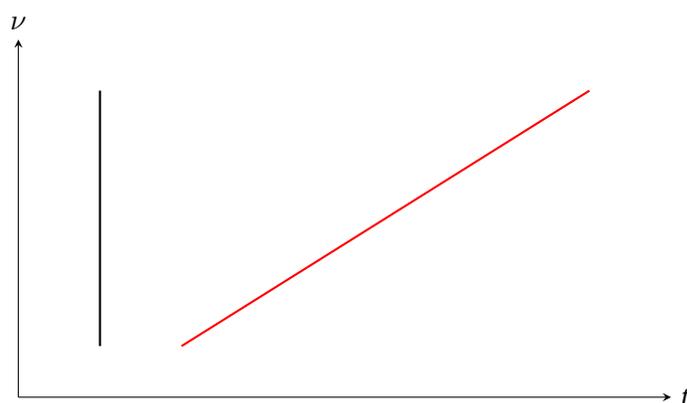


Figure 3.20: Illustration of the effect of chirping. Dispersion from a material (such as glass) causes each spectral component of a pulse to arrive at any given point in the optical path at a different time (red line) instead of all of them arriving at the same time (black line).

but has a spectral width inversely proportional to the pulse duration. Since the refractive index of any material depends on the wavelength, the different spectral components of a pulse will travel at different speeds inside the material. This is called dispersion. Most materials have a refractive index which is a decreasing function of wavelength, so higher-frequency components travel more slowly. This causes a temporal delay at the material's output which is a function of the frequency $\nu = c/\lambda$ (figure 3.20), where c is the speed of light in vacuum. This is called chirping. In coherent Raman microscopy, the pump and Stokes beams need to be equally chirped in order for the frequency difference between them to be the same throughout the pulse; therefore, glass blocks are used to compensate the different chirp caused by the glass components of the setup.

The sample was mounted on the same inverted microscope used for reflectometry and quantitative DIC, and simultaneous SRS (measured as stimulated Raman loss, which is a decrease in the pump intensity; see section 1.3.2) and CARS imaging was performed at room temperature. The microscope optics, as before, consisted of a $60\times$ water-immersion objective whose numerical aperture was 1.27 and a $1.5\times$ tube lens. At the sample, the pump and Stokes powers were, respectively, 4.6 mW and 11.3 mW.

3.4.2. Image acquisition

$15\text{-}\mu\text{m}$ square regions were imaged by raster-scanning the pump and Stokes beams via a galvanometric scanning mirror and detecting the signal in the forward direction. Pixel exposure time was 1 ms, and the spatial step size was 108 nm. The wave number was varied by delaying one of the beams with respect to the other; since they were chirped, this delay resulted in a change in the energy difference between the pump and Stokes spectral components arriving at the sample at any given time. The spectral range was a 200-cm^{-1} range which was centred in the CH_2 stretch region around $2,850\text{ cm}^{-1}$, and the spectral

step size was 5 cm^{-1} ; the acquisition of a full hyperspectral image took around 13 min.

As in the case of interferometric reflectometry, the region of interest was imaged with quantitative DIC immediately before each SRS image was acquired. In this case, the quantitative DIC images served as confirmation for the conclusions drawn from the SRS images (see chapter 5).

Also as with reflectometry, the images were regularised. Further image processing was developed as part of this work and is described in section 5.1.

3.5. Electrophysiology

3.5.1. Pipette tip pulling

150-mm glass pipettes with an outer diameter of 1 mm and an inner diameter of 0.6 mm were used. A microelectrode puller (Narishige PP-830) was used to pull the pipettes in the 2-step configuration. The heat setting was 78.2 for the first step and 62 for the second step. Pipettes were not placed symmetrically in the puller, but rather such that the two resulting pipette tips had a length ratio of 3:1; the smaller pipette tips were discarded. Pipette tips were not polished.

3.5.2. Intracellular and extracellular solution preparation

Intracellular solution

Intracellular solution is an attempt at matching the composition of cytoplasm. It was prepared by mixing 10-mM NaCl, 117-mM KCl, 2-mM MgCl_2 , 1-mM CaCl_2 (Sigma-Aldrich), 2-mM adenosine triphosphate disodium salt ($\text{Na}_2\text{-ATP}$, Sigma-Aldrich), 2-mM guanosine triphosphate sodium salt (Na-GTP , Sigma-Aldrich), 1.2-mM Na_2 -phosphocreatine (Sigma-Aldrich), 11-mM ethyleneglycol-bis(β -aminoethylether)-N,N,N',N'-tetraacetic acid (EGTA, Fisher Scientific) and 11-mM 4-(2-hydroxyethyl)-1-piperazineethanesulfonic acid (HEPES, VWR International) in distilled water. The pH of the solution was adjusted to 7.2 by adding a small amount of potassium hydroxide (KOH, Fisher Scientific).

The solution was sterilised by passing it through a $0.45\text{-}\mu\text{m}$ filter inside a fume cupboard.

When not in use, intracellular solution was stored at -20°C , which is below its freezing temperature. A small amount of solution was thawed immediately before use.

Extracellular solution

Extracellular solution was prepared by mixing 135-mM NaCl, 5-mM KCl, 1.2-mM MgCl_2 , 1.25-mM CaCl_2 , 10-mM D-glucose (Fisher Scientific) and 5-mM HEPES in distilled water.

When not in use, extracellular solution was stored at -4°C .

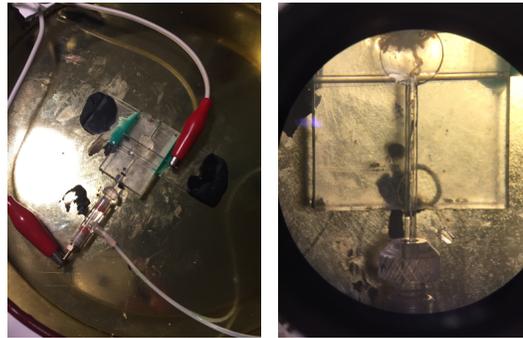


Figure 3.21: Two different chambers employed for extracellular electrophysiology on lobster leg nerve.

3.5.3. Lobster ventral nerve cord

There was insufficient time to perform electrophysiology experiments on individual lobster giant axons. However, a preliminary experiment was performed on a whole lobster connective section running from just below the brain to the first ganglion.

The chamber used for lobster nerve electrophysiology (figure 3.21, left) consisted of two stacked microscope glass slides glued to each other in such a way that their longest sides were perpendicular to each other. The top slide had a wide groove in which the nerve bundle could be placed, as well as two cylindrical tunnels drilled through the centre from end to end. A steel pin, used as the stimulating electrode, was passed through one of these tunnels and glued in place; the glue also sealed the tunnel. Through the other tunnel was passed a glass pipette tip with a silver or platinum wire inserted through it; this was used as the recording electrode. The chamber was filled with extracellular solution, and the nerve bundle was placed in the wide groove in such a way that the end away from the ganglion was touching the mouth of the pipette tip. The chamber was placed in a cylindrical metallic box, which acted both as Faraday cage and as ground (the ground electrode was connected to the cage). The electrodes were connected to an amplifier and function generator via alligator cables. The pipette tip was filled with extracellular solution and passed through a micromanipulator head piece (see section 3.5.4); the head port of the head piece was connected to a screw onto which the recording electrode alligator cable was clamped, while the side port was connected via a flexible plastic tube to a syringe partially filled with air. The nerve was gently sucked into the pipette tip (figure 3.21, right) with the syringe to create a better contact; the amount sucked in was enough for the nerve section at the pipette tip mouth to have the same diameter as the mouth, creating a watertight seal.

3.5.4. Mouse neurons

Coverslips with mouse neurons were mounted on a CSC-13 chamber for round coverslips (Bioscience Tools). This is a circular chamber with an annular metallic base and a silicone

ring clamped down by a second metallic ring to form a seal at the edge of the coverslip. The chamber was mounted on a custom-built annular platform with rectangular shelves which were screwed to the microscope sample stage in place of the usual square platform for rectangular slides. This construction allowed the same microscope which was used for all the other techniques to be used for electrophysiology. The chamber was filled with culture medium, which was replenished periodically to counteract evaporation (see section 4.4.2).

Two micromanipulators (Scientifica PatchStar) were mounted on the microscope stage, one on either side of the microscope (figure 3.22). These were controlled by a single control cube capable of moving one micromanipulator at a time; the precision of the positioning was about $1\ \mu\text{m}$, and the cube had an xyz mode (where the three directions were independently controllable via three knobs on the cube; here, z is vertical and y is perpendicular to the plane in which the axis of the pipette tip is contained) and an approach mode (where the x direction was replaced by movement in the direction on the xz plane parallel to the pipette tip; each micromanipulator contains a sensor which automatically detects its inclination angle and provides the correct approach angle). The micromanipulators and microscope stage were grounded to the optical table. The micromanipulators were connected to an amplifier (Scientifica MultiClamp 700B). A long-working-distance condenser, a $20\times$ dry objective with a numerical aperture of 0.75 and the same $1.5\times$ tube lens as before were employed. Due to the geometry of the setup, pipette tips had to be made about 1.5 times longer than normal so they would reach the sample, since the microscope condenser did not allow the micromanipulators to approach the sample enough if standard-length pipette tips were used (hence the 3:1 ratio used in the pipette pulling).

Each electrode was made by cutting about 9 cm of silver wire with a diameter of 0.25 mm (Goodfellow) and submerging it in bleach. The bleach created an AgCl coating around the wire, preventing Ag^+ ions from contaminating the intracellular solution, which would have been toxic to the cells.³¹ A suitable coating was created within 25 min, but, since AgCl is photosensitive, in practice several 9-cm wires were kept in a bleach bath from one day before measurement to the last day of measurement. A wire typically lasted

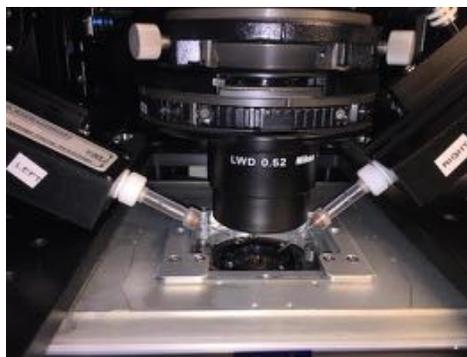


Figure 3.22: Arrangement of the two micromanipulators at the side of the microscope.

3–5 days with the AgCl coating in good condition, after which it was returned to the bleach bath.

Chlorinated wires were inserted into the micromanipulator heads, after which pipette tips filled with intracellular solution were inserted in the heads in such a way that the wire was inside the pipette tip and partially submerged in the intracellular solution. A 1-ml syringe partially filled with air was connected to the side port of one of the micromanipulators via a 3-way valve (figure 3.23).

A cylindrical plastic incubator cover with a circular glass window on the top section (Bioscience Tools TC-1) was available for heated experiments (figure 3.24). Wound around the cylinder was a poorly conducting wire connected to a power source which provided 4.5 V; this allowed heating of the cover and, by conduction, the space inside it. One of the ports on the incubator cover was connected to a plastic centrifuge tube filled with water. Two holes were made in the lid of this tube, one for the tube connecting to the incubator chamber and one for a second tube connected to a CO₂ flow unit, itself connected to a tank which supplied CO₂ at a pressure of about 25 bar. The centrifuge tube was kept in a polystyrene cube with a cylindrical hole made so the tube would fit snugly, and wrapped around the tube was an objective heater (Bioscience Tools MTC-HLS-025) connected to

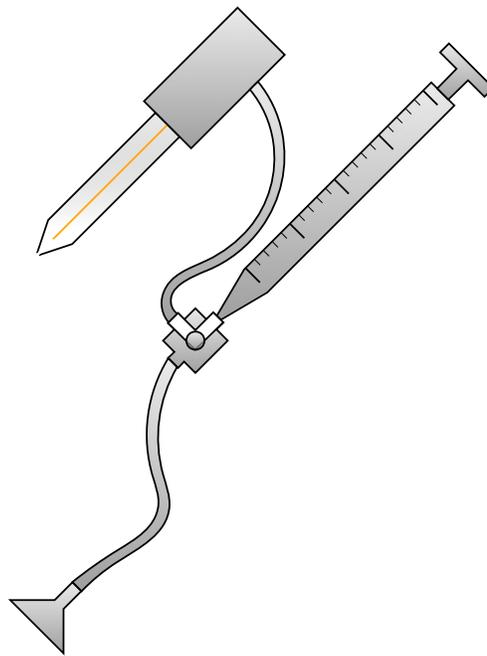


Figure 3.23: Diagram showing the tubing of the pipette tip, syringe and mouthpiece. The pipette tip and micromanipulator head are shown at the top left. A plastic tube connects the pipette tip to a 3-way valve through the micromanipulator head. This valve is also connected to a syringe (top right) and a mouthpiece (bottom left).

one channel of a 2-channel heating control unit (Bioscience Tools TC2-80-150-C). The second port of the incubator cover was left open to act as an air output, allowing the atmosphere inside the incubator chamber to be continually recycled. The other channel of the heating control unit was connected to a second objective heater wound around the microscope objective. This setup permitted the incubator chamber to have a controlled atmosphere with humidified air with a CO₂ concentration of about 5% and a temperature of 37 °C. In order for the pipette tips to have access to the sample, two small holes were drilled on the sides of the incubator lid; the CO₂ loss caused by these additional holes was deemed acceptable, and both the power supply voltage and the temperature setting of the objective channel of the heating control unit could be raised to compensate the additional cooling due to multiple outputs. However, the holes were too small and severely limited the movement of the pipette tips, and larger holes would have resulted in very high CO₂ and temperature losses, so this part of the setup was abandoned; all experiments were carried out at room temperature (about 20 °C). The controlled atmosphere in the incubator would have lengthened the lifespan of the neurons mounted on the microscope, but without it their lifespan was about 2 hr, which was long enough to conduct experiments.

The voltage and current, which were controlled by the MultiClamp software, were set to voltage-clamp mode, which monitors the resistance at the tip of the electrode and controls the current in order to maintain the voltage at a constant level; both electrodes were set to 0 mV.

A small positive pressure was applied to one of the pipette tips, whose electrode we call the recording electrode because it is the one which was in contact with the cells, by injecting about 0.2 ml of air with the syringe; this was to push away any debris floating in the culture medium and prevent them from attaching to the mouth of the pipette tip.

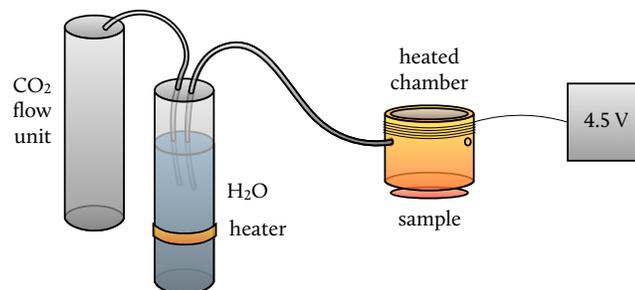


Figure 3.24: Diagram of the incubator cover and its connections. The incubator cover (shown in yellow and labelled “heate chamber”) has two holes, one on either side, as well as a glass window on the top side. One of the holes is open, while the other is connected via a plastic tube to a centrifuge tube containing water (shown on the left). This centrifuge tube is heated by an objective heater (shown in yellow) and connected via a second plastic tube to a CO₂ flow chamber, which is itself connected to a CO₂ tank (not shown). The incubator cover is heated by a poorly conducting wire wrapped around it and connected to a 4.5-V power source (shown on the right).

The syringe valve was then closed.

The pipette tips were submerged in the culture medium and located under the microscope; in order to reduce the likelihood of the pipette tip, which was very fragile, crashing against the coverslip and shattering, the condenser aperture was closed to increase the depth of view and the focal plane was set about 100 μm above the cells. The voltage at the two electrodes was monitored with an oscilloscope. When the pipette tips were submerged, a periodic 15-ms voltage step appeared on both channels of the oscilloscope. The resistance was displayed by MultiClamp; a resistance over 5 $\text{M}\Omega$ indicated that something was blocking the mouth of the pipette tip, while a resistance of less than 3 $\text{M}\Omega$ indicated that the mouth was too wide, usually because it was broken. In either of these cases, the pipette tip had to be replaced.

Once the tips were located (figure 3.25, top), the one not connected to the syringe, whose electrode we call the reference electrode, was brought into focus; this remained 100 μm above the cells throughout the recording and allowed us to determine whether any signal detected by the recording electrode came from the cells or from some external fluctuation.

A voltage offset was applied to the pipette tips with MultiClamp in order to maintain the baseline at 0 mV. The recording electrode was also brought into focus, and it and the

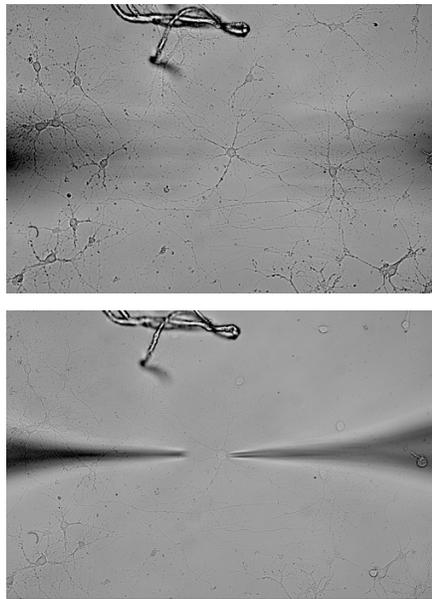


Figure 3.25: Pipette tips out of focus (top, visible as shadows on the left and right sides) and at the recording positions (bottom). The tip on the left contained the reference electrode, and the one of the right contained the recording electrode. Recording was done at the soma due to the increased ease of access compared to the axon; recording at the axon was occasionally attempted, but a good seal could not be created in any case.

focal plane were then slowly and simultaneously lowered until the cells became visible. At this point, a plane slightly above the cells was brought into focus while the pipette tip remained stationary, and a cell was selected. The pipette tip was moved in the x direction until, horizontally, it sat just off the edge of the cell of interest. It then continued to be lowered until it once again came into focus. The micromanipulator control cube was switched to approach mode, and the cell's soma was very slowly approached by the pipette tip until they touched and the cell membrane was slightly deformed. At this point, the voltage step on the oscilloscope decreased in amplitude.

The recording electrode offset was reset so the baseline would be at 0 mV. The pressure was released by turning the syringe valve so the pipette tip was connected to the valve output which was not attached to the syringe; it was attached to a mouthpiece (figure 3.23). The fast capacitance compensation option of MultiClamp was activated, and a seal between the cell and the pipette tip was then created.

To create the seal, short bursts of negative pressure were applied to the cell membrane by using the mouth to briefly but strongly suck air through the mouthpiece in an attempt to tear a small hole in the membrane; this hole must occupy the membrane section surrounded by the mouth of the pipette tip, but it must not extend beyond this area if the cell is to remain alive. It normally took several attempts to create the seal, sometimes a few dozen, and often no seal could be created. When it could, however, the cell interior was directly connected to the electrode through the (electrically conducting) intracellular solution and cytoplasm, allowing intracellular recording of activity.

When a seal was correctly created (figure 3.25, bottom), the resistance at the pipette tip became about 1 G Ω . However, MultiClamp proved not to be very accurate in its resistance measurement, which fluctuated violently between about 500 M Ω and about 5 G Ω . Instead, the disappearance of the voltage step on the corresponding channel of the oscilloscope was used as a reference to know when the seal was created.

The slow capacitance compensation and whole-cell compensation options of the software were then sequentially used, after which the recording electrode was set to current-clamp mode, which maintains the current through the membrane hole at a constant value; this value, in our case, was zero, since stimulation was not required. After this, simultaneous voltage recording at both electrodes began. The voltage was recorded every 100 μ s for 100 s with a gain of 100 in the recording electrode and 10 in the reference electrode.

3.6. References

1. Alberts B, Johnson A, Lewis J, Raff M, Roberts K & Walter P (2002): *Molecular biology of the cell*, Garland Science edition, 4, 584–586
2. Seu K J, Pandey A P, Haque F, Proctor E A, Ribbe A E & Hovis J S (2007): *Effect of surface treatment on diffusion and domain formation in supported lipid bilayers*, Biophysical Journal 92, 2445–2450
3. Sych T, Mély Y & Römer W (2018): *Lipid self-assembly and lectin-induced reorganiza-*

- tion of the plasma membrane*, Philosophical Transactions of the Royal Society B 373, 20170117
4. Rietveld A & Simons K (1998): *The differential miscibility of lipids as the basis for the formation of functional membrane rafts*, Biochimica et Biophysica Acta 1376, 467–479
 5. Ulrich A S, Sami M & Watts A (1994): *Hydration of DOPC bilayers by differential scanning calorimetry*, Biochimica et Biophysica Acta 1191, 225–230
 6. Barenholz Y & Thompson T E (1999): *Sphingomyelin: biophysical aspects*, Chemistry and Physics of Lipids 102, 29–34
 7. Quinn P J & Wolf C (2010): *Egg-sphingomyelin and cholesterol form a stoichiometric molecular complex in bilayers of egg-phosphatidylcholine*, Journal of Physical Chemistry B 114, 15536–15545
 8. Veatch S L & Keller S L (2005): *Miscibility phase diagrams of giant vesicles containing sphingomyelin*, Physical Review Letters 94, 148101
 9. Slotte J P (2013): *Biological functions of sphingomyelins*, Progress in Lipid Research 52, 424–437
 10. Engberg O, Hautala V, Yasuda T, Dehio H, Murata M, Slotte J P & Nyholm T K M (2016): *The affinity of cholesterol for different phospholipids affects lateral segregation in bilayers*, Biophysical Journal 111, 546–556
 11. Melby E S, Mensch A C, Lohse S E, Dehong H, Orr G, Murphy C J, Hamers R J & Pedersen J A (2016): *Formation of supported lipid bilayers containing phase-segregated domains and their interaction with gold nanoparticles*, Environmental Science Nano 3, 45–55
 12. Angelova M I, Soléau S, Méléard P, Faucon F & Bothorel P (2007): *Preparation of giant vesicles by external AC electric fields: kinetics and applications*, in *Trends in colloid and interface science VI: progress in colloid & polymer science* 89, Steinkopff, 127–131
 13. Poole C & Losert W (2007): *Laser tweezer deformation of giant unilamellar vesicles*, in *Methods in membrane lipids* 400, Humana Press, 389–404
 14. Mennicke U & Salditt T (2002): *Preparation of solid-supported lipid bilayers by spin-coating*, Langmuir 18, 8172–8177
 15. Hager R, Lu L, Rosen G D & Williams R W (2012): *Genetic architecture supports mosaic brain evolution and independent brain-body size regulation*, Nature Communications 3, 1079
 16. Carriba P & Davies A M (2017): *CD40 is a major regulator of dendrite growth from developing excitatory and inhibitory neurons*, eLife 6, 30442

17. Brewer G J & Cotman C W (1989): *Survival and growth of hippocampal neurons in defined medium at low density: advantages of a sandwich culture technique or low oxygen*, Brain Research 494, 65–74
18. Evans P D, Kravitz E A, Talamo B R & Wallace B G (1976): *The association of octopamine with specific neurones along lobster nerve trunks*, Journal of Physiology 262, 51–70
19. González-Pérez A, Budvytyte R, Mosgaard L D, Nissen S & Heimburg T R (2014): *Penetration of action potentials during collision in the median and lateral giant axons of invertebrates*, Physical Review X 4, 031047
20. Herrick F H (1911): *Natural history of the American lobster*, Bulletin of the Bureau of Fisheries, plate 33
21. Hecht E (2002): *Optics*, fourth edition, Addison Wesley, 376–378
22. McPhee C, Zorinians G, Langbein W & Borri P (2013): *Measuring the lamellarity of giant lipid vesicles with differential interference contrast microscopy*, Biophysical Journal 105, 1414–1420
23. Williams J B, Langbein W & Borri P (2019): *qDIC: high-sensitivity quantitative analysis of differential interference contrast images*, in preparation
24. Pope I, Payne L, Zorinians G, Thomas E L H, Williams O A, Watson P D, Langbein W & Borri P (2014): *Coherent anti-Stokes Raman scattering microscopy of single nanodiamonds*, Nature Nanotechnology 9, 940–946
25. Stanford Research Systems (2003): *User's manual: model SR844 RF lock-in amplifier*, revision 2.6, Stanford Research Systems, 2-3–2-7
26. Cesari V (2009): *Ultrafast carrier dynamics in p doped InGaAs/GaAs quantum dots*, PhD thesis, Cardiff University, 34–37
27. Langbein W, Regan D, Pope I & Borri P (2018): *Heterodyne dual-polarization epideTECTED CARS microscopy for chemical and topographic imaging of interfaces*, APL Photonics 3, 092402
28. Langbein W, Rocha-Mendoza I & Borri P (2009): *Coherent anti-Stokes Raman microspectroscopy using spectral focusing: theory and experiment*, Journal of Raman Spectroscopy 40, 800–808
29. Rocha-Mendoza I, Langbein W & Borri P (2008): *Coherent anti-Stokes Raman microscopy using spectral focusing with glass dispersion*, Applied Physics Letters 93, 201103
30. Andresen E R, Berto P, Saint-Jalm S & Rigneault H (2012): *Stimulated Raman scattering microscopy by spectral focussing and fiber-generated soliton as Stokes pulse*, Proceedings of SPIE 8226, 822606
31. Sherman-Gold A (2012): *The axon guide*, third edition, Molecular Devices, 27

Chapter 4

Towards standalone optical measurement of neural activity with reflectometry

As discussed at the end of chapter 2, interferometric reflectometry is a technique suitable for the sort of measurements we wish to make, although it requires extremely sensitive measurements.

This chapter begins by deriving the theory of interferometric reflectometry as applied to the experimental configuration described in section 3.3. It then presents our experimental results along with a discussion of possible courses of action to increase the signal-to-noise ratio and reduce the phase drift in the data. Finally, the results of our electrophysiology experiments are presented.

4.1. Interferometric reflectometry theory

Reflectometry is a technique that uses the light reflected by a sample to determine properties of the sample. Interferometric reflectometry uses interference between two beams, one of which is incident on—and reflected back by—a sample and one of which is not, to obtain the complex electric field rather than merely its intensity. Since this interference allows one to retrieve an increased amount of information about the light, it also allows one to obtain more information about the sample, such as the thickness and refractive index of a thin film.

Below is the theory developed as part of this work for the interferometric reflectometry technique.

4.1.1. Reflection by a thin layer

Suppose a sample consists of a thin layer of a homogeneous material of thickness d and refractive index n_s which does not cover the entirety of a flat glass surface of refractive

index n_g and is submerged in water, which has a refractive index n_w (figure 4.1).

For normal incidence, the Fresnel reflection coefficient at the interface between a material with refractive index n_j (through which a light beam travels) and a material with refractive index n_k (which reflects the beam) is

$$r_{jk} = \frac{n_j - n_k}{n_j + n_k}$$

and the Fresnel transmission coefficient at that interface is

$$t_{jk} = \frac{2n_j}{n_j + n_k}.$$

The reflection coefficient of a region of the sample where there is no material between the glass and the water, then, is simply

$$r_{gw} = \frac{n_g - n_w}{n_g + n_w}.$$

For a region where there is a layer of material, some of the light will be reflected at the glass-material interface and some of it will be transmitted. The transmitted light might then be transmitted at the material-water interface, or it may be reflected any number ℓ of times at said interface and either be reflected $\ell - 1$ times at the material-glass interface and eventually transmitted back through this interface or be reflected ℓ times at the material-glass interface and eventually transmitted through the material-water interface. For reflection, we are interested in the first case only. Therefore, the reflection coefficient of such a region is

$$s = r_{gs} + t_{gs}t_{sg} \sum_{\ell=1}^{\infty} r_{sg}^{\ell-1} r_{sw}^{\ell} e^{2\ell i k d n_s} = r_{gs} + \frac{t_{gs}t_{sg}}{r_{sg}} \sum_{\ell=1}^{\infty} \left(r_{sg} r_{sw} e^{2i k d n_s} \right)^{\ell},$$

where the first term corresponds to reflection at the glass-material interface and the exponential in the sum is due to the fact that light reflected ℓ times at the material-water interface and $\ell - 1$ times at the material-glass interface travels 2ℓ times through the material. Here, k is the wave vector of the light in vacuum. Because $|r_{sg} r_{sw} e^{2i k d n_s}| < 1$, this is

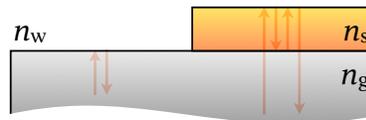


Figure 4.1: Reflection from a thin sample.

equal to

$$\begin{aligned}
 s &= r_{gs} + \frac{t_{gs}r_{sw}t_{sg}e^{2ikdn_s}}{1 - r_{sg}r_{sw}e^{2ikdn_s}} \\
 &= \frac{(n_g + n_s)(n_s - n_w)e^{2ikdn_s} + (n_g - n_s)(n_s + n_w)}{(n_g - n_s)(n_s - n_w)e^{2ikdn_s} + (n_g + n_s)(n_s + n_w)}. \quad (4.1)
 \end{aligned}$$

If $d = 0$ or $n_s = n_w$, this reduces to r_{gw} . If $n_s = n_g$, it instead reduces to $e^{2ikdn_g}r_{gw}$ due to the fact that the light must still travel an additional distance $2d$ through material with refractive index n_g .

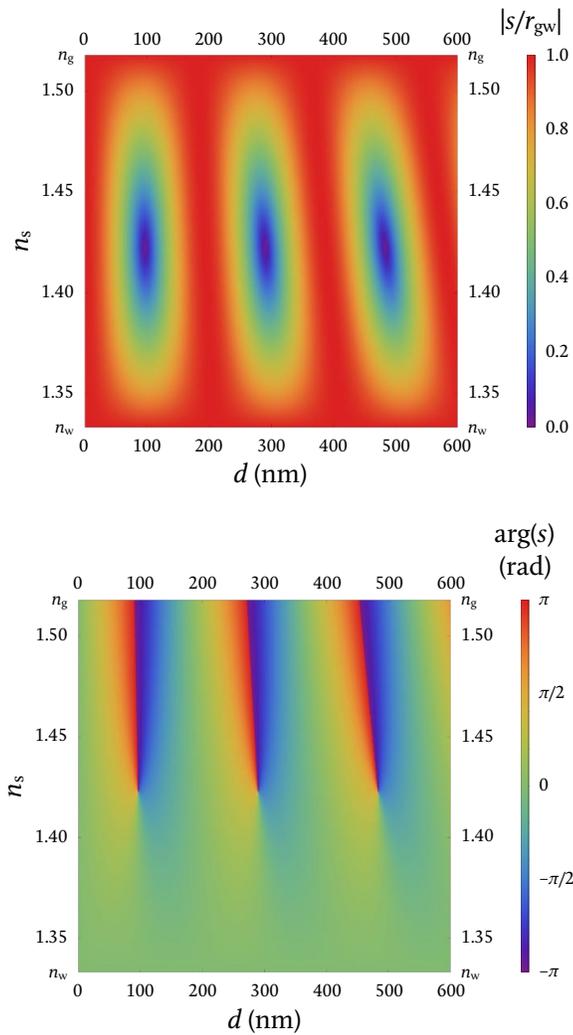


Figure 4.2: Density graphs of $|s/r_{gw}|$ (top) and $\arg(s)$ (bottom) as functions of d and n_s for normal incidence with $n_w = 1.333$, $n_g = 1.518$ and $\lambda = 550$ nm.

Here, we have assumed that the sample does not change the polarisation of the beam (i.e. it is not birefringent) and thus s is a scalar. If it has in-plane birefringence, s is instead a 2×2 matrix and equation 4.1 is no longer appropriate; if it presents out-of-plane birefringence, the calculation of s can become truly complicated. It should be noted that the type of sample we are interested in here, a lipid bilayer, consists of a 2-dimensional array of lipid molecules with their tails approximately perpendicular to said array and can thus be considered approximately isotropic for light travelling parallel (or approximately parallel) to the lipid tails; indeed, the birefringence of a lipid bilayer is negligible for our purposes. We will therefore ignore the effects of birefringence in what follows.

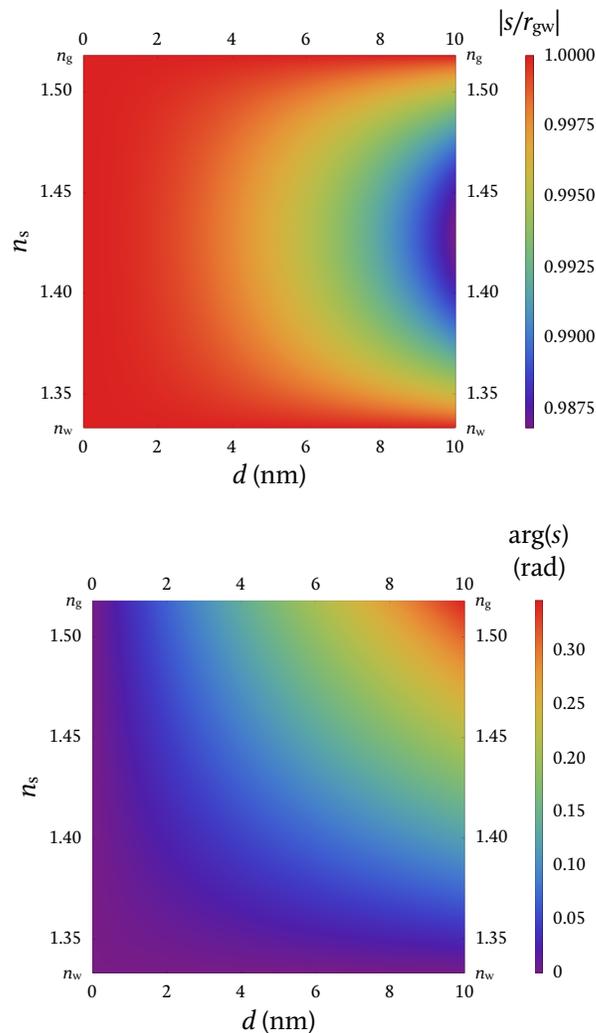


Figure 4.3: Density graphs of $|s/r_{gw}|$ (top) and $\arg(s)$ (bottom) as functions of d and n_s for normal incidence with a reduced d range and with $n_w = 1.333$, $n_g = 1.518$ and $\lambda = 550$ nm.

Figure 4.2 shows graphs of $|s|$ and $\arg(s)$ for d between 0 nm and 600 nm, $n_w = 1.333 \leq n_s \leq 1.518 = n_g$, and $\lambda = 2\pi/k = 550$ nm. These values for the refractive indices were chosen because they correspond to water¹ and the glass which microscope slides and coverslips are typically made of;² the refractive index of a lipid bilayer typically falls between these values.^{3,4} The thickness range was chosen to show three periods of s . It is interesting to note the existence of a value of n_s below which $\arg(s)$ may only take values between $-\pi/2$ and $\pi/2$, meaning the reflection coefficient has a positive real part (see appendix A). Note also that, for $n_w \leq n_s \leq n_g$, we have $|s| \leq r_{gw}$ regardless of the value of d , meaning that the presence of the layer either reduces the amount of reflected light (by spatially distributing the refractive index step from n_g to n_w and giving rise to an interference which is not fully constructive) or does not affect it.

Figure 4.3 shows the same graphs as figure 4.2, but for d between 0 nm and 10 nm only. This range of d covers the thickness one would expect from a lipid bilayer, for example, which is about 4 nm thick. Note that $|s|$ changes very little in this region — the difference between no layer and a 4-nm layer is only about 0.5% for $n_s \approx 1.425$ and even less for other values of n_s (the variation in $\arg(s)$ is also reduced, but much less so — it is a few percent even for intermediate values of n_s). This is due to the fact that the sample is very thin; a small value of d will result in a small value of $2kdn_s$, which in turn means that the first few reflections of the beam within the sample interfere mostly constructively; by the time the number of reflections is large enough for the interference to be destructive, the amplitude of the beam is so small (due to the fact that $|r_{sw}|, |r_{sg}| < 1$) that it contributes very little to the reflected field E_s . In fact, if we only take the first reflection into account, the reflection coefficient becomes

$$s \approx r_{gs} + t_{gs}r_{sw}t_{sg}e^{2ikdn_s}.$$

The difference between this and the complete reflection coefficient given by equation 4.1 is less than about 0.06% throughout the range considered, as might be expected by noting that the denominator of the second term in the first line of equation 4.1 is approximately equal to 1 because $r_{sg}r_{sw} < r_{gw}^2 \approx 0.0042$. Thus, taking only one reflection into account is an acceptable approximation.

4.1.2. Obtaining thickness and refractive index from the reflection coefficient

To obtain the thickness of the sample, we rewrite equation 4.1 as

$$e^{2ikdn_s} = \frac{n_s + n_w}{n_s - n_w} \frac{n_g - n_s - (n_g + n_s)s}{n_g + n_s - (n_g - n_s)s},$$

whereby

$$d = \frac{i}{2kn_s} \log \left(-\frac{n_s - n_w}{n_s + n_w} \frac{n_g + n_s - (n_g - n_s)s}{n_g - n_s - (n_g + n_s)s} \right), \quad (4.2)$$

where \log denotes the complex logarithm. It is now evident that this expression has an infinite number of values and that choosing one equates to choosing a logarithm branch.

This is why the pattern seen in figure 4.2 is periodic in d ; its period is $2\pi/2kn_s$, as is evident from equation 4.2.

To obtain the refractive index of the sample, we note that the imaginary part of d is

$$\text{Im}(d) = \frac{1}{2kn_s} \ln \left(\left| \frac{n_s - n_w}{n_s + n_w} \frac{n_g + n_s - (n_g - n_s)s}{n_g - n_s - (n_g + n_s)s} \right| \right).$$

But d is the thickness of the sample and must thus be a real number. Therefore,

$$\left(\frac{n_s - n_w}{n_s + n_w} \right)^2 \frac{n_g + n_s - (n_g - n_s)s}{n_g - n_s - (n_g + n_s)s} \frac{n_g + n_s - (n_g - n_s)s^*}{n_g - n_s - (n_g + n_s)s^*} = 1.$$

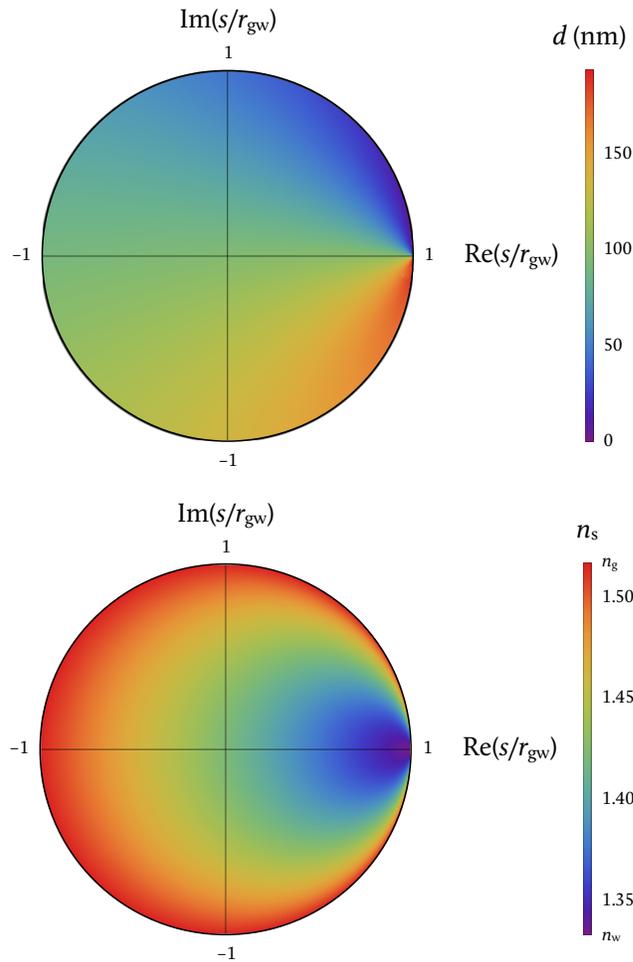


Figure 4.4: Density graphs of the first period of d (top) and n_s (bottom) as functions of the real and imaginary parts of s for normal incidence with $n_w = 1.333$, $n_g = 1.518$ and $\lambda = 550$ nm.

From this expression we finally obtain

$$n_s = \sqrt{n_w n_g \frac{n_g - n_w - 2n_g \operatorname{Re}(s) + (n_g + n_w)|s|^2}{n_g - n_w - 2n_w \operatorname{Re}(s) - (n_g + n_w)|s|^2}}. \quad (4.3)$$

As noted in section 4.1.1, $|s/r_{\text{gw}}| \leq 1$ as long as $n_w \leq n_s \leq n_g$. This is expected because in this refractive index range the difference between n_g and n_w is greater than the difference between n_g and n_s , which results in the glass-water interface being more reflective than the glass-material interface. Taking this into account, we may graph d and n_s in the unit circle of the $\operatorname{Re}(s/r_{\text{gw}}) \times \operatorname{Im}(s/r_{\text{gw}})$ plane (figure 4.4).

4.1.3. Non-normal incidence

For non-normal incidence, the Fresnel reflection and transmission coefficients are

$$r_{jk}^{\parallel} = \frac{n_j \cos(\theta_j) - n_k \cos(\theta_k)}{n_j \cos(\theta_j) + n_k \cos(\theta_k)},$$

$$t_{jk}^{\parallel} = \frac{2n_j \cos(\theta_j)}{n_j \cos(\theta_j) + n_k \cos(\theta_k)}$$

for light polarised parallel to the plane of incidence and

$$r_{jk}^{\perp} = \frac{n_k \cos(\theta_j) - n_j \cos(\theta_k)}{n_k \cos(\theta_j) + n_j \cos(\theta_k)},$$

$$t_{jk}^{\perp} = \frac{2n_j \cos(\theta_j)}{n_k \cos(\theta_j) + n_j \cos(\theta_k)}$$

for light polarised perpendicular to the plane of incidence, where θ_j is the angle of incidence and

$$\theta_k = \arcsin\left(\frac{n_k}{n_j} \sin(\theta_j)\right)$$

is the angle of transmission, given by Snell's law. Obviously, then, s is a function of the angle of incidence as well as of d and n_s . We will henceforth assume that θ_j is small enough and n_k is close enough to n_j to avoid total internal reflection when $n_j > n_k$.

If the incident light is circularly polarised, as in the case of our experimental setup (see section 3.3), the parallel and perpendicular components have equal amplitude, so we have

$$s = \frac{s_{\parallel} + s_{\perp}}{2}.$$

For a distribution $\mathcal{P}(\theta, \varphi)$ of angles of incidence, the reflectometry signal S , which is given by the interference between the probe and reference beams and thus contains the

angular distribution of both beams, must be averaged over all possible angles. This average is given by

$$S = \frac{1}{2\pi} \int_0^{2\pi} \int_0^{\theta_{\max}} \mathcal{P}(\theta, \varphi) s(d, n_s, \theta) \cos(\theta) \sin(\theta) d\theta d\varphi,$$

where θ_{\max} is the maximum angle of incidence of the light incident on the sample; if θ_{\max} is determined by the numerical aperture NA of the microscope objective, for instance, then

$$\theta_{\max} = \arcsin\left(\frac{\text{NA}}{n_g}\right).$$

The cosine in the integral comes from assuming an aplanatic objective is used; the projection of an area element dA of the incident interfered beams onto the aplanatic lens reference sphere is $dA/\cos(\theta)$ (figure 4.5).⁵

It should be noted that the critical angle, the angle at which total internal reflection occurs, is

$$\theta_c = \arcsin\left(\frac{n_w}{n_g}\right) = 61.42^\circ$$

for the glass-water interface and even higher for the glass-layer and layer-water interfaces if the layer has a refractive index between n_w and n_g . For an objective with a numerical aperture of 1.27, $\theta_{\max} = 56.79^\circ$, so total internal reflection is not a problem at the glass-layer and glass-water interfaces. Light travelling through the glass at an angle $\theta \leq \theta_{\max}$ will be transmitted into the layer at an angle

$$\theta_s = \arcsin\left(\frac{n_g}{n_s} \sin(\theta)\right),$$

so at the layer-interface we have

$$\theta_s - \theta_c = \arcsin\left(\frac{n_g}{n_s} \sin(\theta)\right) - \arcsin\left(\frac{n_w}{n_s}\right),$$

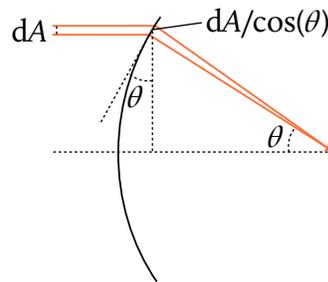


Figure 4.5: The projection of an area element dA of the incident beam onto the reference sphere of an aplanatic objective is $dA/\cos(\theta)$.

which is a monotonically increasing function of both n_s and θ but is negative even for the highest value of n_s we are considering, n_g ; therefore, there will be no total internal reflection at the layer-water interface either.

Since both beams are gaussian and have the same angular distribution, the angular dependence of the detected signal is a function of only θ and is given by⁵

$$\mathcal{P}(\theta) = e^{-\zeta^2 \frac{\sin^2(\theta)}{\sin^2(\theta_{\max})}},$$

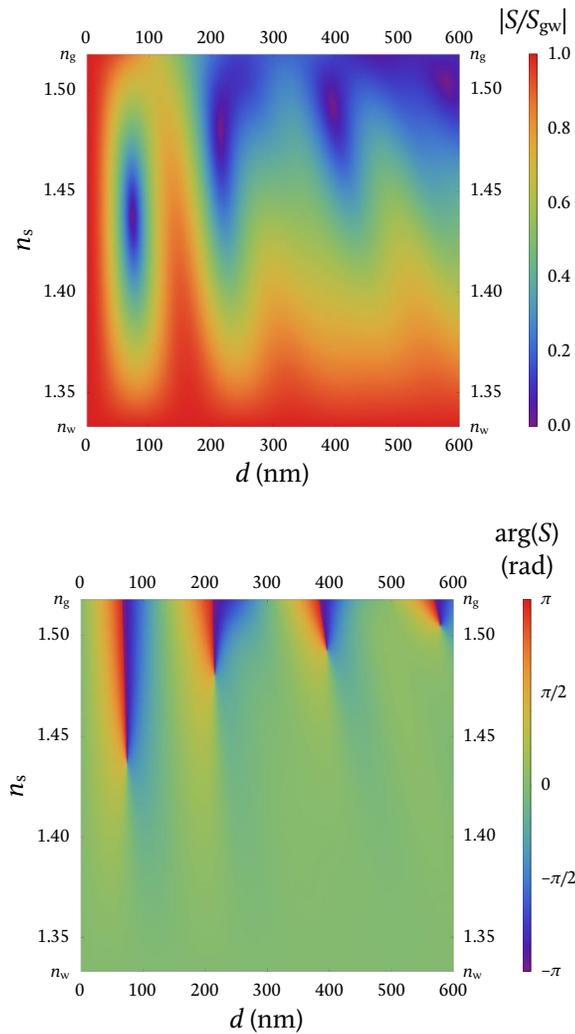


Figure 4.6: Density graphs of $|S/S_{\text{gw}}|$ (top) and $\arg(S)$ (bottom) as functions of d and n_s for the case in which both beams are gaussian, the probe beam emerges from (and is then reflected back through) an objective with numerical aperture 1.27 and fill factor 1, and $n_w = 1.333$, $n_g = 1.518$ and $\lambda = 550$ nm.

where ζ is the objective fill factor. This turns the signal into

$$S = \int_0^{\theta_{\max}} e^{-\zeta^2 \frac{\sin^2(\theta)}{\sin^2(\theta_{\max})}} \frac{s_{\parallel} + s_{\perp}}{2} \cos(\theta) \sin(\theta) d\theta. \quad (4.4)$$

In this case, the expression for $S(d, n_s)$ can no longer be solved analytically.

Figure 4.6 shows the amplitude and phase of the reflectometry signal, S , normalised with respect to the no-layer signal, S_{gw} (given by replacing s with r_{gw} in equation 4.4), for gaussian beams, an objective with a numerical aperture of 1.27 and a fill factor of 1, and

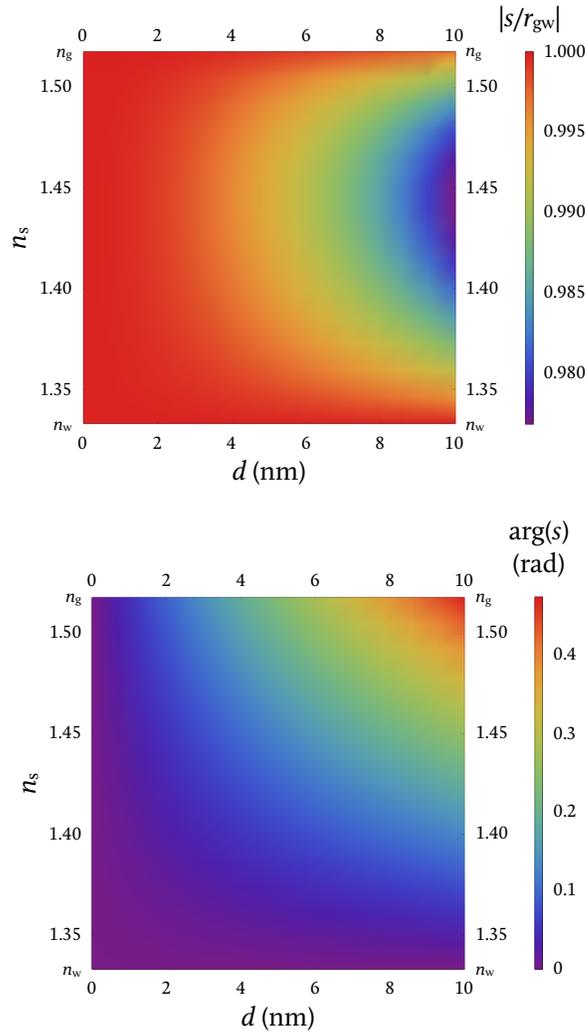


Figure 4.7: Density graphs of $|S/S_{\text{gw}}|$ (top) and $\arg(S)$ (bottom) as functions of d and n_s for the case in which both beams are gaussian, the probe beam emerges from (and is then reflected back through) an objective with numerical aperture 1.27 and fill factor 1, and $n_w = 1.333$, $n_g = 1.518$ and $\lambda = 550$ nm with a reduced d range.

all other parameters as before. With these parameters, $S_{\text{gw}} \approx -0.0063$. Figure 4.7 is as figure 4.6 for d from 0 to 10 nm.

For a small value of θ_{max} (i.e. for a small NA), the relative reflection coefficient is only slightly deformed with respect to that observed for normal incidence; for larger values of θ_{max} , the pattern is no longer periodic (although it retains a partially repetitive behaviour) and its deformation becomes more evident (compare figures 4.2 and 4.6). For single bi-layers, the change is small compared to the normal-incidence case (compare figures 4.3

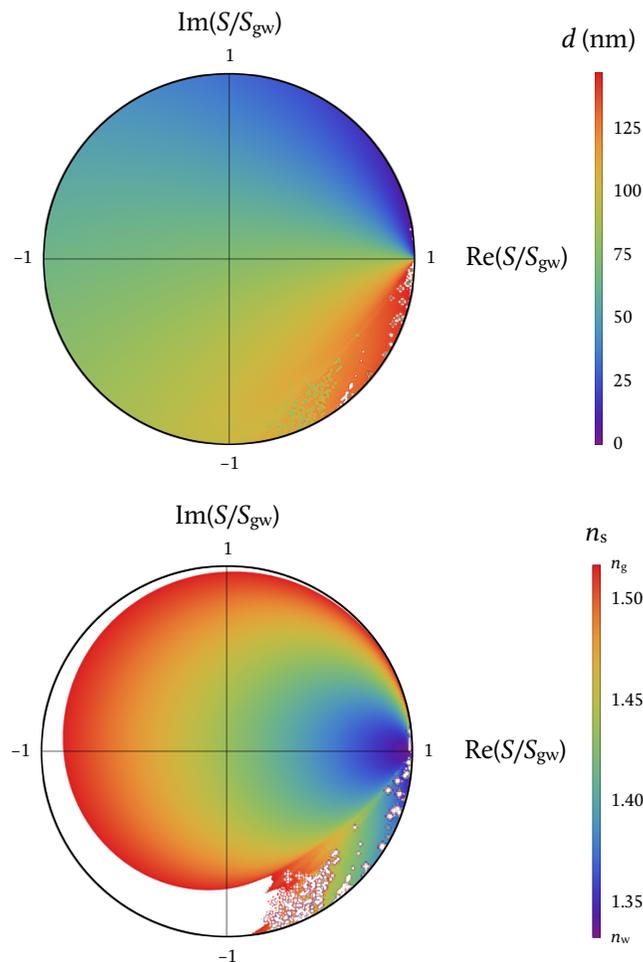


Figure 4.8: Density graphs of the first repetition of d (top) and n_s (bottom) as functions of the real and imaginary parts of S/S_{gw} for the case in which both beams are gaussian, the probe beam emerges from (and is then reflected back through) an objective with numerical aperture 1.27 and fill factor 1, and $n_w = 1.333$, $n_g = 1.518$ and $\lambda = 550$ nm. The irregularly coloured section in the fourth quadrant of the n_s graph is an artefact caused by part of the second repetition having been included in the calculation of n_s , which occurred because a rectangular section of (d, n_s) space was taken and the repetitions are not rectangular.

and 4.7).

It is immediately evident that $|S/S_{\text{gw}}| \leq 1$ as long as $n_w \leq n_s \leq n_g$. We may thus graph d and n_s in the unit circle of the $\text{Re}(S/S_{\text{gw}}) \times \text{Im}(S/S_{\text{gw}})$ plane. To do so, we must take into account only one repetition (as mentioned earlier, the behaviour of S for non-normal incidence is still partially repetitive, although the values of n_s for which S reaches the same value in different repetitions are different from each other).

Figure 4.8 shows d and n_s as functions of $\text{Re}(S/S_{\text{gw}})$ and $\text{Im}(S/S_{\text{gw}})$ for a gaussian beam emerging from an objective with a numerical aperture of 1.27 and a fill factor of 1. Comparing this to figure 4.4, it is immediately evident that, while the entire $(|S/S_{\text{gw}}|, \arg(S)) \in [0, 1] \times [-\pi, \pi]$ space yields values of n_s between n_w and n_g in the case of normal incidence, this is not so in the case of large numerical aperture, as expected from careful examination of the bottom, left and top borders of figures 4.2 and 4.6: whereas in the case of normal incidence $\arg(S)$ takes all possible values between $-\pi$ and π along the top border (where $n_s = n_g$ and $|S/S_{\text{gw}}| = 1$), a non-zero numerical aperture causes $|S/S_{\text{gw}}|$ to no longer equal 1 along the top border, which is the only one of these borders along which $\arg(S)$ takes non-zero values; for a non-zero numerical aperture, $\arg(S) = 0$ all along the left ($d = 0$) and bottom ($n_s = n_w$) borders, which are the only places where $|S/S_{\text{gw}}| = 1$; this translates into figure 4.8 as the white region where $|S/S_{\text{gw}}| \lesssim 1$ and $\arg(S) \neq 0$ simultaneously.

To create figures 4.6 and 4.7, we first numerically calculated S using equation 4.4 for a fine-enough partition of the $d \times n_s$ space and performed a third-order interpolation of the resulting table to obtain a continuous function $S(d, n_s)$. To create figure 4.8, we did this for one repetition in d and then numerically inverted the interpolated function by finding the zeros of $S - S_j$ for different S_j obtained by finely partitioning the $\text{Re}(S) \times \text{Im}(S)$ space (constrained by $|S| \leq S_{\text{gw}}$), after which we performed third-order interpolations to obtain $d(S)$ and $n_s(S)$.

4.2. Dealing with noise

Because the reflection coefficient is not very sensitive to changes in thickness and refractive index in the range we are interested in (recall the discussion at the end of section 4.1.1), a very high signal-to-noise ratio, preferably at least 2,000 (so that 0.5% changes in signal are still an order of magnitude stronger than the noise), is required when measuring the light reflected by the sample. As will be seen shortly, the signal-to-noise ratio of interferometric reflectometry images of lipid bilayers is rather small, often lower than 1. Thus, strategies to reduce the noise to an acceptable level must be implemented before the images are analysed.

Sources of noise include fluctuations in the intensity and polarisation of the incident light, variations in the relative phases of the probe and reference beams due to them having separate beam paths (i.e. any changes in the optical path length of only one of the beams, such as those caused by vibrations of the sample stage), dark noise of the detectors and shot noise. The first of these noise sources will affect both the amplitude and the phase

of the reflected field, while the second will affect only the phase and the third and fourth will affect only the amplitude.

4.2.1. Phase filtering

One way to deal with noise is to use a filter to remove certain frequencies in the data. Since the phase noise dominated the images and the amplitude noise was comparatively low, this was primarily done with the phase of the data. This must be done on the raw time traces (i.e. before regularisation).

This method consists of taking the Fourier transform of the data and removing (by multiplying by $1 - f(\nu)$, where f is a gaussian function of the frequency ν ; the amplitude of f is 1) frequencies unrelated to the structures in the sample (figure 4.9) before taking the inverse Fourier transform.

At the very least, frequencies under 5 Hz, as well as 15 harmonics of the fast-axis scan frequency, were filtered out. Because adjacent rows of an image were obtained by scanning the sample stage in alternating directions along the fast axis (the x axis), the fast-axis scan period is $2N_x t_{\text{exp}}$, where N_x is the number of pixels in the fast scan axis and t_{exp} is the pixel dwell time; the fast-axis scan frequency is thus $1/2N_x t_{\text{exp}}$. There was also a collection of frequency peaks around 10 Hz present in all images; these peaks were also filtered out in every case. Several versions of the images, each with a different degree of additional phase filtering, were analysed; these ranged from no additional filtering (only the aforementioned frequencies were removed) to severe (any peaks visible in the Fourier transform were removed), with several intermediate versions.

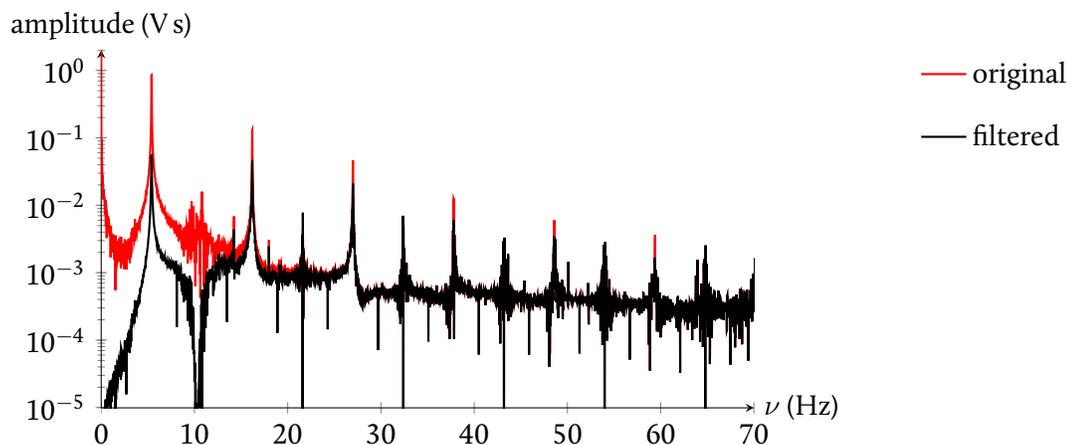


Figure 4.9: Frequencies of the phase of the co-polarised channel of a typical image of a DC₁₅PC bilayer (seen in figure 4.10) before and after filtering. Only a small frequency range is shown. The sharp, equally spaced peaks are the harmonics of about 5.4 Hz, the fast-axis scan frequency in this particular image.

Phase filtering removed the enormous phase drift and some of the noise, though it was found to be impossible to remove all of the noise without also destroying some structural information about the sample, since often it was difficult to discriminate between noise and signal in the frequency domain. In general, any filtering other than the weakest version resulted in blurred structure edges and loss of information about the sample (see figure 4.10), so in the end only the aforementioned frequencies were removed. This meant that much noise remained, making it impossible to accurately determine the thickness and refractive index at each pixel of each image.

Despite the small amount of filtering it was possible to do without loss of information and the high amount of noise that remained as a result, it was possible to observe not only single bilayers, as evidenced in figure 4.10, but also the coexistence between different phases of the lipid in ternary mixtures (DOPC:sm:ch = 11:5:4, figure 4.11). These appeared as faint changes in reflected signal (the thicker LO domains reflected less light than the thinner LD domains). Figure 4.11 shows the reflectometry amplitude and phase signal from three images from the same ternary sample alongside quantitative DIC images of the same sample. The boundaries between the LD and LO domains are visible as slight changes in amplitude, although the phase remains almost the same; they are marked with arrows in the figure so it is easier to see them. They correspond to precisely the signal changes observed in the quantitative DIC images. The change in thickness between

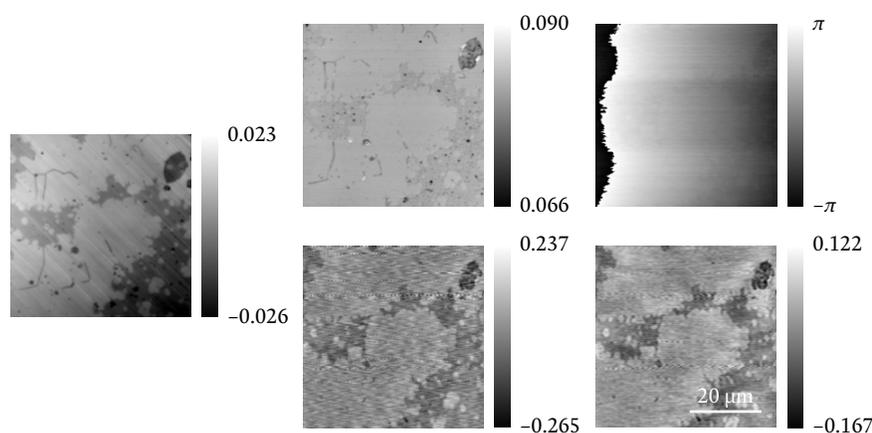


Figure 4.10: Top centre: Amplitude of the co-polarised channel of the image from figure 4.9. Bottom centre and right: Phase of the co-polarised channel of the same image: unfiltered (top right); with frequencies below 5 Hz, 15 harmonics of the fast-axis scan frequency and the collection of peaks around 10 Hz filtered out (bottom centre); and with all frequencies below 300 Hz filtered out (bottom right). The image on the bottom right is overfiltered, resulting in loss of information. Left: Quantitative DIC image of the same field of view. In all cases, the dark patches are lipid bilayers (typically a few μm across) and the lighter regions contain no lipid. The dark structure at the top-right corner of each image is an irregular multilamellar lipid patch. Greyscale ranges are in radians for the phase and quantitative DIC images and in volts for the amplitude image.

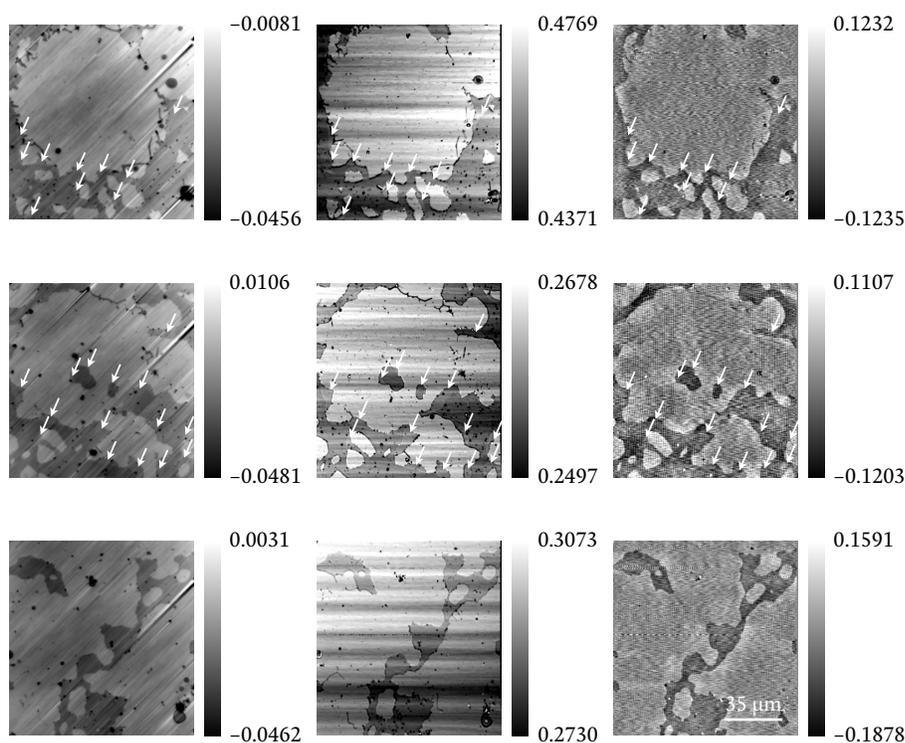


Figure 4.11: Amplitude (centre) and phase (right) of three interferometric reflectometry images of the same ternary sample alongside quantitative DIC (left) images of the same fields of view. The composition (DOPC: sphingomyelin: cholesterol) of the sample is 11:5:4. Boundaries between the LD and LO domains can be seen as faint changes in quantitative DIC signal and reflectometry amplitude but not reflectometry phase; they are marked with white arrows. The bottom image shows no domain segregation; the lipid in it is homogeneous. In the other two images, it is evident that LD domains prefer being at the bilayer edge over being completely surrounded by LO lipid. All the images have the same size and scale. Greyscale ranges are in radians for the phase and quantitative DIC images and in volts for the amplitude images.

an LD domain and an LO domain is about 0.9 nm, so interferometric reflectometry can resolve sub-nanometre structures in the axial direction.

4.2.2. Referencing

All of the sources of noise listed at the beginning of section 4.2, with the exception of dark noise and shot noise, which we will ignore because they are much smaller than the other sources of noise, can be removed with a technique called referencing, which consists of comparing two sets of data with either identical noise or identical signal (but not both).

One way to achieve this is to place a Wollaston prism in the probe beam's path before it reaches the sample. Recall that the probe beam is otherwise circularly polarised at the sample; a Wollaston prism converts this into two linearly polarised beams with orthogonal

polarisations (\vec{H} and \vec{V}) and slightly different beam paths. The two beams are incident on slightly different points on the sample, which we may call \vec{r}_H and \vec{r}_V , and then reflected back into the prism, where they recombine into a single beam before being detected. At the point of recombination, both \vec{H} and \vec{V} carry information of the sample at \vec{r}_H and \vec{r}_V (figure 4.12, left), so this information must be separated (figure 4.12, right). Once it has been separated, however, it is suitable for referencing. If the (complex) noiseless signal from point \vec{r}_j would be S_j , then the detected quantity is $\sigma_j S_j$. Because \vec{H} and \vec{V} come from the same source, are incident on the sample at the same time and are detected at the same time by the same detectors, the noise on one is identical to that on the other: $\sigma_H = \sigma_V$. By dividing one detected quantity by the other, then, we can get rid of the noise. If we know, for example, what S_V should be (for instance, because we know exactly the sample thickness and refractive index at \vec{r}_V), we may retrieve S_H and thus know what the sample thickness and refractive index at \vec{r}_H are; this is easily achieved by ensuring that there is no material (and thus $s = r_{gw}$) at \vec{r}_V . It is important to note that, since we no longer have circular polarisation, equation 4.4 needs to be modified to turn the average of s for the two polarisations into the corresponding linear polarisation.

This solution introduces the problem of separating the two polarisation components. The problem would be trivial (due to the simple relationship between circular and linear polarisations) were it not for the fact that the beams are incident on different points of

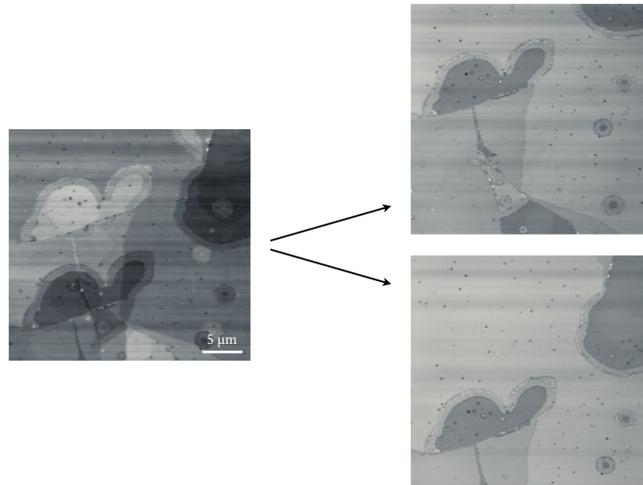


Figure 4.12: Amplitude of an interferometric reflectometry image taken with a Wollaston prism for referencing (left). Note that all structures appear twice and that the shear (the separation distance and direction between a structure and its copy) is constant throughout the image. In this case, the shear is approximately vertical. The use of a Wollaston prism introduces the problem of needing to separate the information from the two polarisations. Here, approximately separated amplitude images are shown (right); unfortunately, the corresponding phase images (not shown) were not separated at all (the separated amplitude images are shown for illustrative purposes only).

the sample and the fact that one beam travels a somewhat longer optical distance within the Wollaston prism than the other due to the refractive index difference, which makes it acquire a phase, and might have a slightly different amplitude due to the Fresnel coefficients at the prism's surfaces being different. We thus have the problem of separating two components from a linear combination of them where we do not know the coefficients.

We know that the measured data sets D_1 and D_2 are linear combinations of the true orthogonally polarised images S_H and S_V :

$$D_1 = a_1 S_H + a_2 S_V,$$

$$D_2 = a_3 S_H + a_4 S_V,$$

where we have ignored the noise terms for clarity. We also know, as explained previously, that the effect of the polarisation information combination on the amplitude is very small, which means that $|a_1| \approx |a_2| \approx |a_3| \approx |a_4|$. Furthermore, since we wish to divide one image by the other once they are separated and we know the glass-water reflection coefficient, it does not matter if after separation we are left with multiples of the S_j instead of with the S_j themselves. The number of variables is thus reduced from 4 to 3. However,

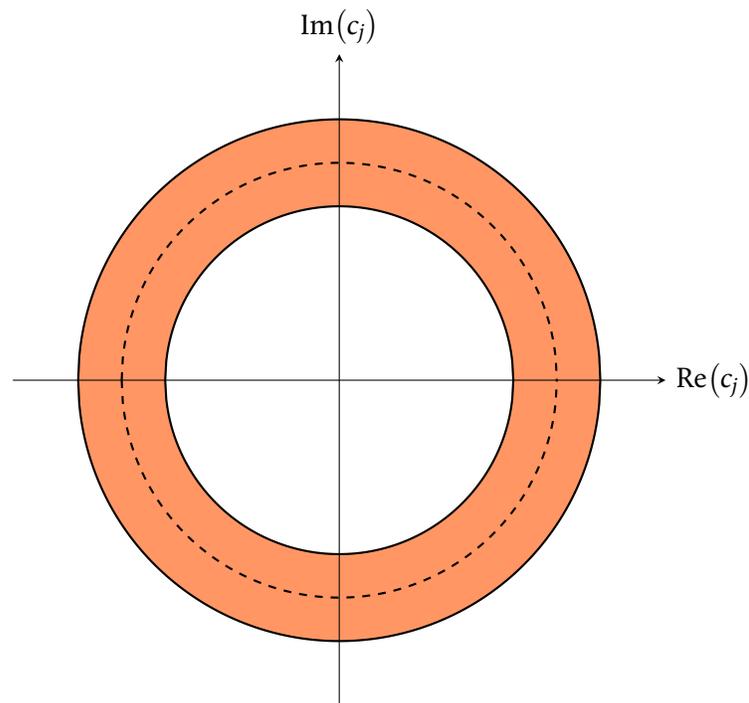


Figure 4.13: Narrow ring (shaded region) composed of the values each of the c_j can take, which lie close to the unit circle on the complex plane. The 6-dimensional space inhabited by the solution we seek consists of the outer product of three such rings, one for each of the c_j .

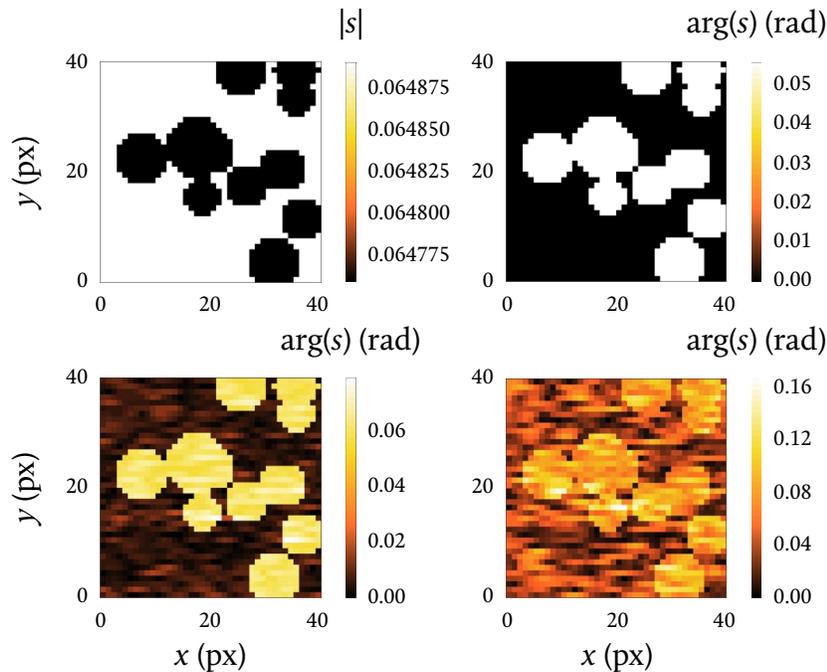


Figure 4.14: An example of simulated data: amplitude (top left) and phase (rest of images) of the reflection coefficient s a 4-nm-thick layer of material with refractive index 1.44 would have for normal incidence. The parameters used in the calculation (n_w , n_g and λ) are the same as in all previous figures. The phase shown here has different amounts of gaussian noise of amplitude a and standard deviation $\sigma = \text{Im}(s)$: $a = 0$ (top right), $a = 2,000 \text{Im}(s)$ (bottom left) and $a = 10,000 \text{Im}(s)$ (bottom right).

since they are complex, the problem is 6-dimensional. The space occupied by the parameters we seek is the outer product of three narrow rings on the complex plane with inner and outer radii close to 1 and centred at $(0, 0)$ (figure 4.13).

We may calculate the following linear combinations of the measured quantities:

$$A_1 = D_1 + c_1 D_2, \quad (4.5a)$$

$$A_2 = c_3 (c_2 D_1 + D_2), \quad (4.5b)$$

where c_1 , c_2 and c_3 are unknown complex numbers. If we obtain the correct coefficients c_j , then the A_j will be multiples of the S_j .

We have attempted to achieve separation with a correlation method (appendix B), a gradient descent algorithm (appendix C) and a genetic algorithm (appendix D), in all three cases treating this as a minimisation problem; these, as well as the obtained results, will be discussed below. In all cases, the computer code of the algorithm was written by the author of this work in its entirety.

The methods described here were tested on both simulated and experimental data. The simulated data consisted of an array of $N_x \times N_y = 40 \times 40$ pixels divided into two types

of homogeneous sections. For one section type, the reflection coefficient s was calculated for a 4-nm-thick bilayer with refractive index 1.44 using equation 4.1; for the other type, r_{gw} was used. Noisy data was also simulated. In order to simulate the noise observed in experimental images, a $1 \times N_x N_y$ array of gaussian white noise with amplitude a , mean 0 and standard deviation $\sigma = \text{Im}(s)$ was created, low-pass-filtered, wrapped into an $N_x \times N_y$ array and added to the phase of the simulated data. Figure 4.14 shows an example of simulated data with different noise amplitudes: $a = 0$, $a = 2,000\sigma$ and $a = 10,000\sigma$. The latter case approximates the amount of noise observed in experimental data.

Correlation

The correlation method consists of minimising the normalised cross-correlation of A_1 and A_2 at zero displacement, $(A_1 \star A_2)(0)$, where

$$(A_1 \star A_2)(\vec{r}) = \frac{(A_1(0) - \langle A_1 \rangle)^* (A_2(\vec{r}) - \langle A_2 \rangle)}{\mathcal{S} \text{var}(A_1) \text{var}(A_2)}; \quad (4.6)$$

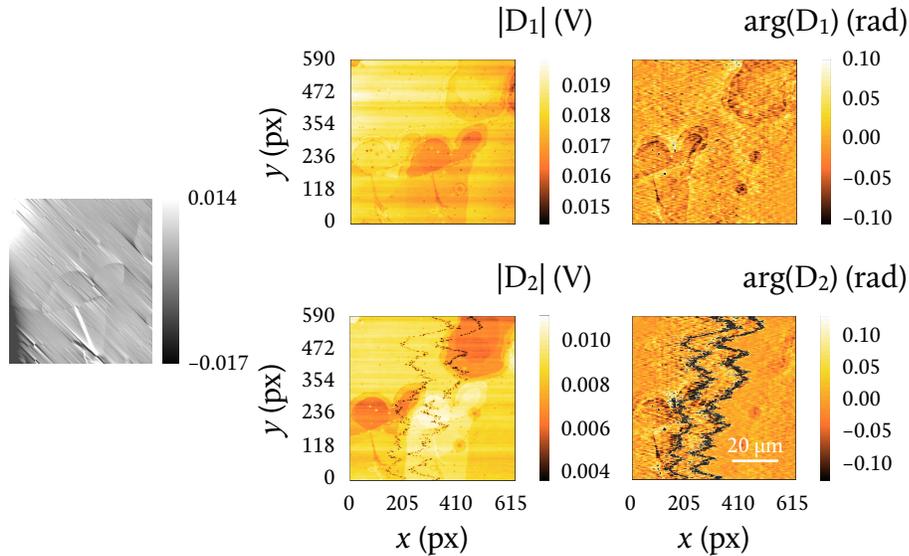


Figure 4.15: Amplitude (centre) and phase (right) of the co- (top) and cross-polarised (bottom) components of the signal of a ternary (DOPC:sm:ch = 11:5:4) bilayer detected using a Wollaston prism for referencing. The bilayer composition (DOPC:sm:ch) is 11:5:4. The shear direction is approximately horizontal; the combination of polarisations is visible as a duplication of every feature in the direction of the shear (some features are only visible once because their copies lie outside the field of view). The phase has been minimally filtered as described in section 4.2.1 to remove the phase drift of several radians, which completely obscured the features, and the fast-axis scan frequency (3.858 Hz), but most of the noise remains because excessive filtering results in loss of information. A quantitative DIC image of the same field of view (without the displacement induced by the Wollaston prism) is shown on the left.

here, \vec{r} is the displacement, $\langle A_j \rangle$ is the average of A_j over the overlap region of the image, S is the overlap area between the images (which for $\vec{r} = 0$ is the entire image) and var is the variance of a data set; the product is the sum of the pixel-by-pixel product of the images over the overlap region. The method relies on the fact that, since the A_j will contain two copies of every structure at the same positions unless the c_j are the correct ones for image separation (in which case each A_j will contain a single copy of every structure at a different position), the cross-correlation at zero displacement has a minimum at the correct values of the c_j in (c_1, c_2, c_3) space (recall figure 4.12).

The data was minimally phase-filtered as described in section 4.2.1 prior to the correlation analysis so the phase drift would not completely overwhelm the spatial features. Figure 4.15 shows an example of a phase-filtered image of a ternary lipid bilayer. Features are duplicated along the shear direction (approximately horizontal) due to the presence of the Wollaston prism.

Since extended structures can contribute to the cross-correlation at zero displacement even in the case in which the images are completely separated, we actually calculated the A_j using sufficiently high-pass-filtered versions of the D_j for only the edges of structures to remain. Figure 4.16 shows the result of such a high-pass filtering of the D_j , obtained by

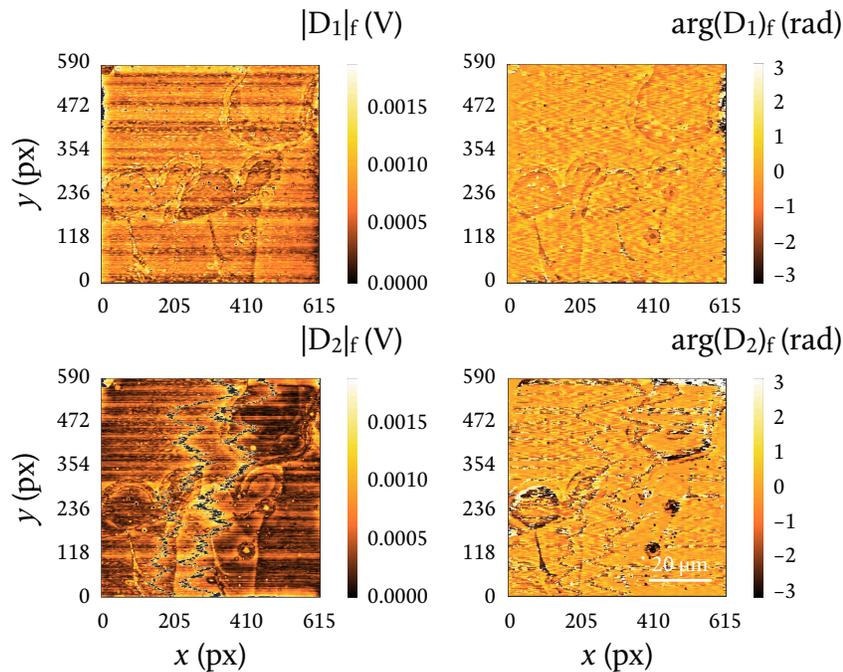


Figure 4.16: Amplitude (left) and phase (right) of the co- (top) and cross-polarised (bottom) components of the signal after applying a gaussian high-pass filter to the data with a “cutoff” frequency (the frequency at which the amplitude of the gaussian is $1/\sqrt{e}$) of 0.02 px^{-1} . The noise is strong, particularly in the cross-polarised data, but the edges of the lipid features can still be seen.

taking the Fourier transform of the D_j , multiplying it by a function which is approximately 1 for high frequencies and close to 0 for low frequencies, and taking the inverse Fourier transform of the result; in this case, the filtered data is

$$(D_j)_f = \mathcal{F}^{-1} \left[\left(1 - e^{-\frac{\xi^2 + \eta^2}{2\sigma^2}} \right) \mathcal{F}[D_j] \right],$$

where \mathcal{F} is the Fourier transform, ξ and η are the x - and y -direction spatial frequencies (respectively), and $\sigma = 0.4 \text{ px}^{-1}$.

Figure 4.17 shows several projections into 2-dimensional space of the correlation of $(A_1)_f$ and $(A_2)_f$, which are as A_1 and A_2 from equations 4.5 but using the $(D_j)_f$ instead of the D_j , at zero displacement: $\vec{r} = (0, 0)$. Although the range of the $|c_j|$ in the figure is $[0, 2]$,

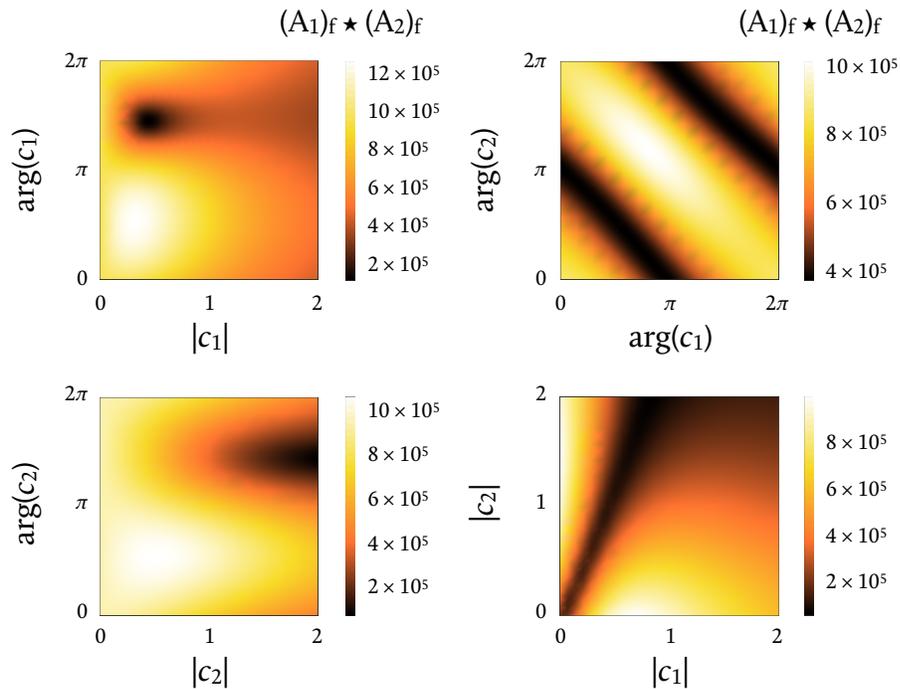


Figure 4.17: Several 2-dimensional projections of the correlation at zero displacement of the $(A_1)_f$ and $(A_2)_f$ (setting $c_3 = 1$) from the experimental data shown in previous figures in this appendix. In $(|c_1|, \arg(c_1), |c_2|, \arg(c_2), |c_3|, \arg(c_3))$ space, the projections are as follows:

$$\begin{aligned} & \left(|c_1|, \arg(c_1), 1, \frac{3\pi}{2}, 1, 0 \right) \quad (1, \arg(c_1), 1, \arg(c_2), 1, 0) \\ & \left(1, \frac{3\pi}{2}, |c_2|, \arg(c_2), 1, 0 \right) \quad \left(|c_1|, \frac{3\pi}{2}, |c_2|, \frac{3\pi}{2}, 1, 0 \right) \end{aligned}$$

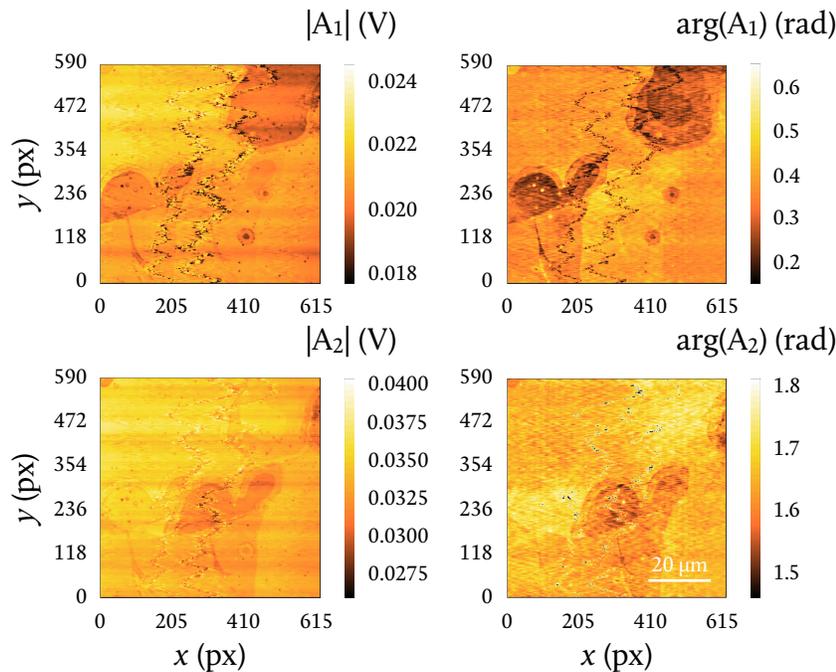


Figure 4.18: Amplitude (left) and phase (right) of the two linear combinations of D_1 and D_2 with the coefficients c_j obtained by minimising the correlation at zero displacement. The separation at the minimum-correlation point is imperfect, indicating that the correlation approach is inappropriate due to the D_j being complex.

the range they can realistically occupy is much smaller, but those boundaries were chosen to avoid having to guess how small the range could realistically be.

For this set of data, the minimum correlation is about 3.31×10^4 and is reached at $(c_1, c_2, c_3) = (-0.64 + 1.89i, 0.20 + 0.87i, 1)$. Figure 4.18 shows the result of substituting these coefficients and the D_j (*not* the $(D_j)_F$; these are only used to prevent contributions from spatially extended features to the correlation) into equations 4.5. Although in the co-polarised images (top) the left copy of each feature is much stronger than the right copy and in the cross-polarised images (bottom) the reverse is true, both copies are still present in all images, indicating that the separation is imperfect. It should also be noted that $|c_1| = 2$ and $|c_2| = 0.90$; while the latter might be reasonably close to 1 (without knowing how much light is lost by reflection at the Wollaston prism's surfaces, it is impossible to estimate how close to 1 the magnitudes of the c_j can realistically be), the former is clearly not.

The correlation algorithm was tried on multiple experimental data sets with similar results. In every case, the separation achieved at the lowest-correlation point was imperfect.

Because the A_j are complex, for certain combinations of the c_j there can be negative contributions to the cross-correlation even when the images are not separated. This can

result in a situation where there is no minimum at the position in (c_1, c_2, c_3) space where perfect separation would be achieved. Indeed, the cross-correlation method was unable to locate the correct values of the c_j even with simulated data (see appendix B). Experimental data brings the additional complication that noise in the data results in noise in the correlation. This explains the difficulties encountered with this method.

Gradient descent

The gradient descent and genetic algorithms both rely on a different method of determining whether the images are properly separated or not. This method consists of comparing the original data sets with reconstructed data sets created under the assumption that the images are separated and then changing the parameters used in the reconstruction to attempt to improve it.

First, the A_j are calculated from the D_j and the c_j as before. Since the two separated images should be identical (up to a multiplicative constant) but displaced with respect to each other by the Wollaston prism shear, it is assumed this is the case for the A_j and a single A is calculated as the average of the A_j , each displaced by half the shear towards the other:

$$A(\vec{r}) = \frac{A_1\left(\vec{r} + \frac{\vec{r}_0}{2}\right) + A_2\left(\vec{r} - \frac{\vec{r}_0}{2}\right)}{2},$$

where \vec{r}_0 is the shear vector. If the c_j are the correct ones and the polarisations are indeed properly separated, A will have a single copy of every structure at the centre of where the two copies are in the D_j (figure 4.19, left); otherwise, it will have three copies of every structure (figure 4.19, right). Then, a reconstruction of the original data sets is attempted

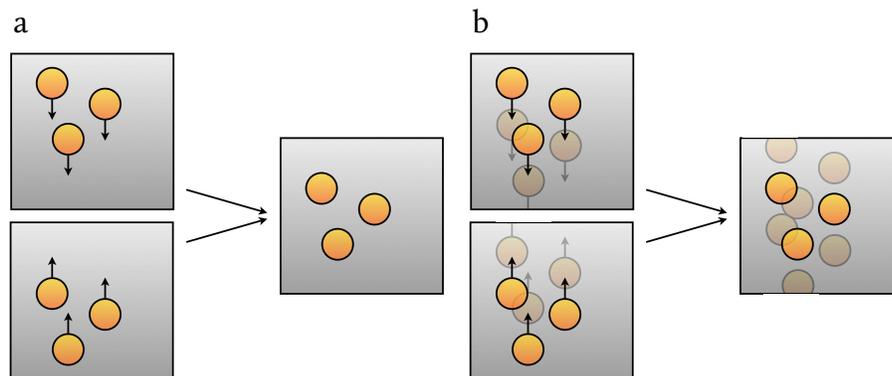


Figure 4.19: Construction of A from A_1 and A_2 in the cases of perfect (a) and imperfect (b) separation. In the former case, a single image is formed; in the latter case, three copies of the image are present.

using A:

$$D'_1(\vec{r}) = \frac{c_3 A\left(\vec{r} - \frac{\vec{r}_0}{2}\right) - c_1 A\left(\vec{r} + \frac{\vec{r}_0}{2}\right)}{(1 - c_1 c_2) c_3},$$

$$D'_2(\vec{r}) = \frac{A\left(\vec{r} + \frac{\vec{r}_0}{2}\right) - c_2 c_3 A\left(\vec{r} - \frac{\vec{r}_0}{2}\right)}{(1 - c_1 c_2) c_3}.$$

Finally, the images are compared and the error, which quantifies the difference between the original and reconstructed images and is denoted here by E, is calculated as

$$E = \sum_{\vec{r}} \frac{\operatorname{Re}^2\left(\log\left(\frac{D_1}{D'_1}\right)\right) + \operatorname{Re}^2\left(\log\left(\frac{D_2}{D'_2}\right)\right)}{\sigma_r} + \sum_{\vec{r}} \frac{\operatorname{Im}^2\left(\log\left(\frac{D_1}{D'_1}\right)\right) + \operatorname{Im}^2\left(\log\left(\frac{D_2}{D'_2}\right)\right)}{\sigma_\varphi}, \quad (4.7)$$

where $1/\sigma_r$ and $1/\sigma_\varphi$ are weights added to the sum of the real parts and the sum of the imaginary parts, respectively, of the logarithms; recall that the real part of a logarithm is related to the argument's amplitude and the imaginary part of a logarithm is related to the argument's phase. In our case, these weights are calculated simply as the average of the amplitude and phase of the original data sets in the regions containing no layer and serve to take into account differences in the amplitude and phase noise levels; for simulated data, $\sigma_r = \sigma_\varphi = 1$.

Since the information outside the image is unknown, the values at the edge of A (and thus those at the edge of the reconstructed images D'_j) are unknown. Therefore, the edge of the D_j is ignored when comparing the original and reconstructed images.

After calculating E, the gradient descent algorithm approximates its gradient, ∇E , by calculating E at small positive and negative displacements in each of the six parameters—the $|c_j|$ and the $\arg(c_j)$ —individually. It then moves the vector $\vec{c} = (c_1, c_2, c_3)$ a step of size $\alpha\gamma$ in the direction opposite to this gradient:

$$(c_1, c_2, c_3)_{N+1} = (c_1, c_2, c_3)_N - \alpha_N \gamma_N \nabla E|_{(c_1, c_2, c_3)_N}.$$

Finally, it calculates E at this new point. If it is smaller, it assumes an improvement has been made; otherwise, it discards the attempt and retries with a smaller step size. The value of γ on the N -th iteration, γ_N , is calculated using the Barzilai-Borwein method⁶ unless this results in an increase in step size (in which case it simply equals γ_{N-1}):

$$\gamma_N = \frac{|(\vec{c}_N - \vec{c}_{N-1})^\top \cdot (\nabla E_N - \nabla E_{N-1})|}{|\nabla E_N - \nabla E_{N+1}|^2},$$

where \vec{c}_N is \vec{c} on the N -th iteration, \top indicates the transpose of a vector or matrix, and $E_N = E(\vec{c}_N)$. α is a parameter which tells the algorithm when to stop; $\alpha_{N+1} = 2^{\frac{1-m}{2}} \alpha_N$ (where m is the number of failed, or discarded, attempts at improvement on the current iteration), and the algorithm stops when $\alpha < \alpha_{\text{th}}$ for some threshold α_{th} . Changing the value of α allows the algorithm to dynamically increase the step size if it improves on its first try on an iteration, which means it has yet to reach the minimum of E but it is moving in the right direction, and to decrease it if it has trouble making an improvement, which means it is taking steps which are too large, in turn potentially signifying that the algorithm is close to the minimum, where it is easy to overshoot. The algorithm stops when it has failed to make an improvement with an extremely small step size. In our case, we set $\alpha_0 = 1$ and $\alpha_{\text{th}} = 10^{-9}$.

Here we present an example of the gradient descent algorithm applied to experimental data. Figure 4.20 shows the data, which, as previously, was minimally phase-filtered before undergoing the gradient descent analysis in order to remove the phase drift, which completely occluded most of the spatial features. This time, the shear was approximately vertical; it was found to be 4 pixels horizontally and 218 pixels vertically.

Like the correlation between the $(A_j)_f$, the error E is a function of 6 variables: the amplitudes and phases of the c_j . Figure 4.21 shows projections of E , calculated using equa-

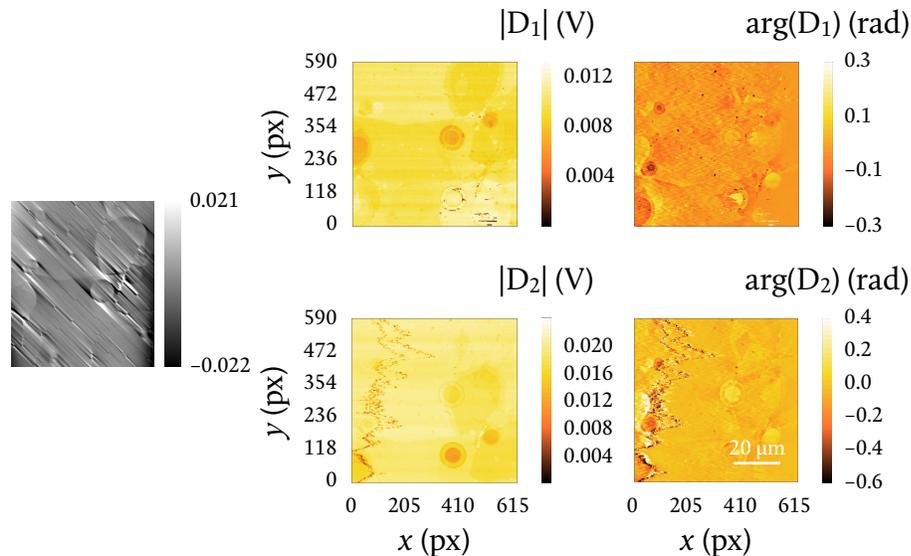


Figure 4.20: Amplitude (centre) and phase (right) of the co- (top) and cross-polarised (bottom) components of the signal of a ternary (DOPC:sm:ch = 11:5:4) bilayer detected using a Wollaston prism for referencing. The bilayer composition (DOPC:sm:ch) is 11:5:4. The shear direction is approximately vertical. The phase has been minimally filtered as described in section 4.2.1 to remove the phase drift and the fast-axis scan frequency (3.858 Hz), but most of the noise remains. A quantitative DIC image of the same field of view is shown on the left.

tion 4.7, into various 2-dimensional subspaces of (c_1, c_2, c_3) space. These projections are the same as those shown in figure 4.17. In this case, there appears to be a minimum around $(-i, -2i, 1)$.

The main drawback of gradient descent algorithms is that they can become trapped in local minima, as often happened with our data. Unlike that of the correlation analysis, which is completely deterministic, the result of the gradient descent analysis depends on the starting point $(c_1, c_2, c_3)_0$, so it may need to be run several times for the global minimum of E to be found. Figure 4.22 shows the evolution of E and of the six

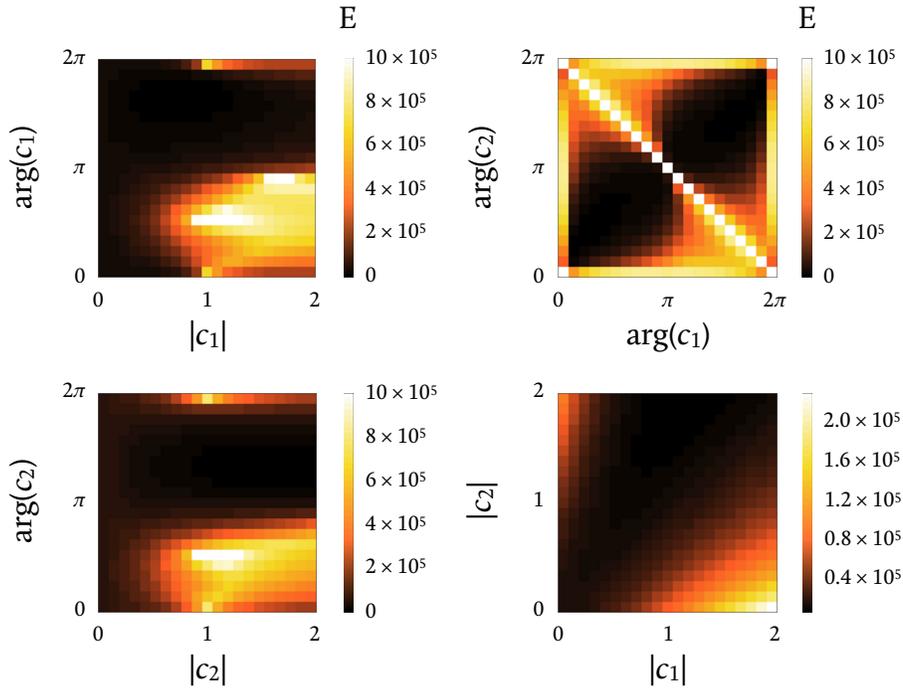


Figure 4.21: Several projections of the data reconstruction error in the gradient descent method with the experimental data shown in previous figures in this appendix. In $(|c_1|, \arg(c_1), |c_2|, \arg(c_2), |c_3|, \arg(c_3))$ space, the projections are as follows:

$$\begin{aligned} & \left(|c_1|, \arg(c_1), 1, \frac{3\pi}{2}, 1, 0 \right) \quad \left(1, \arg(c_1), 1, \arg(c_2), 1, 0 \right) \\ & \left(1, \frac{3\pi}{2}, |c_2|, \arg(c_2), 1, 0 \right) \quad \left(|c_1|, \frac{3\pi}{2}, |c_2|, \frac{3\pi}{2}, 1, 0 \right) \end{aligned}$$

With the exception of $(1, \arg(c_1), 1, \arg(c_2), 1, 0)$, the projections look qualitatively similar to the same projections for simulated data (figure C.2), but the error is much greater for experimental data, partly because the images are almost 227 times larger by area and partly because the noise present in the data makes the reconstruction less perfect.

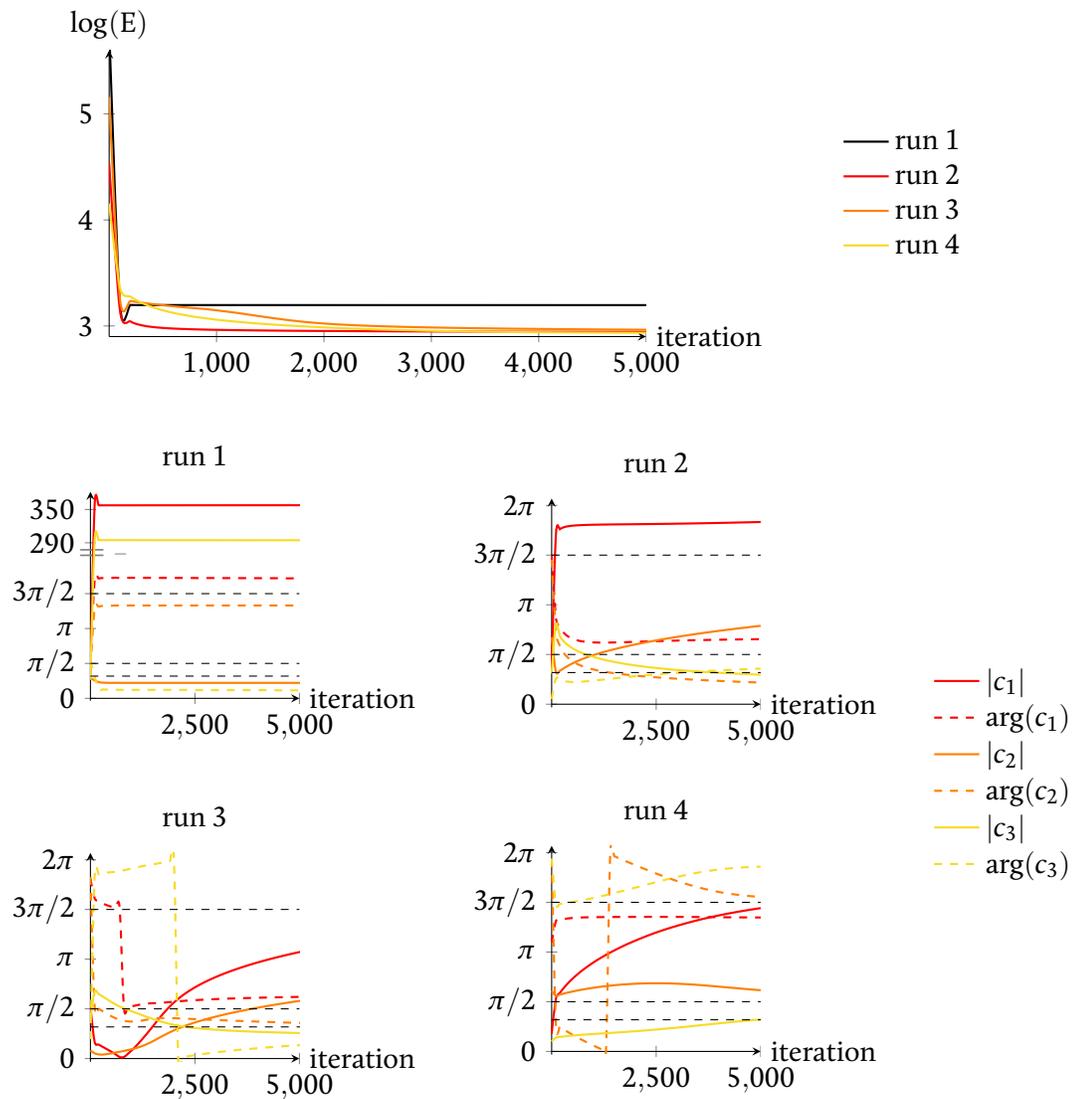


Figure 4.22: Evolution of the error E (top) and the parameters (centre and bottom) in four different runs with the same simulated data. The starting points were random. In runs 1 and 2, the parameters started with $|c_1|, |c_3| \in [0.8, 1.2]$; in run 1, $|c_2|$ started in the same range, while in run 2 it started in the range $[1.8, 2.2]$; in run 1, the starting point had $\arg(c_j) \in [\pi/2 - 0.2, \pi/2 + 0.2]$ for all j , while in run 2 it had $\arg(c_1), \arg(c_2) \in [3\pi/2 - 0.2, 3\pi/2 + 0.2]$ and $\arg(c_3) \in [-0.2, 0.2]$. In runs 3 and 4, the starting points had $|c_j| \in [0, 2]$ and $\arg(c_j) \in [0, 2\pi]$ for all j . Ironically, in runs 1 and 2 $|c_1|$ and $|c_3|$ reached unrealistically large values for numbers which are meant to be candidates for the correct amplitudes of the parameters; it is expected that these are much closer to 1. All runs were stopped after approximately 5,000 iterations. In neither case was perfect separation achieved (see figure 4.23). The black dashed lines in the centre and bottom graphs mark the values $1, \pi/2$ and $3\pi/2$.

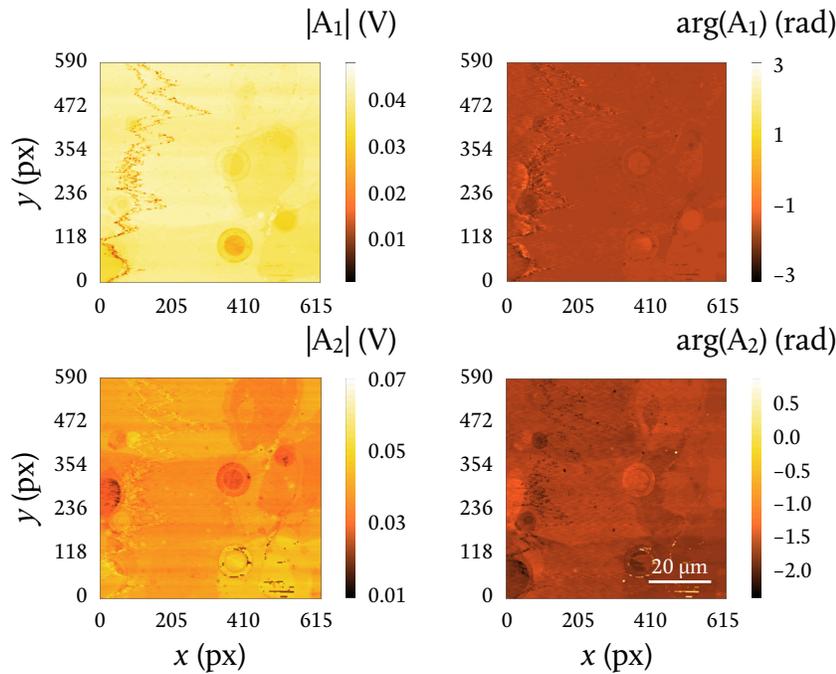


Figure 4.23: Amplitude (left) and phase (right) of A_1 (top) and A_2 (bottom) calculated using equation 4.5 and the final parameters of run 4. The imperfect separation of the polarisations indicates that the algorithm became trapped in a local minimum of the error. The fact that this happened on all runs with all experimental data indicates that this approach is also unsuitable.

parameters in four runs of the algorithm on the same experimental data. The starting point of the first run had $|c_j| \in [0.8, 1.2]$ and $\arg(c_j) \in [\pi/2 - 0.2, \pi/2 + 0.2]$ for all j . The starting point of the the second run had $|c_1|, |c_3| \in [0.8, 1.2]$; $|c_2| \in [1.8, 2.2]$; $\arg(c_1), \arg(c_2) \in [3\pi/2 - 0.2, 3\pi/2 + 0.2]$; and $\arg(c_3) \in [-0.2, 0.2]$. The third and fourth runs had starting points in the extended region defined by $|c_j| \in [0, 2]$ and $\arg(c_j) \in [0, 2\pi]$. All runs were stopped after approximately 5,000 iterations, as the error had ceased to decrease; the $|c_j|$ had not yet converged, but it seemed unlikely that they would — letting the algorithm run for 40,000 iterations did not result in the error reaching lower values than about 850, and separation was not achieved.

Figure 4.23 shows the final result of run 4 from figure 4.22. As can be seen in the figure, separation was not achieved. As with the example shown for the correlation method, this is representative of all runs attempted with all experimental data sets; perfect separation was never achieved with any data set.

As mentioned earlier, gradient descent algorithms are vulnerable to becoming trapped in a local minimum of the function to be minimised and consequently being unable to find the global minimum. The fact that we did not know which region of (c_1, c_2, c_3) space the starting point of a run should be in (i.e. which region the global minimum was in) made the problem worse, as starting closer to a local minimum than to the global minimum has

a high chance of resulting in the algorithm falling into the local minimum. With no way to properly visualise E as a function of all 6 variables, it is impossible to know how many local minima exist and where they are.

Genetic algorithm

One way to overcome the problem of local minima is to have multiple starting points and repeat the algorithm for each of them, hoping at least one of them will reach the global minimum. Another way is to use an algorithm that is less vulnerable to local minima, such as a genetic algorithm.

Instead of taking a single initial point, the genetic algorithm requires a fairly large collection of them (250 were used in our case). They could be chosen randomly, as in the gradient descent algorithm, but, because the genetic algorithm is less vulnerable to local minima, it can instead use said local minima as starting points, making convergence faster. In the case of simulated data (see appendix D), the algorithm partitions a region in (c_1, c_2, c_3) space and finds the local minima of the error in the partition, then takes the

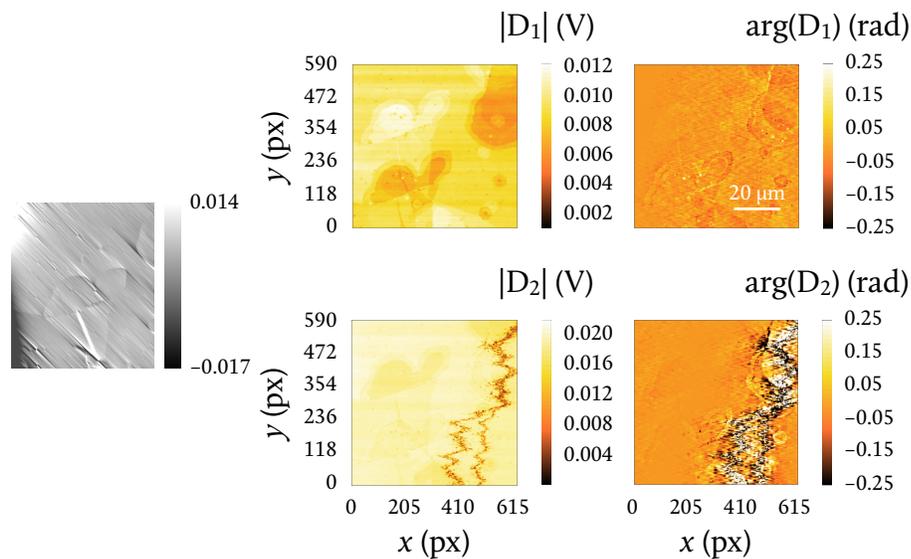


Figure 4.24: Amplitude (left) and phase (right) of the two polarisation components of the signal from a sample which was measured using the Wollaston prism for referencing. The sample is the same as that shown in figures 4.15, 4.16 and 4.18, but this is a different measurement with approximately vertical, instead of horizontal, shear. As before, the combination of polarisations is visible as a duplication of every feature in the direction of the shear (some features are only visible once because their copies lie outside the field of view). The phase has been minimally filtered to remove the phase drift of several radians, which completely obscured the features, and the fast-axis scan frequency (3.858 Hz), but most of the noise remains because excessive filtering results in loss of information.

250 best local minima (or, if there are $N < 250$ local minima in the partitioned region, it takes all N and then the $250 - N$ next-best points in the partition) as the initial population. For experimental data, the starting population is selected at random, since the position of the global minimum of E is unknown and partitioning the (c_1, c_2, c_3) space to find the approximate positions of the local minima would take an excessive amount of time.

A genetic algorithm receives its name from the way it simulates a living population and its genetics. Members of the population “reproduce” and “die”, with some selection mechanism ensuring that only the fittest members survive. The children’s traits are combinations of their parents’ traits and can make them fitter or less fit. Fitter individuals have a higher chance to reproduce and a lower chance to die than individuals which are less fit.

Each member of the initial population of points in (c_1, c_2, c_3) space, which we call the parent generation, has a fitness. In our case, if E is the error (given by equation 4.7, as in the case of the gradient descent procedure) of the point in question, its fitness is given by $F = -\log_{10}(E)$.

After determining the parent population’s fitness, a number of reproduction events (in our case, 1; we also tried having 10 reproduction events per generation, but the results were similar) takes place. A reproduction event results in two individuals of the parent generation combining their genetic material to produce one or more children (in this case, we have a single child, each of whose coordinates in (c_1, c_2, c_3) space is a random linear combination of the parents’ corresponding coordinates, meaning the child can be any point in the hyper-cuboid whose opposite vertices are the child’s two parents); a number of members of the parent generation equal to the number of children produced then dies. The likelihood P_r that an individual with fitness F will reproduce in each reproduction event is

$$P_r = \frac{F - F_{\min}}{\sum_{j=1}^{250} (F_j - F_{\min})},$$

where the F_j are the fitnesses of the parents and F_{\min} is the fitness of the least-fit individual; the least-fit individual thus has no chance of reproducing. The likelihood P_d that an individual with fitness F will die after a reproduction event is

$$P_d = \frac{F_{\max} - F}{\sum_{j=1}^{250} (F_{\max} - F_j)},$$

where F_{\max} is the fitness of the fittest individual; the fittest individual thus has no chance of dying, which ensures that each generation is at least as good as the previous one (using the fitness of the fittest individual as a metric, since we only need to reach the global minimum of E once).

After the desired number of reproduction events has occurred, the surviving parents and all the children form the new parent generation. In this way, every generation has the same number of members.

It is also possible to introduce mutations. This serves to diversify the population and make it more resistant to becoming trapped in a small space around a local minimum. Mutations occur when a child is created, before the individual who dies to make place for it is determined; in this way, a particularly unfit child has itself a chance to die, rendering that particular reproduction event void. A mutation is a change from the point the child would have been to a random point and occurs with probability $P_m = \min(1, P_0 e^{-\alpha \delta})$, where P_0 and α are positive numbers and δ is the cartesian distance between the “proposed” mutation and the point the child would be if unchanged. Parents are immune to mutations; only children can mutate, and only at the moment of creation.

The algorithm stopped after 100,000 consecutive iterations without improvement.

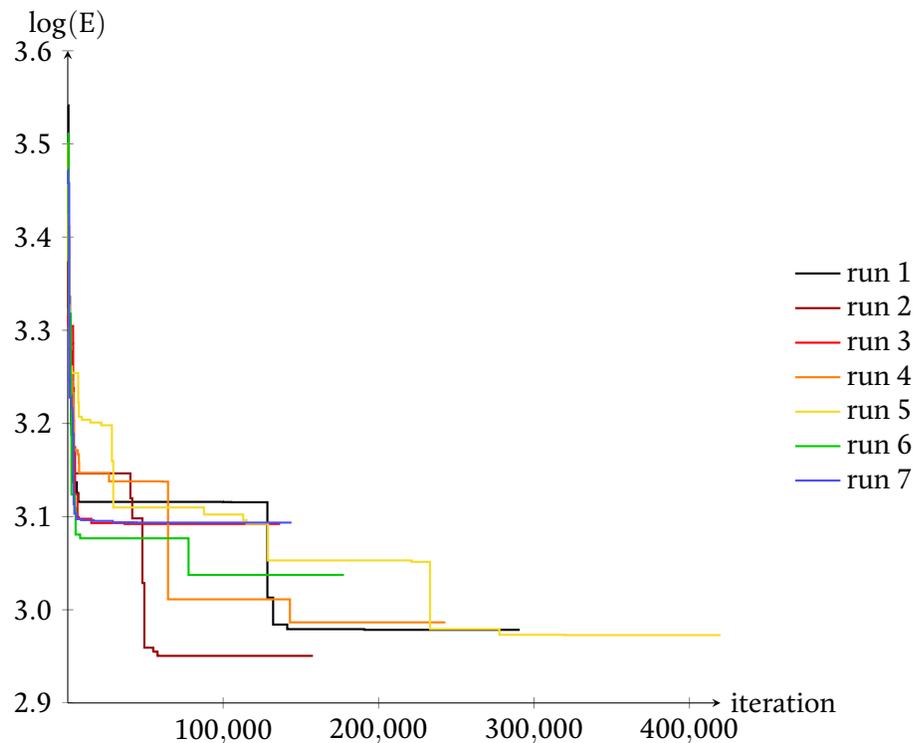


Figure 4.25: Evolution of the error E in seven different runs with the same experimental data shown in figure 4.24. The initial population was randomly selected and was different in each run due to the vast amount of time computing the 250 best local minima would have taken (unless the partition of (c_1, c_2, c_3) space were very coarse, which would defeat the purpose of partitioning it anyway). In this case, run 2 resulted in the lowest error; the result of that run (the A_j calculated with the point obtained in that run) is not properly separated, as shown in figure 4.26.

Figure 4.24 shows an example of experimental data, which was again minimally phase-filtered. The shear was approximately vertical.

The random nature of reproduction and mutation means that different runs of the algorithm on the same data with the same initial population (i.e. the same 250 points in the first generation) can produce different results. Figure 4.25 shows the results of 7 runs of the genetic algorithm on the same data. Most runs took between 150,000 and 300,000 iterations and reached errors just under 10^{-3} , similar to the best results reached with the gradient descent analysis. Similar errors were obtained with other data sets.

Figure 4.26 shows the amplitude and phase of the A_j calculated using the result from run 2 with equation 4.5. Unfortunately, the polarisations are not adequately separated; the residual structures with the wrong polarisation information have an amplitude of about 0.079 and 0.036 (a difference from the no-layer amplitude of 3.5% and 11.7% of the no-layer amplitude) in A_1 and A_2 , respectively. This lack of adequate separation of the polarisations was the case with all the experimental data analysed.

Our genetic algorithm was, therefore, also unable to separate the polarisations. It would reach a generation where improvement, defined as an increase in the highest fit-

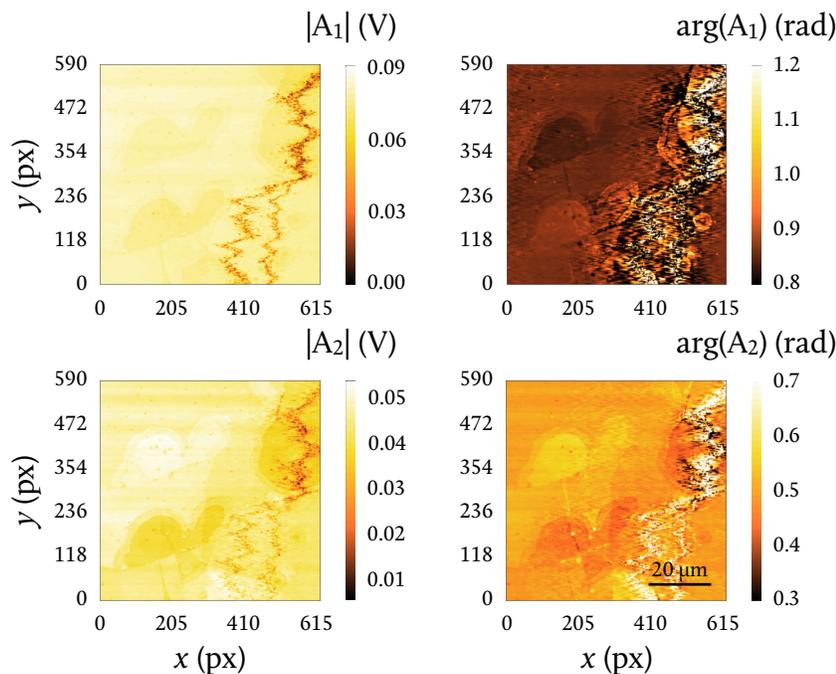


Figure 4.26: Amplitude (left) and phase (right) of the A_j obtained from the result of run 2 from figure 4.25. The polarisation is imperfectly separated; the residual error from the structures corresponding to the wrong polarisation in A_1 (top) is 0.002643, which constitutes about 3.5% of the no-layer amplitude (i.e. it will produce an error of about 3.5% in any future calculations this data is used for), while that in A_2 is 0.004808, or about 11.7% of the no-layer amplitude.

ness out of all the individuals in the generation, was difficult, and the maximum number of iterations without improvement would elapse. Increasing this number threefold, from 100,000 to 300,000, did not help, and increasing it any further would have made the algorithm take way too long to run.

This is believed to have happened due either to failing to reach the global minimum of the error (where the images are perfectly separated) or to the minimum not corresponding to the point where the polarisations are separated (it is possible that the presence of strong noise, which is complex and predominantly affects the phase, caused the error to increase unevenly at all points in (c_1, c_2, c_3) space, resulting in the point where separation would be achieved no longer having the smallest error; if this is the case, the selection criterion in the algorithm would have favoured the deletion of the separation point from the population, making it impossible to attain separation). In truth, because this doesn't happen for simulated data and the separation point is unknown for experimental data, it is not possible to know for certain what the reason is.

There was no appreciable difference between the results of running the genetic algorithm in cartesian coordinates and those of working in polar coordinates.

4.3. Line trace analysis

Unfortunately, the failure of all three methods at disentangling the polarisations meant the referencing approach was unviable. Similarly, the noise remaining after an adequate amount of phase filtering caused very large artefactual spatial variations in the measured thickness and refractive index of the sample. The idea of a 2-dimensional map of thickness and refractive index had to be abandoned in favour of a more modest approach, that of line traces.

After phase filtering, the data was regularised, and the real and imaginary parts of the co-polarised channel were saved for line trace analysis.

In general, the phase noise observed is low-frequency noise. Therefore, the noise is much less along the fast scan axis. 61-pixel line traces perpendicular to bilayer edges were selected, ensuring that roughly half of each trace lay in a region without bilayer and half passed through a bilayer (figure 4.27). In principle, at the centre of each line trace should be a step in both amplitude and phase. Bilayer edges are frequently lined by vesicles or debris, as can be seen, for example, in the amplitude image in figure 4.10 (they are visible as dark spots and are prominent along the edge of the bilayer at the top-left corner of the image); when present, these result in a dip in the signal on the lipid side of the step. The function

$$\Theta(x) = \frac{a}{2} \tanh\left(\frac{x-b}{c}\right) + dx + e + \frac{f}{2} \operatorname{sech}\left(\frac{x-g}{c}\right)$$

was fitted to the N_p central pixels of the amplitude and phase traces simultaneously, with N_p varying between 21 and 51 in steps of 2 and the optimal N_p for each line trace then being selected by least-squares error of the fit. All of the traces were oriented so they

started ($x = 0$ px) within the bilayer and ended ($x = 60$ px) in the no-bilayer region in order to ensure $a > 0$ for both amplitude and phase. The hyperbolic tangent in the step function Θ above models the step itself, while the hyperbolic secant models the dip caused by edge vesicles. The spatial parameters of the hyperbolic functions (b , c and g) were shared between the amplitude fit and the phase fit of any given trace, since the step due to the bilayer edge and the dip caused by any present edge vesicles should be at the same positions and have the same shapes in the amplitude and phase data. The linear term represents slow drift due to low-frequency noise or the sample not being perfectly

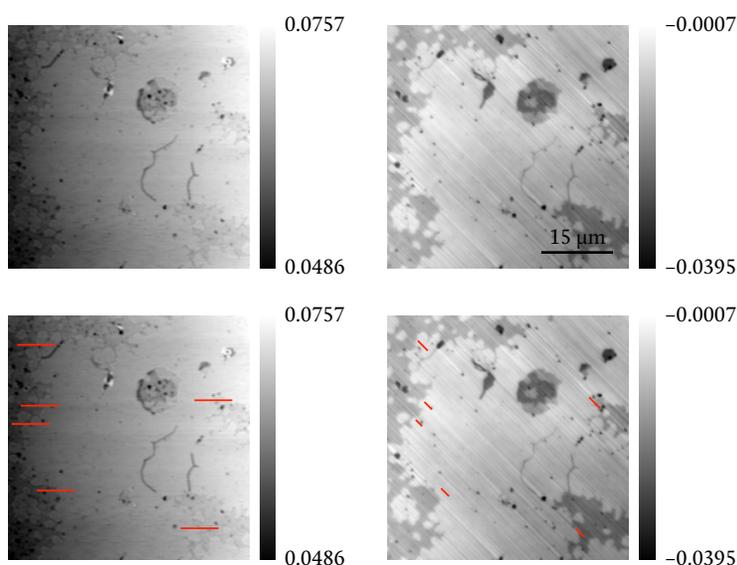


Figure 4.27: Interferometric reflectometry amplitude (left) and quantitative DIC phase (right) images of one of the $DC_{15}PC$ measurements (top). The images are repeated below with red lines indicating the positions and approximate lengths of the line traces taken. Because the reflectometry traces had to be parallel to the slow scan axis and the quantitative DIC traces had to be parallel to the DIC shear, the traces were not parallel between the two techniques; however, the position of the lipid bilayer step was preserved as closely as possible between the techniques. The grey value of the reflectometry images is in volts. The top images are included so the structure at the trace positions is clear and to facilitate comparison between reflectometry and quantitative DIC data. Structures, including vesicles and bilayer steps, are identically visible in both imaging modalities; in both cases, darker regions are thicker ones, with the lightest regions being the no-layer regions. However, in reflectometry large vesicles or otherwise non-planar objects are sometimes visible as alternating light and dark regions, while in quantitative DIC they are homogeneous (see, for example, the medium-sized structure at top-centre, left of the large patch which is partially unilamellar and partially bilamellar); this is likely because the non-planar structures reflect light at angles different from the angle of incidence with respect to the focal plane, resulting in interference patterns. Because the reflectometry model presented in section 4.1.1 and appendix A assumes the sample consists of planar layers, no attempt has been made to characterise these features.

flat against the mounting stage of the microscope. In the case of the phase traces, the constant term, e , takes care of any offset present; we were only interested in the step height a because $S_{\text{gw}} \in \mathbb{R}$. In the case of the amplitude traces, e was used alongside a to reference the no-bilayer signal to S_{gw} in order to obtain the relative amplitude of the normalised reflectometry signal at the bilayer:

$$\left| \frac{S}{S_{\text{gw}}} \right| = 1 - \frac{a}{e + a/2}.$$

This enabled the full determination of s at the bilayer (figure 4.28).

The $[0 \text{ nm}, 30 \text{ nm}] \times [n_w, n_g]$ region of the (d, n_s) space was partitioned in steps of 0.1 nm in d and 0.001 in n_s . S was calculated using equation 4.4 at each point of the partition. The $|S/S_{\text{gw}}|$ and $\arg(S)$ values calculated as described above were then compared to the partition, and the d and n_s corresponding to the closest partition point were assigned

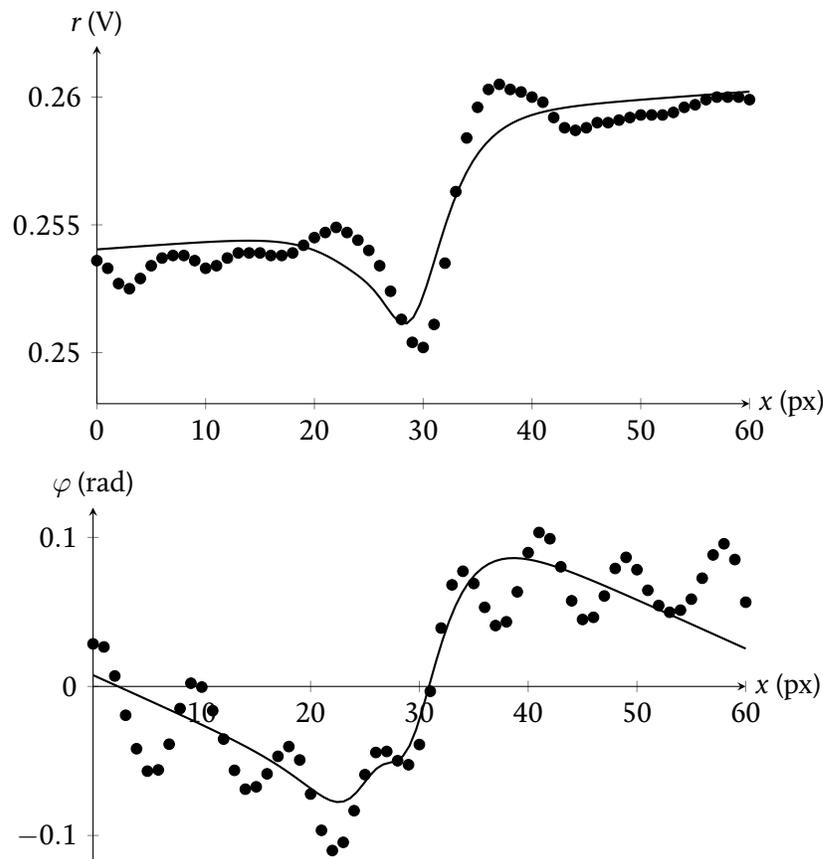


Figure 4.28: Example of a line trace. r is shown at the top, and φ is shown at the bottom, each with its step function fit Θ .

to that line trace. The computer code used for the analysis can be found in appendix E; like the codes in the other appendices, it was written entirely by the author of this work.

Line traces from seven different images were analysed; the images taken with the Wollaston prism were, of course, unsuitable for line trace analysis for the reasons described in section 4.2. Four of the images were from a pure DC₁₅PC supported lipid bilayer, while the other three were from a sample consisting of a ternary mixture of DOPC, egg sphingomyelin and cholesterol in an 11:5:4 ratio. The refractive index of DC₁₅PC is³ 1.440, and that of the LD domains of a DOPC+sphingomyelin+cholesterol mixture (which are simply DOPC; see also chapter 5) is⁴ 1.445. Using these values for the refractive indices, the thicknesses of DC₁₅PC and DOPC bilayers have been measured with quantitative DIC to be (5.3 ± 0.1) nm and (4.1 ± 0.03) nm, respectively⁷ (uncertainties are the standard error). Samples were imaged at room temperature, about 20 °C. All line traces were parallel to the fast scan axis; in the case of the ternary sample, care was taken to ensure the traces did

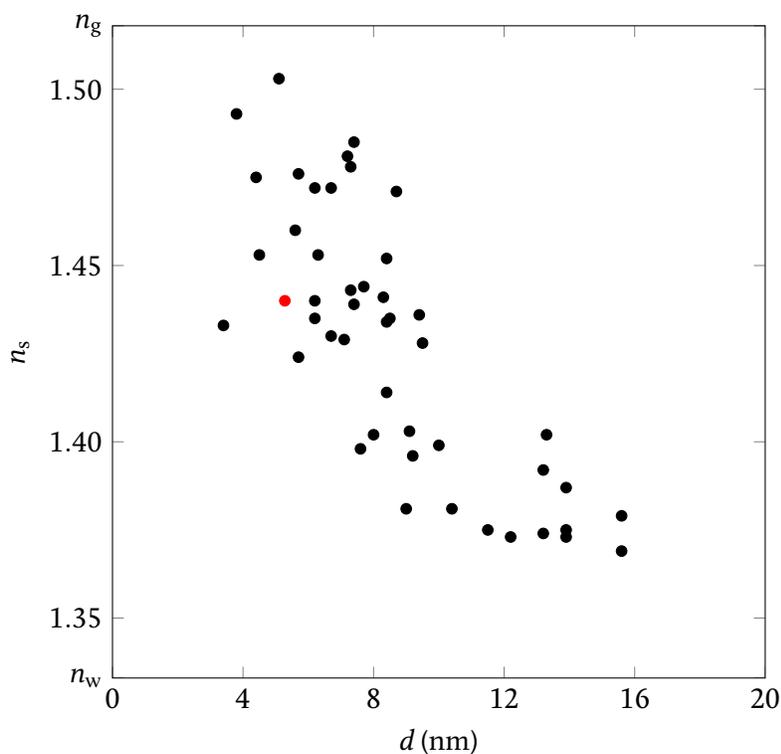


Figure 4.29: Results from the DC₁₅PC sample. Each point is an individual measurement (i.e. the result of a single line trace). The results show great variability in both thickness and refractive index. There seems to be some correlation between d and $1/n_s$. The red circle indicates the nominal thickness and refractive index as measured by quantitative DIC by others;⁷ the error bar (± 0.1 nm) is too small to be seen at the scale shown.

not cross LO domains. Any line traces where the sum of the squares of the residuals of the fit was greater than 0.01 were discarded; this almost always coincided with the fits which resulted in very large thicknesses and very small refractive indices. Line traces where it was difficult to see the step clearly by eye in both the amplitude data and the phase data were also discarded.

For the DC₁₅PC sample, the measured thickness was (8.6 ± 3.4) nm and the measured refractive index was 1.428 ± 0.041 (46 traces, figure 4.29). For the ternary mixture, the thickness was measured as (12.4 ± 3.0) nm and the refractive index was measured as 1.429 ± 0.030 (36 traces, figure 4.30). Errors reported are the standard deviation. In both cases, the results are roughly what would be expected from a bilamellar or trilamellar lipid instead of a unilamellar one, which quantitative DIC measurements suggest the samples were. As mentioned earlier, a unilamellar DC₁₅PC sample should have a thickness of 5.3 nm (assuming a refractive index of 1.440) and the LD domains of a unilamellar ternary sample with the aforementioned composition should have a thickness of 4.1 nm

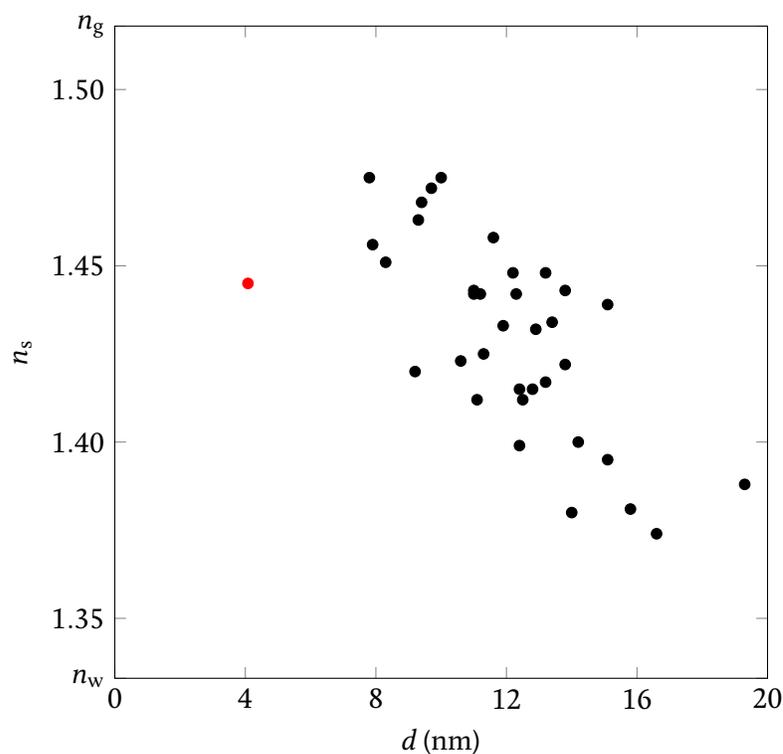


Figure 4.30: Results from the ternary sample. As in the case of the DC₁₅PC sample, the results show great variability in both thickness and refractive index. The red circle indicates the nominal thickness and refractive index as measured by quantitative DIC by others;⁷ the error bar (± 0.03 nm) is too small to be seen at the scale shown.

(assuming a refractive index of 1.445).

Two tests were carried out to determine whether the error was due to the noise or there was some systematic error causing the measured thicknesses to be too large.

The first of these tests was to compare the results with similar results for simulated data. Data was constructed as described in section 4.2.2, 50 line traces were taken parallel to what would have been the slow scan axis had the data been experimentally obtained, and the function Θ described above was fitted to the traces with the same constraints and procedure as before; the only differences were that the traces were 41 pixels in length and N_p was varied between 21 and 31. This was done for both noiseless and noisy simulated data. It was observed (figure 4.31) that without noise (figure 4.14, top right) the thickness was (4.1 ± 0.1) nm and the refractive index was 1.443 ± 0.002 , with noise about 10% as strong as seen experimentally (figure 4.14, bottom left) the thickness was (4.1 ± 0.1) nm and the refractive index was 1.442 ± 0.007 , and with noise of similar strength as the experimental noise (figure 4.14, bottom right) the thickness was (4.7 ± 1.4) nm and the

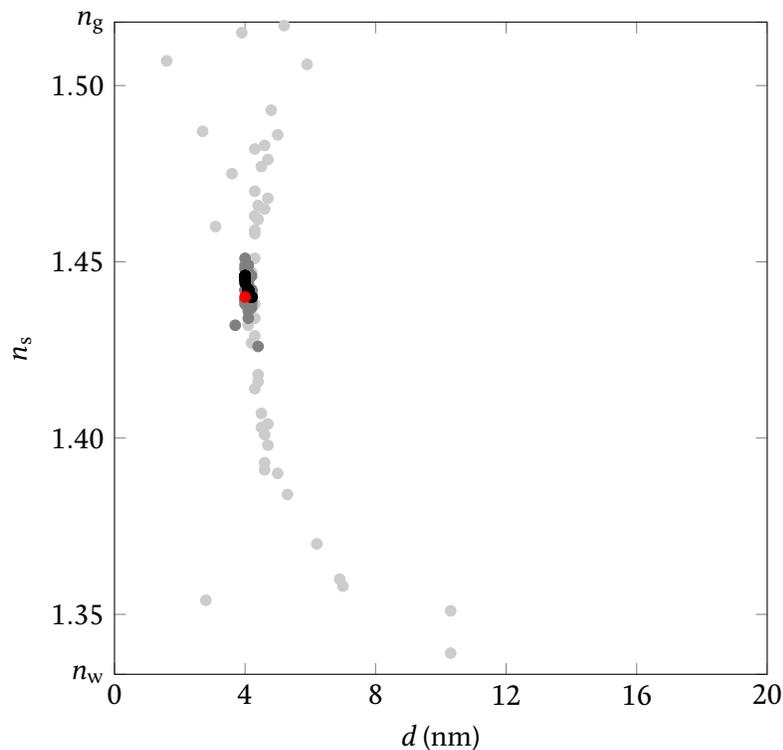


Figure 4.31: Results from simulated data without noise (black), with 10% phase noise (dark grey) and with 100% phase noise (light grey), where the noise percentage refers to the approximate strength of the experimentally observed noise. The red circle indicates the correct thickness and refractive index.

refractive index was 1.437 ± 0.049 (one point was discarded in this last case because it fell outside the range described previously). These values are all much closer to the nominal values than in the case of the two sets of experimental data, but the spread of the results becomes larger, the average thickness increases and the average refractive index decreases with increasing noise, as seen in the experimental case. The smaller deviation from the nominal result and the smaller standard deviation in the simulated data could be due to the simulated noise not representing the experimental noise perfectly; to begin with, the experimental noise is certainly not gaussian. Nevertheless, this result suggests that the large standard deviation and inaccuracy of the experimental results stem at least partially from the noise.

The second test consisted of comparing the reflectometry results to quantitative DIC results. To do this, quantitative DIC measurements of the same fields of view of the same two samples were taken and line traces were taken in such a way that they crossed from a

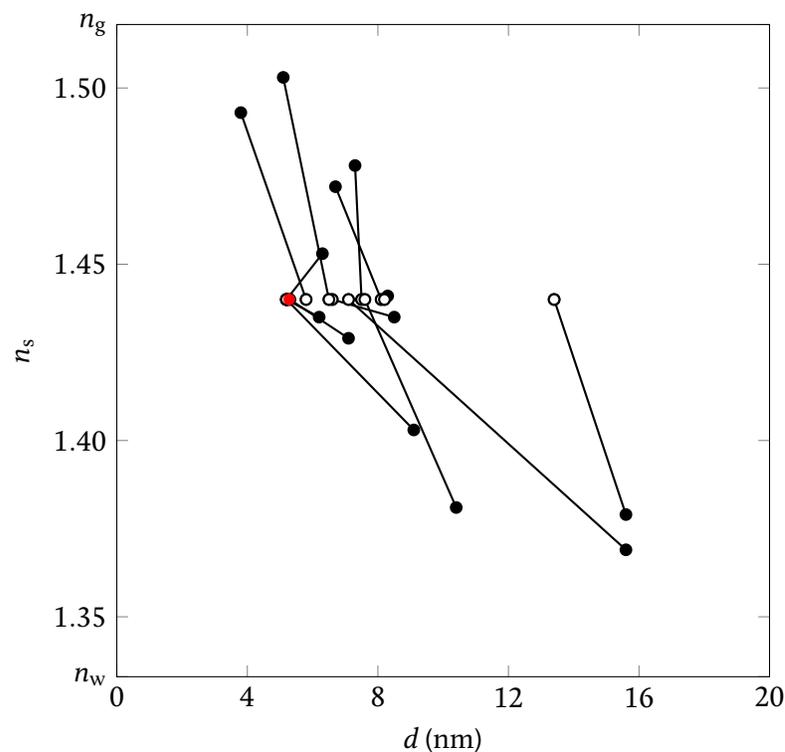


Figure 4.32: Interferometric reflectometry (solid circles) and quantitative DIC (hollow circles) results of one of the DC₁₅PC images. Each line joins the reflectometry and quantitative DIC results from the line traces at the same position. There does not seem to be a correlation between the magnitudes of the reflectometry inaccuracies and those of the corresponding quantitative DIC inaccuracies. The red circle indicates the nominal thickness and refractive index as measured by quantitative DIC by others;⁷ the error bar (± 0.1 nm) is too small to be seen at the scale shown.

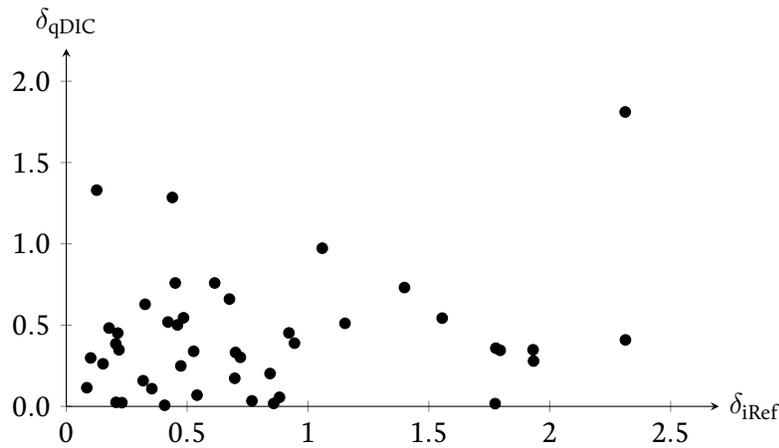


Figure 4.33: Weighted cartesian distances between the experimental and nominal results with interferometric reflectometry and quantitative DIC for the DC₁₅PC sample. The correlation coefficient is very low.

lipid bilayer to a water-only region, as with the interferometric reflectometry line traces; the function Θ shown previously was fitted to them to determine the step height. The quantitative DIC line traces were centred as close to the centres of the reflectometry line traces as possible given that they had to be taken parallel to the DIC shear direction, which was at 45° with respect to the reflectometry traces. This resulted in a thickness of (6.9 ± 1.9) nm for the DC₁₅PC sample and (5.6 ± 1.6) nm for the LD domains of the ternary sample. This is about 31% too large for the DC₁₅PC sample and 38% too large for the ternary sample due to a presently uncharacterised systematic error in the quantitative DIC measurements taken with the microscope with which reflectometry imaging was done. Assuming the quantitative DIC results are self-consistent despite this systematic error, a high correlation coefficient between the accuracy of the result of each quantitative DIC line trace and the accuracy of the result of the corresponding interferometric reflectometry line trace would indicate that some systematic error is responsible for the inaccuracy of the reflectometry measurements, while a low correlation coefficient would indicate that the inaccuracies are due to noise, which is different for reflectometry and DIC. Figure 4.32 shows the interferometric reflectometry and quantitative DIC results for one of the DC₁₅PC images; reflectometry results are marked by solid circles, quantitative DIC results are marked by hollow circles, and a line joins the reflectometry result of each line trace with the quantitative DIC result of the corresponding line trace. The correlation coefficient⁸ is calculated as

$$\rho = \frac{\sum_j (\delta_{iRef,j} - \langle \delta_{iRef} \rangle) (\delta_{qDIC,j} - \langle \delta_{qDIC} \rangle)}{\sqrt{\sum_j (\delta_{iRef,j} - \langle \delta_{iRef} \rangle)^2} \sqrt{\sum_j (\delta_{qDIC,j} - \langle \delta_{qDIC} \rangle)^2}};$$

the sums in the above expression are taken over all the line traces except those which were discarded from the calculation of the mean and standard deviation for the reasons explained in section 4.3, $\langle \delta \rangle$ denotes the average of δ over all j , and $\delta_{\ell,j}$ is the weighted cartesian distance between result j obtained with technique ℓ and the nominally correct (d, n_s) and is given by

$$\delta_{\ell,j} = \sqrt{\sigma_d(d_{\ell,j} - 5.282 \text{ nm})^2 + \sigma_{n_s}(n_{s,\ell,j} - 1.44)^2}.$$

The weights σ_d and σ_{n_s} ensure that the thickness and refractive index differences, which are of different orders of magnitude, have similar bearing on the calculation of δ ; if σ_d and σ_{n_s} were equal in magnitude, the inaccuracies in the refractive index would be negligible compared to the inaccuracies in the thickness. $\sigma_d = 0.05 \text{ nm}^{-1}$ and $\sigma_{n_s} = 1/(n_g - n_w)$ were chosen here. Taking into account all images, the correlation coefficient was 0.258 for the DC₁₅PC sample (figure 4.33) and -0.087 for the ternary sample, which is low in both cases and thus suggests, once again, that the reflectometry inaccuracies are caused by noise.

4.4. Electrophysiology

4.4.I. Lobster ventral nerve cord

As mentioned in chapter 3, the work on lobster ventral nerve cord was performed during a stay in Thomas Heimburg's laboratory in Denmark.

The lobster nerve was stimulated at room temperature with 5-V steps programmed using the MultiClamp software (Molecular Devices), and the nerve's response was visualised in the Axon pClamp software (Molecular Devices). The noise was about $0.95 \mu\text{V}$,

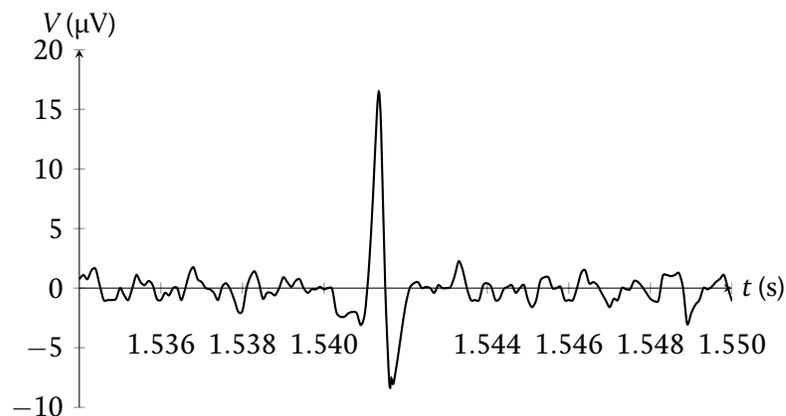


Figure 4.34: Action potential recorded extracellularly with lobster ventral cord. The amplitude of this particular action potential was $16.5 \mu\text{V}$.

and action potentials (see figure 4.34), which lasted about 1.25 ms and occurred at an average rate of one every 120 ms, had an average amplitude of 17.0 μV (recall that this was extracellular recording).

4.4.2. Mouse neurons

Evaporation experiment

Since the incubator chamber could not be used (see section 3.5.4), an experiment to determine how quickly the culture medium evaporated (and thus at which rate it needed to be replenished) was undertaken. This consisted of placing an empty coverslip in the CSC-13 chamber and filling the chamber with culture medium. One of the pipette tips was lowered until it just touched the water, which caused a change in contrast in the microscope eyepiece. This was repeated every 10 min, recording the position of the micromanipulator each time using the LinLab (Scientifica) software, which could also control the micromanipulators but was discarded in favour of manual control. The volume of the evaporated liquid was then calculated.

An evaporation rate of 132 $\mu\text{l/hr}$ was obtained; this was the average of four measurements, two of them lasting 120 min and the other two lasting 170 min. Therefore, the culture medium of the cell samples was replenished approximately every 30 min by adding 65 μl of fresh medium.

Electrophysiology experiments

Of the 36 mouse neuron samples studied, spontaneous neural activity was only detected in one: DIV-15 with GlutaMAX and a cell density of 25,000 cells/coverslip (18,800 cells/cm²) at the moment of plating (figure 4.35). It is likely that cell cultures with similar parameters (e.g. DIV-14 at the same cell density or DIV-15 at a lower cell density) also displayed spontaneous activity but it was not detected due to electrophysiology techniques being difficult to master for non-electrophysiologists. The voltage trace recorded by the reference electrode was filtered to remove slow drift, after which it was determined that the noise was about 0.3 mV.

Due to the time required to learn the electrophysiology techniques employed (section 3.5), establish the interferometric reflectometry theory (section 4.1.3), become familiar with the interferometric reflectometry equipment and experimental setup (section 3.3) and attempt to deal with the noise in the reflectometry measurements (section 4.2), the latter of which was not satisfactorily achieved due to the extremely low signal-to-noise ratio of single lipid bilayers, there was insufficient time remaining in this project to carry out the intended simultaneous electrophysiological and optical experiments. This means that our final objective, attempting to observe action potentials optically and determine whether or not they involve a phase transition of the cell membrane, was not realised. Nevertheless, future work might be able to achieve this once the signal-to-noise ratio of interferometric reflectometry measurements on single lipid bilayers is increased to at least 2,000.

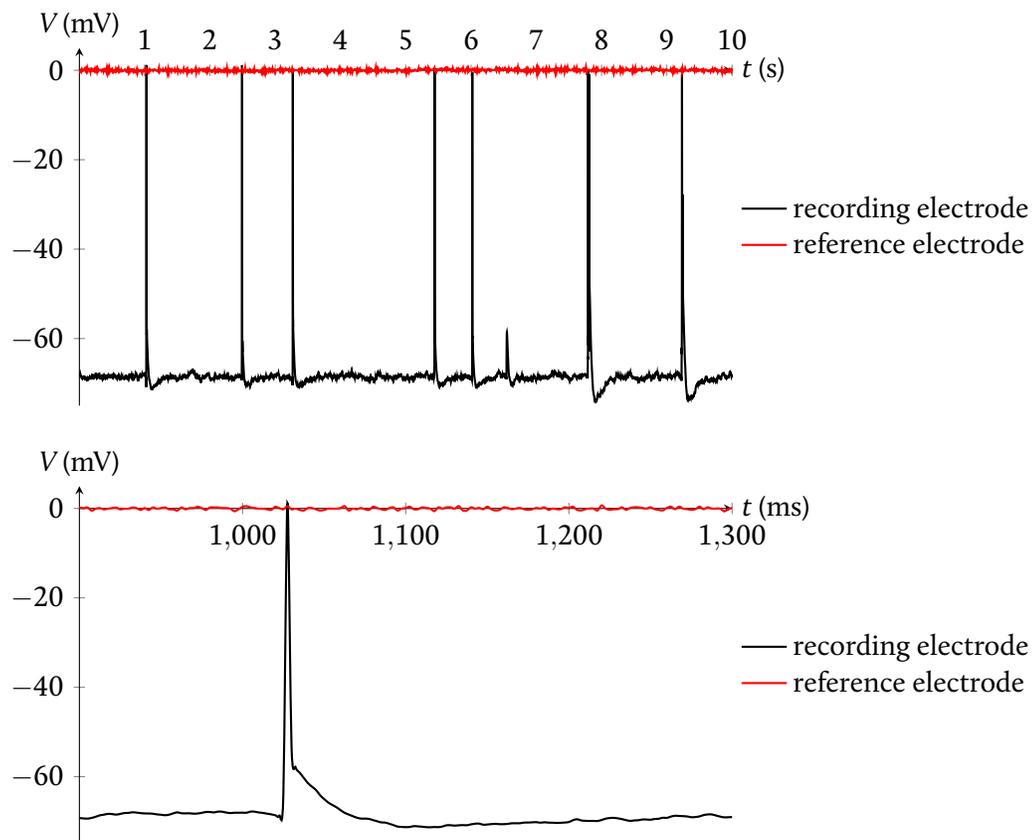


Figure 4.35: First 10 s of spontaneous activity recorded from one of the cells in the DIV-15 sample with GlutaMAX and a cell density of 25,000 cells/coverslip (top) and a 400-ms period of the same recording showing an action potential (bottom). The reference-electrode voltage, filtered to remove drift, is also shown. The voltages shown here have been divided by the gain to retrieve the actual values. The cell had a membrane resting potential of -68 mV. After the action potential at 5.44 s, there are several voltage peaks which do not cross the $V = 0$ threshold and are thus not true action potentials. Action potentials continued until about 20 s, although the fraction of peaks which did not reach $V = 0$ became larger; after 20 s, no action potentials were recorded and the peaks began to decrease in amplitude until they were only able to reach about -30 mV (an amplitude of about 40 mV). Further recordings on the same cell resulted in very small peaks for a few seconds followed by cessation of activity. This was likely due to the degradation and eventual death of the cell.

4.5. Conclusions

We have presented a complete model for the reflection of light by an isotropic, homogeneous layer placed between two semi-infinite materials. The model, which is based on Fresnel's expressions for reflection and transmission at an interface, takes into account not only the light reflected at the layer's first interface, but also the interference of the light

that enters the layer and is reflected at its interfaces any number of times. This theory is not new; it has been used, for example, for coating films or interferometers.⁹ We showed how the reflection coefficient s , which is a complex quantity, changes as a function of the thickness d and refractive index n_s of the layer; we also showed that it is periodic in d for any given value of n_s (with a period inversely proportional to n_s). More importantly, we showed how to compute d and n_s from a known value of s .

The model was generalised to the case of non-normal incidence and to the case of an angular distribution. An exact expression for the signal detected by the interferometric reflectometry imaging setup described in section 3.3, which employs a “reference” beam that never interacts with the sample but interferes with the light reflected from the layer at a collection of detectors (two for each polarisation component of the detected light) in order to allow independent measurement of the amplitude and phase of the reflected light, was obtained. The model here makes the assumptions that the layer’s refractive index is intermediate between the refractive indices of the surrounding media, that incidence occurs from the high-refractive-index side, and that the maximum angle of incidence is not large enough for total internal reflection at any of the interfaces. In this case, the media on either side of the layer no longer need to be semi-infinite, merely large enough for the focal volume to fall entirely within them. We showed how the detected signal changes in the specific case of a distribution of angles of incidence corresponding to gaussian beams and an aplanatic objective of numerical aperture 1.27, as employed in the experimental setup.

The model can be generalised to an arbitrary number of layers (appendix A), again with the conditions that all the refractive indices of the layers be between those of the external media. In this case, the mathematical expressions quickly become unwieldy, but taking only one reflection per layer into account is an acceptable approximation resulting in an error smaller than 0.1%. Since interference is employed to measure the amplitude and phase of the reflected light, any two quantities (two thicknesses, two refractive indices, or one of each where one does not necessarily belong to the same layer as the other) can be determined from the signal as long as all the others are known.

Two different techniques for dealing with noise were shown, since the changes in the signal caused by changes in d and/or n_s are of the order of a few percent and thus very clean, high-precision measurements are required to detect phase changes in single lipid bilayers. If the noise can be reduced enough, interferometric reflectometry offers the opportunity to detect such changes in a noninvasive manner, contrary to other techniques such as fluorescence and atomic force microscopy, and with enough time resolution, contrary to techniques like quantitative DIC (which requires two separate measurements to be taken with different polariser angles and can thus not achieve the required imaging speed).

One of these, called referencing, involves using a Wollaston prism placed before the microscope to split the incident beam into two spatially separated, orthogonally polarised beams which nevertheless have approximately identical noise; the images, each containing the signal corresponding to light of one polarisation, are then to be divided by each other in order to remove most of the noise. The problem this introduces is that each

image does not contain the signal from one of the polarisations, but rather combined information from both; they thus need to be separated before division. Since each image is a linear combination of the “true” (separated) images, the inverse is also true: each separated image is a linear combination of the measured images. We showed how finding the parameters for these linear combinations is a 6-dimensional minimisation problem which can theoretically be solved in a variety of ways. In practice, however, this is not necessarily true. We presented three approaches to the minimisation problem—an approach which used the cross-correlation between the images for different values of the parameters to attempt to find the correct values, an algorithm which attempted to combine the images and then reconstruct them and used gradient descent to minimise the error caused by using the wrong coefficients for the combination and reconstruction, and a genetic algorithm which attempted to minimise the same error without becoming trapped in local minima—and demonstrated the results obtained with them for both simulated data (in appendices B–D) and experimental data (in this chapter); unfortunately, even though the results of the second and third approaches with simulated data were promising, their results with experimental data were inadequate. Possible reasons for this inadequacy were suggested.

The second technique consists of filtering out spatial frequencies in the data which correspond to noise and to phase drift in an attempt to reduce the noise. This worked as expected until a certain point, after which removing more frequencies resulted in loss of information, so not all of the noise could be removed. This meant that, since images consisted of hundreds of thousands of pixels and the noise and residual phase drift in them were considerable, it was impossible to accurately determine d and n_s at each pixel.

Nevertheless, it was shown that our interferometric reflectometry setup is sensitive enough to detect variations in thickness of less than half a nanometre, which occur at the boundaries between LO and LD domains in unilamellar ternary lipid mixtures.

Line traces were also taken from the phase-filtered data and individually analysed. Each trace was taken parallel to the fast-scan axis to further decrease the noise along the trace, and each one was taken at the edge of a bilayer. The amplitude and phase data from each trace were simultaneously fitted with a step function whose height was allowed to change between the amplitude data and the phase data but whose spatial parameters were not. The height of the amplitude and phase steps gave the relative change in the reflectometry signal, while the absolute value of the amplitude in the region without a bilayer was used to normalise this relative change and obtain d and n_s .

Our results indicate that the noise remains too large for such small structures. While it is possible to see extremely small changes in the images, quantitatively computing the thickness and refractive index of the bilayer from such noisy data cannot be done with high accuracy. The measured thickness of a DC₁₅PC bilayer was (8.6 ± 3.4) nm, while its nominal thickness is 5.3 nm; for a DOPC bilayer (the LD domain of a ternary mixture of DOPC, sphingomyelin and cholesterol), the measured thickness was (12.4 ± 3.0) nm and the nominal thickness is 4.1 nm. The reflectometry measurements, then, are 2–3 times larger than they should be. Tests we made with simulated data and with quantitative

DIC measurements of experimental data from the same positions of the same samples indicate that these errors are due to the residual noise and not to a systematic error in the measurement or analysis procedure of the interferometric reflectometry technique.

Interferometric reflectometry has previously been used to measure snow depth from satellites.¹⁰ A similar technique, imaging ellipsometry (which uses changes in the polarisation of the reflected light), has been used to measure the thickness of artificial lipid bilayers and detect thickness differences between different thermodynamic phases of the lipid;¹¹ however, this requires previous knowledge of the refractive index of the lipid and the assumption that it does not change from one phase to another, which we have seen is not generally true. To our knowledge, this is the first time imaging of layers only about 5 nm in thickness and measurement of nanometric thickness differences have been attempted using interferometric reflectometry.

Spontaneous neural activity was measured electrophysiologically in cultured mouse hippocampal neurons, indicating that these neurons provide an adequate system for the study of action potentials. However, optical observation of action potentials was not achieved because it would have required the interferometric reflectometry measurements to accurately provide quantitative information of the thickness and refractive index (the latter of which can be used as a measure of lateral density) of lipid bilayers. Since the characteristics the neurons must possess have been determined and the interferometric reflectometry theory has been developed for the special case of the experimental configuration described in chapter 3, this is the only obstacle remaining between the current state of affairs and an attempt to observe action potentials with reflectometry. Other than the low signal-to-noise ratio, this technique is adequate for the problem; it can measure at a rate of 5,000 Hz, higher than the minimum rate of 2,000 required for Nyquist sampling of millisecond-long action potentials, and it is noninvasive.

In future, we plan to attempt the measurement of the thickness and refractive index of thicker samples, where the contrast between the layer and no-layer regions is stronger and thus the signal-to-noise ratio is higher. If the inaccuracies in the measurements presented here are indeed due to the noise, as we believe they are, then the retrieval of d and n_s for a sufficiently thicker sample (of the order of several tens of nanometres) should be much more accurate.

4.6. References

1. Hecht E (2002): *Optics*, fourth edition, Addison Wesley, 66
2. Ochei J & Kolhatkar A (2008): *Medical laboratory science: theory and practice*, Tata McGraw-Hill, 446
3. Howland M C, Szmodis A W, Sanii B & Parikh A N (2007): *Characterization of physical properties of supported phospholipid membranes using imaging ellipsometry at optical wavelengths*, Biophysical Journal 92, 1306–1317

4. Devanathan S, Salamon Z, Lindblom G, Gröbner G & Tollin G (2006): *Effects of sphingomyelin, cholesterol and zinc ions on the bringing, insertion and aggregation of the amyloid A β_{1-40} peptide in solid-supported lipid bilayers*, FEBS Journal 273, 1389–1402
5. Novotny L & Hecht B (2012): *Principles of nano-optics*, second edition, Cambridge University Press, 56–61
6. Barzilai J & Borwein J M (1988): *Two-point step size gradient methods*, Institute of Mathematics & its Applications Journal of Numerical Analysis 8, 141–148
7. Regan D, Williams J, Borri P & Langbein W (2019): *Lipid bilayer thickness measured by quantitative DIC reveals phase transitions and effects of substrate hydrophilicity*, Langmuir 35, 13805–13814
8. Wright S (1921): *Correlation and causation*, Journal of Agricultural Research 20, 557–585
9. Born M & Wolf E (1999): *Principles of optics*, seventh edition, Cambridge University Press, 360–361
10. Gutmann E D, Larson, K M, Williams M W, Nievinski F G & Zavorotny V (2012): *Snow measurement by GPS interferometric reflectometry: an evaluation at Niwot Ridge, Colorado*, Hydrological Processes 26, 2951–2961
11. Howland M C, Szmodis A W, Sanii B & Parikh A N (2007): *Characterization of physical properties of supported phospholipid membranes using imaging ellipsometry at optical wavelengths*, Biophysical Journal 92, 1306–1317

Chapter 5

Lipid domain discrimination by stimulated Raman scattering microscopy

The imaging equipment and configuration used for SRS were described in section 3.4; this chapter will describe the post-acquisition analysis and discuss the results obtained.

5.1. Data analysis

The SRS signal of a single lipid bilayer is very low; after all, it is the nonlinear response of a mere 4–5 nm of material. A typical regularised (section 3.3.4) SRS image (figure 5.1) has very little contrast and a low signal-to-noise ratio. The impedance of the detector used is approximately $10 \mu\Omega$, and the DC signal impedance is 36Ω .¹ The DC signal is of the order of 100 mV, but the SRS signal is only about 0.5 mV (see, for example, figure 5.1), which means that the relative modulation of the transmission is about 1.8×10^{-6} .

This section describes the development of the data analysis procedure we used to extract information of the phases of lipid bilayers from such low-signal images.

5.1.1. Spectral factorisation

Factorisation into spectra and concentrations of chemical components² (FSC³) was performed on the denoised, gradient-corrected SRS images from the ternary mixture. This procedure consists of finding matrices C , S and E such that $D = CS + E$ and $|E|^2$, the sum of the squares of the elements of E , is minimised. Here, D is the data, C is a non-negative matrix containing the component concentration distributions, S is a non-negative matrix containing the spectral components and E is the residual (or error) matrix. The algorithm can be guided by one or more existing locked spectra, which means incorporating the locked spectra into S , or it can be unguided. In both cases, the number of spectra to be used is given to the algorithm, and the remaining spectral components are all treated as

a systematic error in the factorisation. Guiding the algorithm with a number of locked spectra equal to the number of spectra to be used simply gives the spatial distribution of the guide spectra and the residual.

Each hyperspectral image was subjected to unguided FSC³ with 5 components. Figure 5.2 shows the results of FSC³ of the images shown in figure 5.1. One component was assigned to lipid and one was assigned to PBS; the assignment was based on the spatial distribution of the components and on prior knowledge of which regions consist of lipid and which have no lipid, and this prior knowledge itself came from the quantitative DIC images. The lipid component is shown in red, while the PBS component is shown in blue. The images in figure 5.2 show the spatial distribution of these two components. The remaining three components were generally systematic noise; they invariably had spatial distributions which looked like either white noise or a gradient rather than reflecting the spatial distribution of the sample's components. One of them is shown in green in each of the images. The colours of the spectra on the graphs on the right correspond to those of the components shown in the corresponding images; the two spectral components whose spatial distributions are not shown in each image are also omitted from the corresponding graph. The full results with all five components are shown in appendix F.

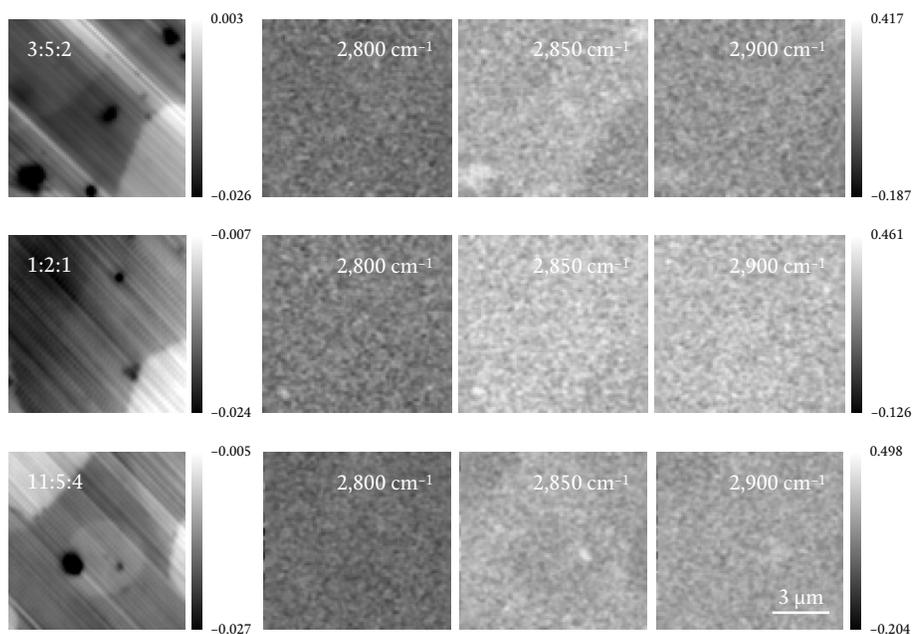


Figure 5.1: SRS images of three ternary samples: DOPC:sm:ch = 3:5:2 (top), DOPC:sm:ch = 1:2:1 (centre) and DOPC:sm:ch = 11:5:4 (bottom). The SRS data is hyperspectral and is shown at three different wave numbers: $2,800\text{ cm}^{-1}$ (centre-left), $2,850\text{ cm}^{-1}$ (centre-right) and $2,900\text{ cm}^{-1}$ (right). Quantitative DIC phase images of the same fields of view (left) are also shown. The scale is in radians for the quantitative DIC images and in mV for the SRS images.

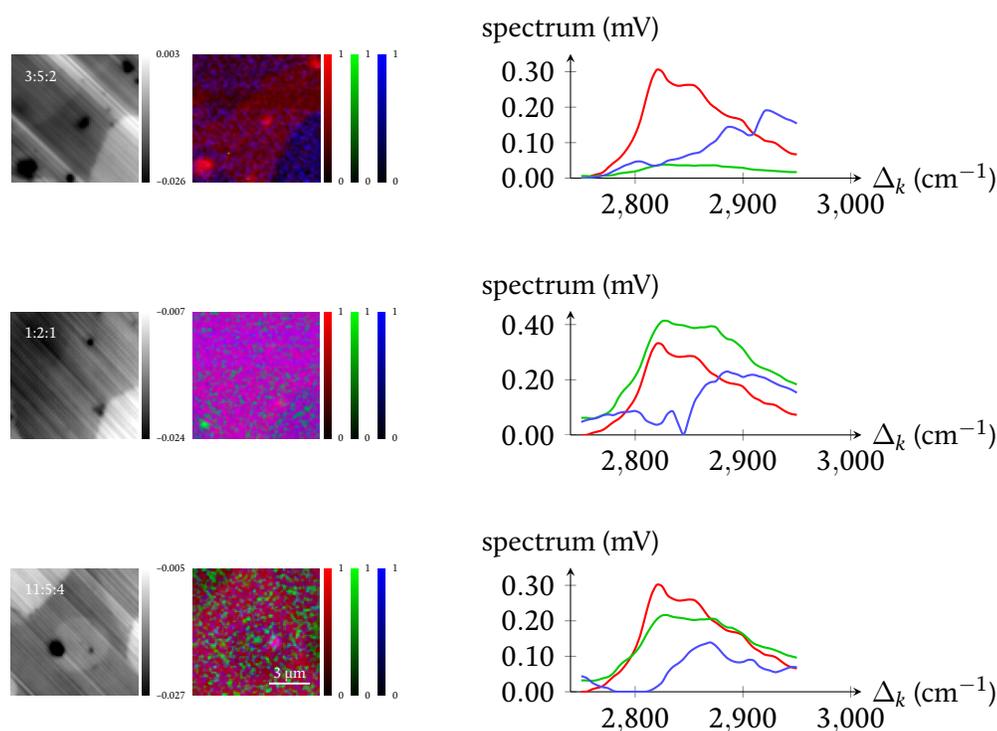


Figure 5.2: Quantitative DIC images (left), FSC³ images (centre) and FSC³ spectra of the regions shown in figure 5.1. The composition (DOPC:sm:ch) of the samples is 3:5:2 (top), 1:2:1 (centre) and 11:5:4 (bottom). The scale bars are in radians for the quantitative DIC images and in concentration for the colour components of the FSC³ images.

Unlike in the SRS images (figure 5.1), the structures in the sample can be seen fairly clearly in the FSC³ images, especially in those of the 3:5:2 and 11:5:4 samples. However, the different lipid domains (LO and LD) cannot be resolved; only one lipid component has been factorised in the case of each sample, meaning the two lipid components have SRS spectra too similar to discriminate by this method.

The vesicle seen inside the circular feature at the centre of the image of the 11:5:4 sample (seen as a large black spot in the quantitative DIC image and as a bright spot in the 2,850-cm⁻¹ and 2,900-cm⁻¹ SRS images) moved between the acquisition of the quantitative DIC and SRS images, which implies it was not attached to the lipid underneath.

Attempts to improve the spectral retrieval to the point where the LO and LD domains can be distinguished from each other are described in the remainder of section 5.1.

5.1.2. PBS signal subtraction

The thickness of the layer of lipid was negligible compared to the focal depth (about 1 μm), and the lower half of the focal volume was occupied by the coverslip glass. Therefore, the volume occupied by PBS was approximately half of the focal volume, some 100 times more

than the volume occupied by the lipid layer within the focal volume. The spectrum at each point of an image thus had a very significant contribution from the spectrum of PBS despite PBS having an $\text{Im}(\chi_3)$ of only 0.01–0.1 in the wave number range studied (compared to as much as 5 for lipid³), masking any differences in the spectra of the domains.

To overcome this, the spectrum of PBS was subtracted from each image prior to FSC³. Since this spectrum (which is very similar to that of water) is non-resonant at the wave numbers studied, it was deemed sufficient to subtract a constant from the spectrum at each pixel of the image. In order to obtain this offset for each pixel, the portion of the spectrum from 2,650 cm^{-1} to 2,670 cm^{-1} at the corresponding pixel was averaged.

5.1.3. Pump-Stokes overlap correction

The overlap of the pump and Stokes beams varied over the spectral scan range, resulting in the signal having a bell-shaped envelope, which shows the SRS spectra at two different points of an SRS image of a ternary sample.

The CARS intensity⁴ is proportional to $|\chi_3|^2 |E_p|^4 |E_s|^2$, where E_p is the amplitude of the pump beam and E_s is that of the Stokes beam. The SRS intensity,^{4,5} on the other hand, is proportional to $|\text{Im}(\chi_3)| |E_p|^2 |E_s|$. $|\chi_3|$ is dominated by the non-resonant response of the material in the sample, which is approximately constant. Therefore, the strength of the lipid signal relative to the background is much smaller for CARS images than for SRS images. The CARS images were thus much more uniform; while structure could occasionally be seen very faintly, the images were dominated by the PBS.

Since PBS, like water, is mostly non-resonant at wave numbers smaller than 3,000 cm^{-1} ,⁶ the shape of the CARS and SRS spectra in the studied region corresponded almost exactly to the overlap between the pump and Stokes pulses, with the lipid having

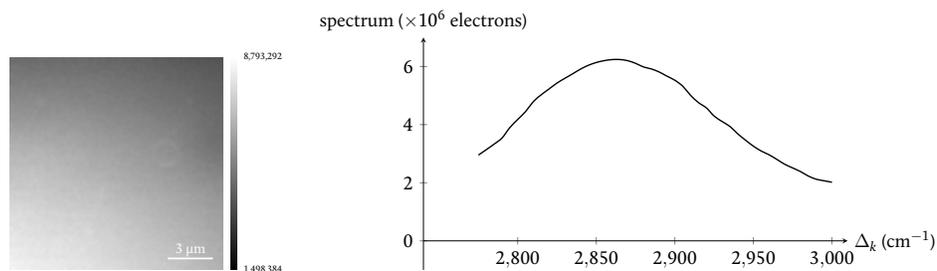


Figure 5.3: Spatially averaged CARS spectrum of an image from one of the ternary samples (DOPC:sm:ch = 11:5:4). This spectrum is enveloped by the pump-Stokes overlap but is otherwise approximately uniform. The image to the left shows the CARS signal from the same image at 2,850 cm^{-1} ; the greyscale range is in electrons.

only a very small contribution, especially in the CARS images for the reason explained above. This fact was used to compensate for the spectral unevenness of the pulse overlap. The CARS spectrum of each image was averaged over the entire image (figure 5.3), and the SRS spectra of each image were divided by the square root of the average CARS spectrum of the image (the square root was used because, as mentioned previously, the CARS intensity is proportional to $|\chi_3|^2$ and the SRS intensity is proportional to $|\text{Im}(\chi_3)|$).

With this correction and the PBS spectrum subtraction, the structures were slightly more visible in the SRS images even without spectral factorisation (figure 5.4). The noise is also visibly reduced, since the noise caused by laser fluctuations (which is present in both the CARS signal and the SRS signal) is divided by itself. Additionally, while the structure remains faint, the different domains are now visible as different signal intensities at $2,850\text{ cm}^{-1}$, although the LD domains are difficult to distinguish from PBS.

The improvement is also visible after spectral factorisation (figure 5.5). While the factorisation can still not distinguish the lipid domains from each other very well (the green

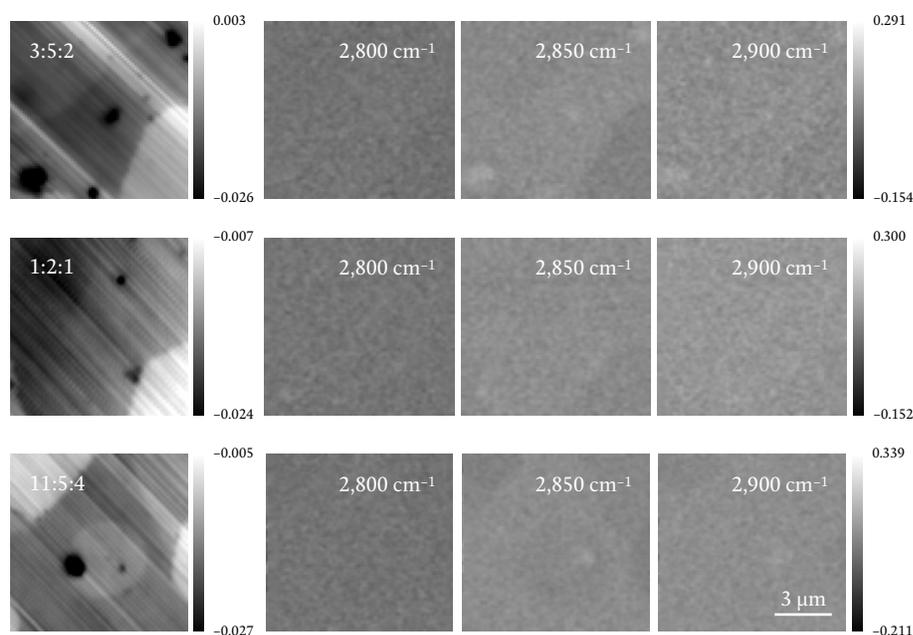


Figure 5.4: SRS images of the samples shown in figure 5.1 after subtraction of the PBS signal using the $(2,650\text{--}2,670)\text{-cm}^{-1}$ range and after pump-Stokes overlap correction. The images are shown in the same order and at the same wave numbers as in figure 5.1; the composition (DOPC:sm:ch) of the samples is 3:5:2 (top), 1:2:1 (centre) and 11:5:4 (bottom). The SRS data is shown at the same wave numbers as before: $2,800\text{ cm}^{-1}$ (centre-left), $2,850\text{ cm}^{-1}$ (centre-right) and $2,900\text{ cm}^{-1}$ (right). Also as before, quantitative DIC phase images of the same fields of view (left) are shown. The scale is in radians for the quantitative DIC images and in $\text{nV electrons}^{1/2}$ for the SRS images, since the CARS signal was detected with a photomultiplier tube and the CARS spectrum thus has units of electrons.

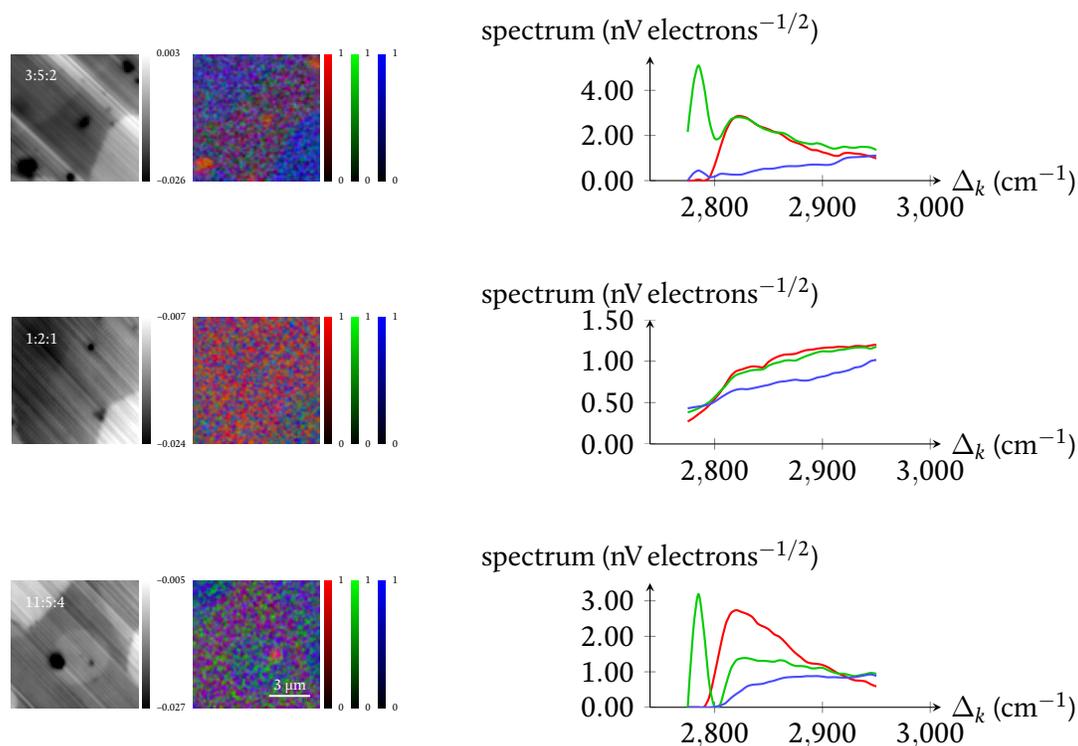


Figure 5.5: Quantitative DIC images (left), FSC³ images (centre) and FSC³ spectra of the regions shown in figure 5.4. The composition (DOPC:sm:ch) of the samples is 3:5:2 (top), 1:2:1 (centre) and 11:5:4 (bottom). The FSC³ was performed after the PBS signal was subtracted from the SRS spectra and the uneven pump-Stokes overlap was corrected by dividing the SRS spectrum of each image by the square root of the average CARS spectrum of the same image. The scale bars are in radians for the quantitative DIC images and in concentration for the FSC³ images.

component is still noise, as are the two components which are not shown), the features are more visible than without the PBS and overlap correction (compare figures 5.2 and 5.5). Additionally, the spectrum of the PBS component is more reminiscent of the spectrum of water (or PBS) at the wave numbers shown in figure 5.5 than it is in figure 5.2. As before, the full results are shown in appendix F. The second field of view was only subjected to 3-component FSC³ because using more components resulted in some of the components being different spectral halves of the same component: they would have the same concentration map, the spectrum of one would be zero from about 2,850 cm⁻¹ to 2,950 cm⁻¹, and the spectrum of the other would be zero from 2,750 cm⁻¹ to about 2,850 cm⁻¹.

5.1.4. PBS spectrum subtraction

In order to remove the contribution of PBS more effectively from the SRS data and further improve the SRS spectra (by which is meant further highlight the differences in the domain

spectra by reducing other contributions to the SRS spectrum of the images), the spectrum of PBS, rather than a constant number, was subtracted from the SRS spectra.

This SRS spectrum (figure 5.6) was obtained from an image of a pure DOPC sample. The spectrum at each point of the subregion of the DOPC image which contained no lipid (which once again was determined with the aid of the quantitative DIC image of the same field of view of the same sample) was averaged, and this average spectrum was used as the PBS spectrum. This, instead of a constant number, was subtracted from the SRS spectra of the images of the ternary samples.

The remainder of the procedure was the same as before, although there is one exception: the pump-Stokes overlap's unevenness was corrected after spectral factorisation rather than before. This was done because this unevenness does not affect factorisation, merely the shape of the resulting spectra, and because the factorisation assumes the data has white noise and, as seen earlier, this correction affects the noise.

Figure 5.7 shows the SRS images after PBS spectrum subtraction. The contrast is improved with respect to figure 5.1, though the noise remains; dividing by the square root of the average CARS spectrum removes the noise which comes from fluctuations in the laser intensity, so the noise in figure 5.7 is the same as that in figure 5.1. However, the noise is not strong enough to obscure the features completely.

It is now possible to distinguish the LD domain from PBS at $2,850\text{ cm}^{-1}$. With the exception of vesicles (which appear as very bright spots along the diagonal from bottom-left to top-right in the top image, near the bottom-left corner in the central image and

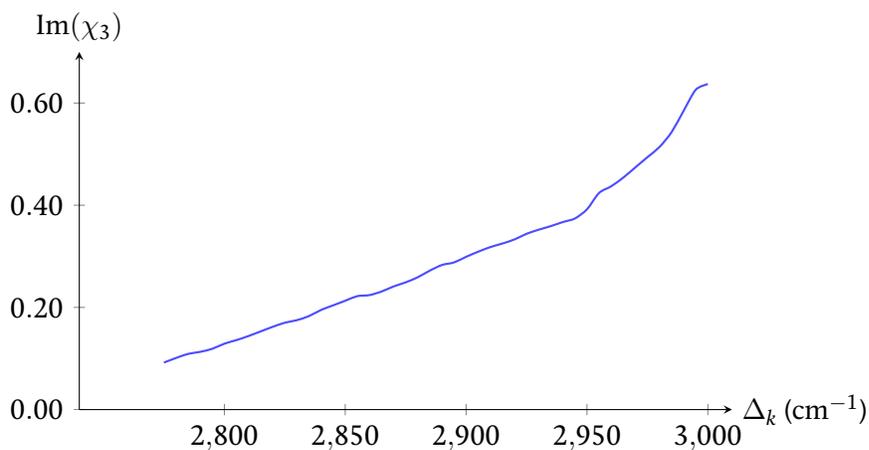


Figure 5.6: PBS spectrum from a pure-DOPC sample after division by the square root of the spatially averaged CARS spectrum of the same image from which the PBS spectrum was taken. The spectrum was obtained by spatially averaging the spectrum of the regions of the image which contained no lipid. Note that the PBS spectrum was subtracted from the images of the ternary samples prior to division by the square root of the CARS spectrum.

inside the circular LD domain at the centre of the bottom image), the brightest regions are LO domains. Meanwhile, the intermediate regions (for example, the top-left third of the top image and the top-right corner and circular feature in the bottom image) are LD domains and the darkest regions (the bottom-right third of the top image and the small regions at the top-left and centre-right of the bottom image)) are PBS.

Figure 5.8 shows the results of FSC³ of the SRS images after subtraction of the PBS spectrum. It is immediately evident that the image of the 3:5:2 sample has a horizontal feature near the top which is the result of a fluctuation of the laser intensity. However, the factorisation is better in that it now retrieves two lipid components (with the exception of the image of the 1:2:1 sample), even if they are not properly separated. These two components, shown in red and green in the 3:5:2 and 11:5:4 images, are spatially separated from the PBS component, shown in blue in all images. While the images are grainy, the lipid is now properly separated from the PBS and there are sometimes two lipid components. Once again, the full results are shown in appendix F. For the same reason as before, the

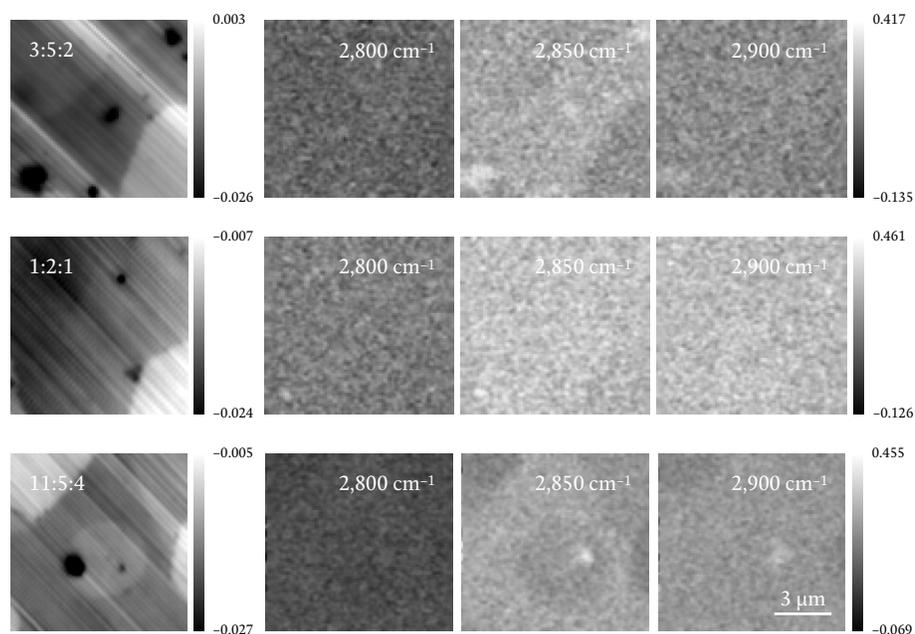


Figure 5.7: SRS images of the samples shown in figure 5.1 after subtraction of the PBS spectrum. The images are shown in the same order and at the same wave numbers as in previous figures; the composition (DOPC:sm:ch) of the samples is 3:5:2 (top), 1:2:1 (centre) and 11:5:4 (bottom). The SRS data is shown at the same wave numbers as before: 2,800 cm⁻¹ (centre-left), 2,850 cm⁻¹ (centre-right) and 2,900 cm⁻¹ (right). Also as before, quantitative DIC phase images of the same fields of view (left) are shown. The scale is in radians for the quantitative DIC images and in mV for the SRS images. The subtracted PBS spectrum was obtained by averaging the spectrum of a pure DOPC sample over the region which contained no lipid.

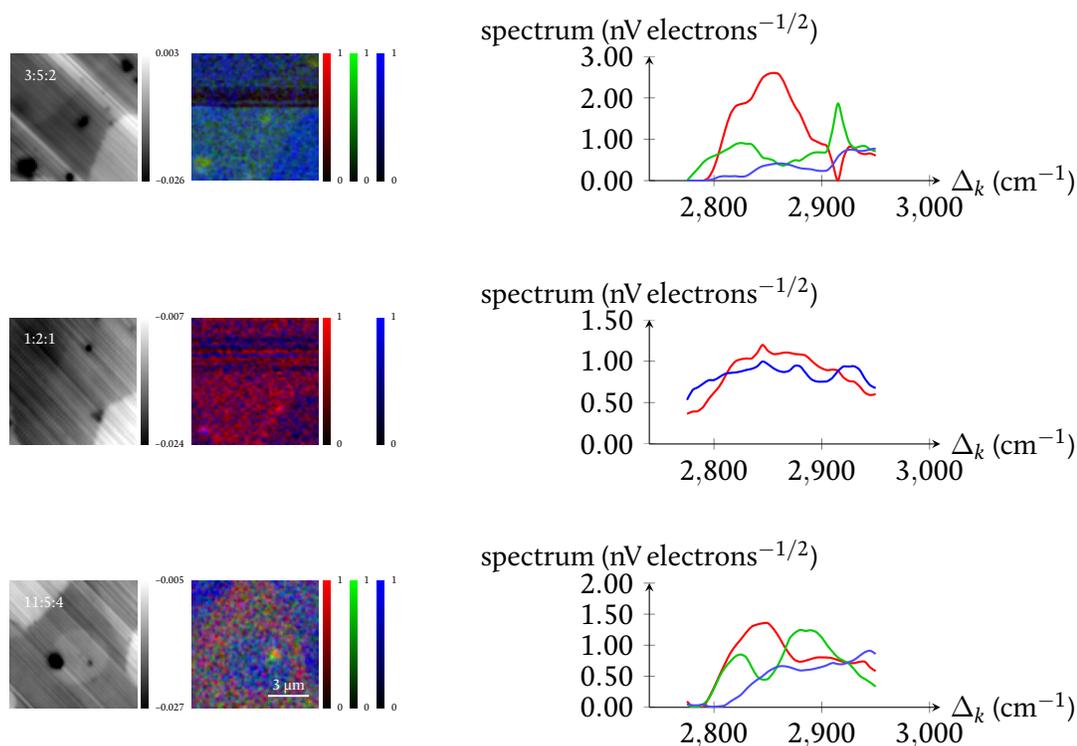


Figure 5.8: Quantitative DIC images (left), FSC³ images (centre) and FSC³ spectra of the regions shown in figure 5.7. The composition (DOPC:sm:ch) of the samples is 3:5:2 (top), 1:2:1 (centre) and 11:5:4 (bottom). The FSC³ was performed after the PBS spectrum (taken from a DOPC-only sample) was subtracted from the SRS spectra. The scale bars are in radians for the quantitative DIC images and in concentration for the colour components of the FSC³ images.

second field of view was only subjected to 4-component FSC³.

5.1.5. Local PBS subtraction

Due to noise and to variations in the exact amount of PBS in the focal volume across different images, the spectrum which needs to be subtracted in order to remove the contribution of the PBS is not necessarily the same for all images. This is evident in the fact that some PBS remains in the images.

Rather than attempt to scale the PBS spectrum individually for each image, which is a process which is not straightforward and prone to error, the local PBS spectrum was obtained in the same way as described above: the SRS spectrum was averaged over the PBS-only regions, but this was done for each image; the average PBS spectrum of each image was subtracted from that image.

It is obvious that this can only be done if there is at least one PBS-only region in the image, and it soon became evident that it only yields good results if the PBS-only re-

gion(s) amount to a fairly large fraction of the total image. Thus, a new set of images of a new ternary (DOPC:sm:ch = 11:5:4) sample was acquired and care was taken to include enough PBS-only area in each image.

5.1.6. Spatial binning

In order to reduce the noise present in the images and obtain better FSC³ results, the images were subjected to a 4×4 spatial binning in both the x direction and the y direction. $N_x \times N_y$ binning consists of partitioning the image into rectangular regions N_x pixels long in the x direction and N_y pixels long in the y direction and then averaging the signal across all pixels in each of these rectangular regions, turning each rectangle into a single pixel. This results in a smaller, less-smooth image, but one with less noise.

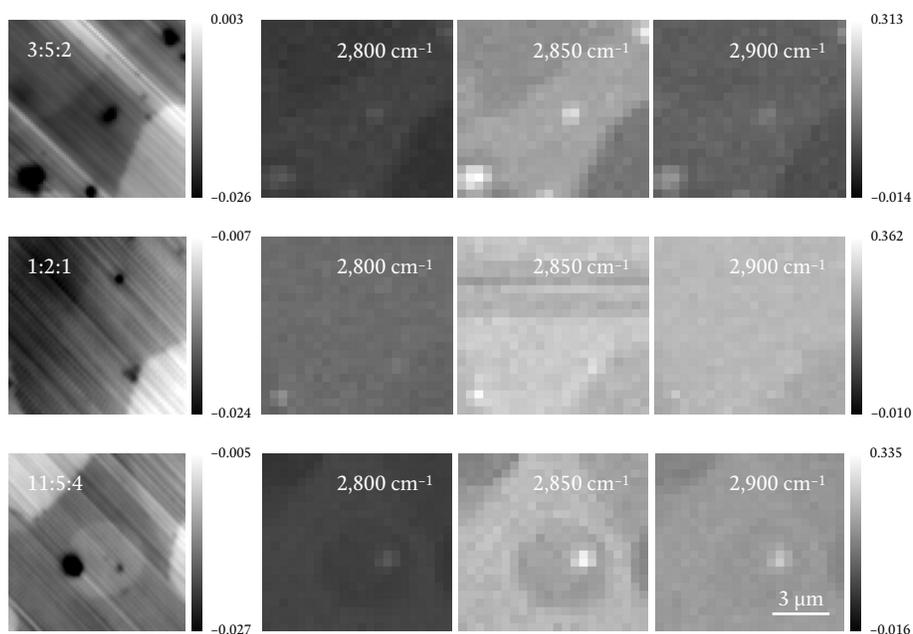


Figure 5.9: SRS images of the samples shown in figure 5.1 after subtraction of the local PBS spectrum and 4×4 spatial binning. The images are shown in the same order and at the same wave numbers as in previous figures; the composition (DOPC:sm:ch) of the samples is 3:5:2 (top), 1:2:1 (centre) and 11:5:4 (bottom). The SRS data is shown at the same wave numbers as before: $2,800 \text{ cm}^{-1}$ (centre-left), $2,850 \text{ cm}^{-1}$ (centre-right) and $2,900 \text{ cm}^{-1}$ (right). Also as before, quantitative DIC phase images of the same fields of view (left) are shown. The scale is in radians for the quantitative DIC images and in mV for the SRS images. The PBS spectrum subtracted from each image was obtained by averaging the spectrum of the PBS-only region(s) of the same image.

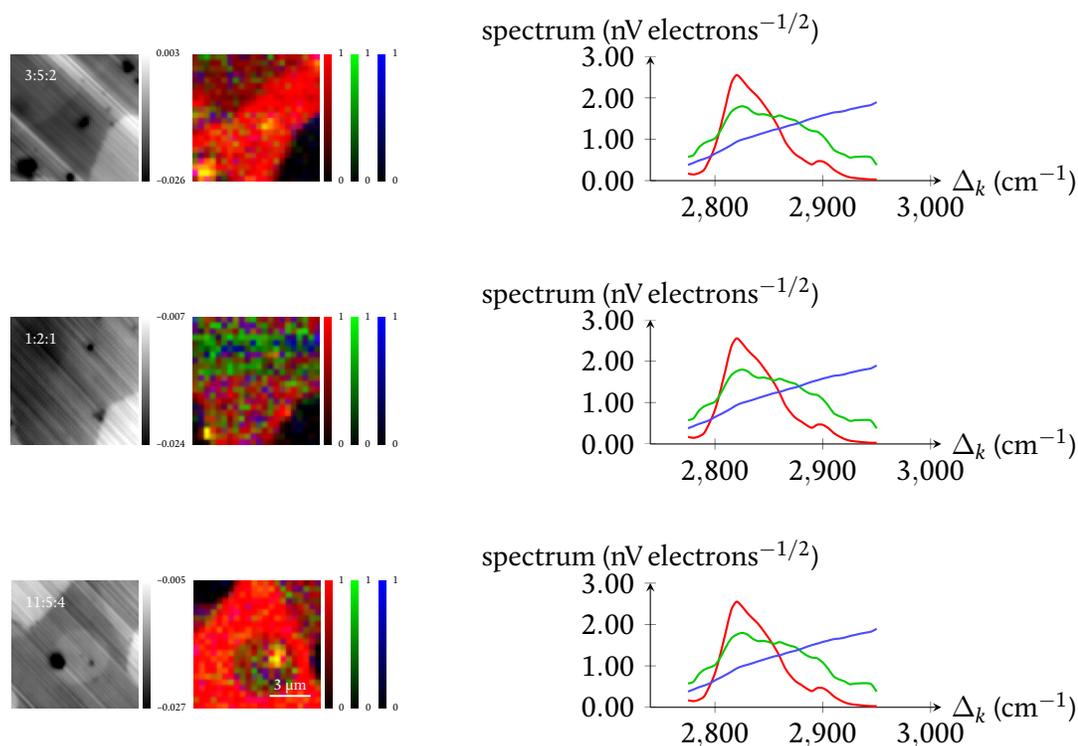


Figure 5.10: Quantitative DIC images (left), FSC³ images (centre) and FSC³ spectra of the regions shown in figure 5.9. The composition (DOPC:sm:ch) of the samples is 3:5:2 (top), 1:2:1 (centre) and 11:5:4 (bottom). The FSC³ was performed after the local PBS spectrum was subtracted from the SRS spectra and the images were binned to reduce noise. The scale bars are in radians for the quantitative DIC images and in concentration for the colour components of the FSC³ images.

5.1.7. Simultaneous spectral factorisation

Since the composition of the LO and LD domains is expected to be approximately equal for all images (see section 3.1.1), 3-component FSC³ was performed simultaneously on the images of the 3:5:2 and 1:2:1 samples and the first 11:5:4 sample. This means that the spectra retrieved by the factorisation algorithm are the same for all the images. The same was done for the images of the second 11:5:4 sample.

Figure 5.9 shows the SRS images of the 3:5:2 and 1:2:1 samples and the first 11:5:4 sample after local PBS spectrum subtraction and spatial binning. Even in the image of the 1:2:1 sample, the spatial features of the sample are now visible at all three wave numbers and especially at the higher two wave numbers.

Figure 5.10 shows the result of simultaneous FSC³ of all the images shown in figure 5.9. The spectra shown on the graphs have been divided by the square root of the average CARS spectrum of all the images. The red and green components correspond to lipid, while the blue component is the residual PBS. The images appear black in the PBS-only

regions, indicating that the vast majority of the PBS's contribution has been subtracted. The LO domains are clearly visible in red, but the LD domains are visible as dark-red regions with some green scattered across them, meaning the separation is imperfect; the red spectrum is the spectrum of the LO domains, but the green spectrum is the difference between the spectrum of the LD domains and the red spectrum, such that the spectrum of the LD domains is some linear combination of the red and green spectra.

Figure 5.11 shows the SRS images of the second 11:5:4 sample after local PBS spectrum subtraction and spatial binning. The quantitative DIC images clearly show the existence of multilamellar patches in addition to unilamellar LO and LD domains. The bottom field of view contains no LD domains. The SRS images of the top field of view are dominated by the extremely strong signal of the multilamellar patches, which obscures the boundaries between the unilamellar domains.

Figure 5.12 shows the result of simultaneous FSC³ of all the images shown in figure 5.11. While the red component has the same spatial distribution as the LO domains

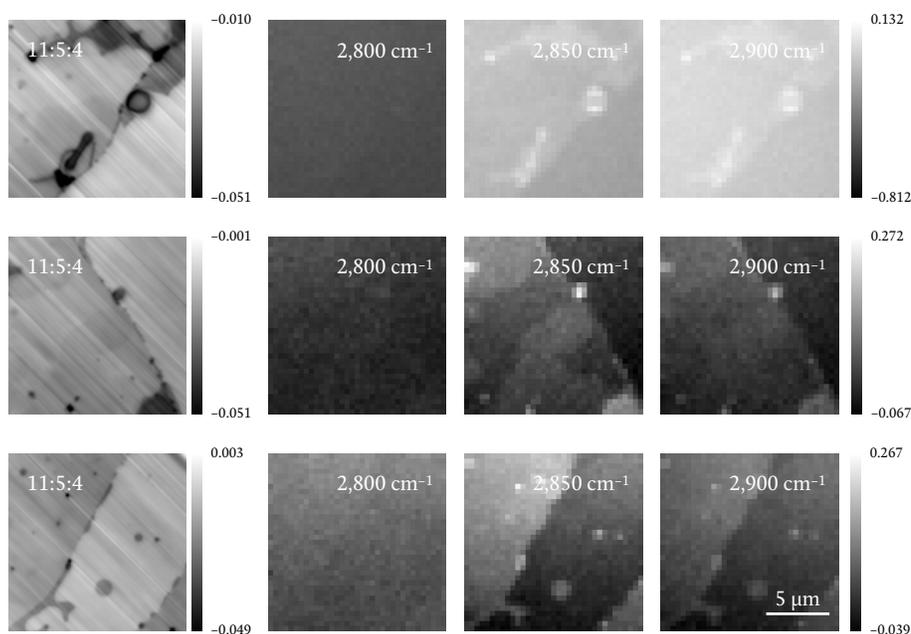


Figure 5.11: SRS images of a new ternary sample after subtraction of the local PBS spectrum and 4×4 spatial binning. The composition (DOPC:sm:ch) of the sample is 11:5:4. The images show different fields of view, all of which contain large PBS-only regions. The SRS data is shown at the same wave numbers as in figures 5.1, 5.4, 5.7 and 5.9: $2,800 \text{ cm}^{-1}$ (centre-left), $2,850 \text{ cm}^{-1}$ (centre-right) and $2,900 \text{ cm}^{-1}$ (right). Also as in those figures, quantitative DIC phase images of the same fields of view (left) are shown. The scale is in radians for the quantitative DIC images and in mV for the SRS images. The PBS spectrum subtracted from each image was obtained by averaging the spectrum of the PBS-only region(s) of the same image.

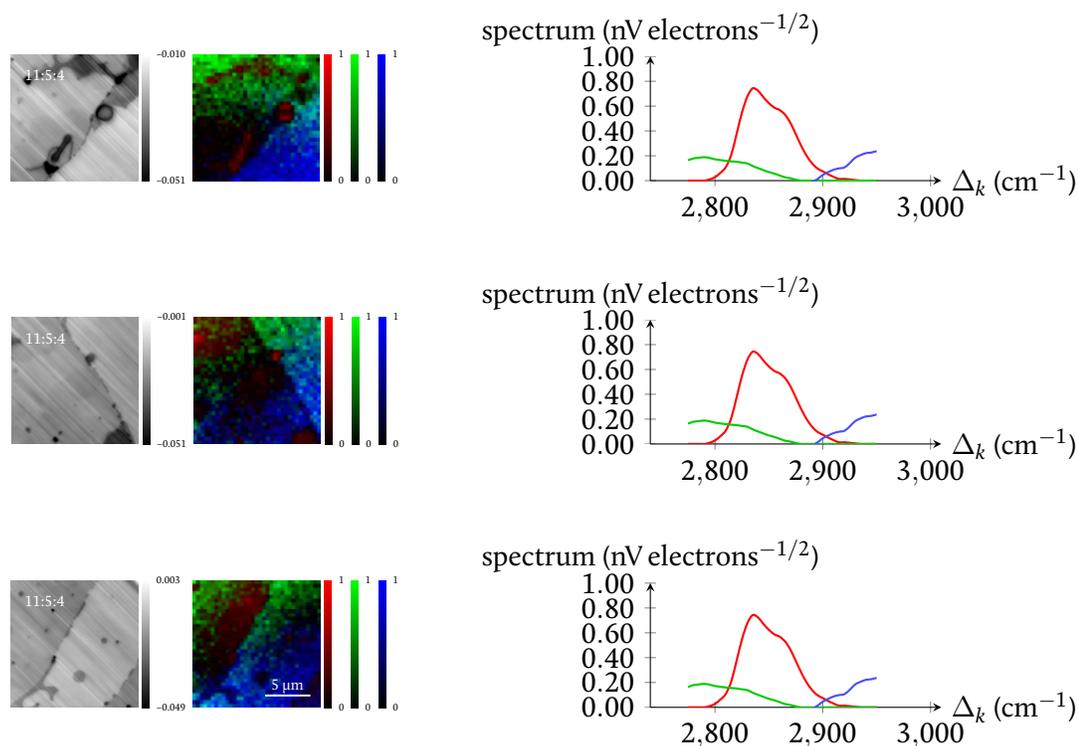


Figure 5.12: Quantitative DIC images (left), FSC³ images (centre) and FSC³ spectra of the regions shown in figure 5.11. The composition (DOPC:sm:ch) of the sample is 11:5:4; the top, central and bottom rows show different fields of view of the same sample. The FSC³ was performed after the local PBS spectrum was subtracted from the SRS spectra and the images were binned to reduce noise. The scale bars are in radians for the quantitative DIC images and in concentration Δ_k for the colour components of the FSC³ images.

seen in the quantitative DIC images, the green and blue components seem to be the result of the strong gradients along the y direction visible in the SRS images shown in figure 5.11, particularly the central and bottom images. This gradient is addressed in section 5.1.9.

5.1.8. Denoising

A way to reduce noise without binning is to perform a singular-value decomposition (SVD) of the spectral components and discarding all components which were not seen to contribute to the spatial structures in the image (figure 5.13).

Singular-value decomposition consists of finding matrices U , C and S such that the data D can be written as $D = UCS$. If D is an $N_s \times N_k$ matrix, where N_s is the number of spatial points and N_k is the number of spectral points, then U is a unitary $N_s \times N_s$ matrix which contains the spectral components, C is a non-negative diagonal $N_s \times N_s$ matrix which contains the component concentrations, and S is an $N_s \times N_k$ matrix which

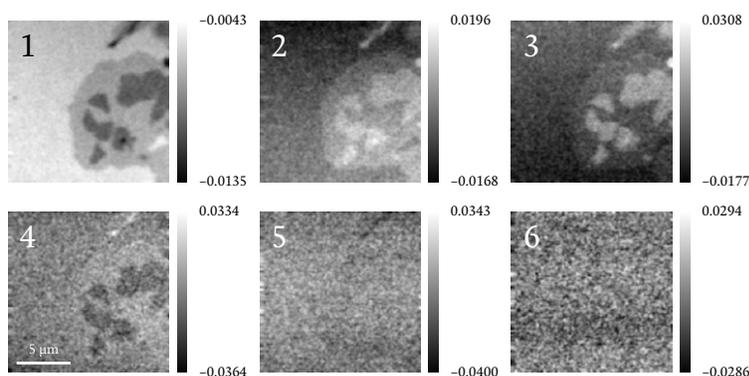


Figure 5.13: First six SVD components of one of the SRS images of the 2:2:1 sample. The first four components display varying degrees of spatial structure and were thus kept; all of the remaining components, as exemplified by the fifth and sixth ones in this figure, have no spatial structure and were thus discarded as noise. The greyscale ranges are in mV.

contains the spectral components' spatial distributions.

Figure 5.13 shows the first six SVD components of an image of a ternary sample (DOPC:sm:ch = 2:2:1). In this case, the first four SVD components show spatial structure, indicating that they contribute to the signal from the sample, and all remaining SVD components (including the fifth and sixth ones) do not, indicating that they contribute only to the noise. All components from the fifth one were discarded to reduce the noise.

5.1.9. Gradient correction by background fit

Thermal drift of the focal plane during image acquisition resulted in the amount of PBS in the focal volume changing as a function of time (and thus of the position on the image), which in turn resulted in a gradient along the slow scan axis, which we call the y axis (figure 5.14). This gradient was very pronounced in some of the images (see, for example, figure 5.11) and influenced the spectral factorisation.

In addition, small spatial variations in the overlap of the pump and Stokes beams often resulted in a correspondingly small gradient across the image.

The PBS signal and both gradients were simultaneously removed from the SRS images by obtaining a linear fit of the PBS-only regions and subtracting the fitted function from the entire image for each wave number; like denoising, this was done individually for each image, since each image had a different gradient. The function was independent for each wave number, so, in fact, one linear function per wave number was fitted and subtracted from the data at the corresponding wave number. This resulted in the PBS spectrum effectively being subtracted as well. This was done instead of local PBS spectrum subtraction.

Like the subtraction of the local PBS spectrum, this can best be done on images which have large PBS-only regions.

Figure 5.15 shows the SRS images of the second 11:5:4 sample after the data was de-

noised via SVD and the gradients were removed by fitting and subtracting a linear function of y to the PBS-only regions. A weak gradient remains in the central image, probably due to the PBS-only region being too small for the fit to be very accurate, but the features of the samples are more visible than in figure 5.11. Nevertheless, the images remain dominated by the multilamellar features, so the contrast remains too low for the domains in the top image to be visible.

The images were subjected to 2-component FSC³ independently of each other. Only two components were used because the samples were not expected to have more than this amount (one for each lipid phase; the PBS had been removed by the subtraction of the linear function).

Figure 5.16 shows the result of FSC³ of the images from figure 5.15.

The top image remains dominated by the multilamellar features, so the factorisation resulted in these features being separated from the rest of the sample; therefore, the spectra are not representative of the LO and LD spectra and the two domains could not be separated in this case.

The central image, as mentioned previously, still had a gradient, which is represented here by the green component; the gradient was strong enough for the factorisation to favour isolation of it rather than of the LD domains, so, while the red component corresponds to the LO domains, the green component does not correspond to the LD domains.

The field of view seen in the bottom image, as mentioned previously, does not contain any LD domains. Therefore, the factorisation resulted in two components with different spectra but the same spatial distribution. The weaker of these components is shown in grey, while the stronger one is shown in red. It is believed that the factorisation algorithm

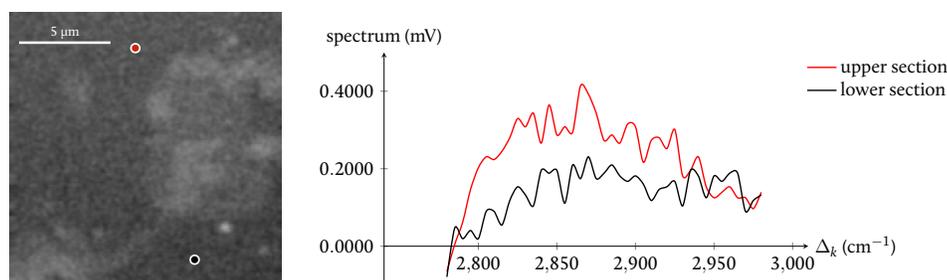


Figure 5.14: An SRS image from one of the ternary samples (DOPC:sm:ch = 2:2:1). The image has a gradient along both spatial axes. The graph on the right shows the spectra at one point in the upper section of the image and one point in the lower section. The red and black circles on the image indicate where the spectra were taken. The spectrum at the point in the upper section is considerably stronger than that at the point in the lower section, revealing the presence of the gradient.

split the spectrum of the same component into two complementary spectra, since two factorisation components were used, so the actual spectrum of the regions shown in red in the FSC³ image is the sum of the two spectra shown; this has been seen previously, and more dramatically, with images with several different chemical components which are subjected to FSC³ with more spectra than there are components in the sample.

5.1.10. Remarks

It can be concluded, then, that this procedure (SVD denoising followed by fitting and subtraction of the PBS and gradients) increases the signal-to-noise ratio to the point where it is realistic to expect FSC³ to be able to separate the LO and LD domains from each other.

However, it must be cautioned that this requires the absence (or masking to remove from the analysis) of multilamellar or other high-contrast features in the images (so they do not obscure the more subtle differences between the spectra of the two phases) and the presence of a large PBS-only region (in order for the fitting to be accurate).

In the following section, the results of the analysis of new images from several samples —pure-LO (DOPC:sm:ch = 0:7:3), pure-DOPC (DOPC:sm:ch = 1:0:0) and ternary

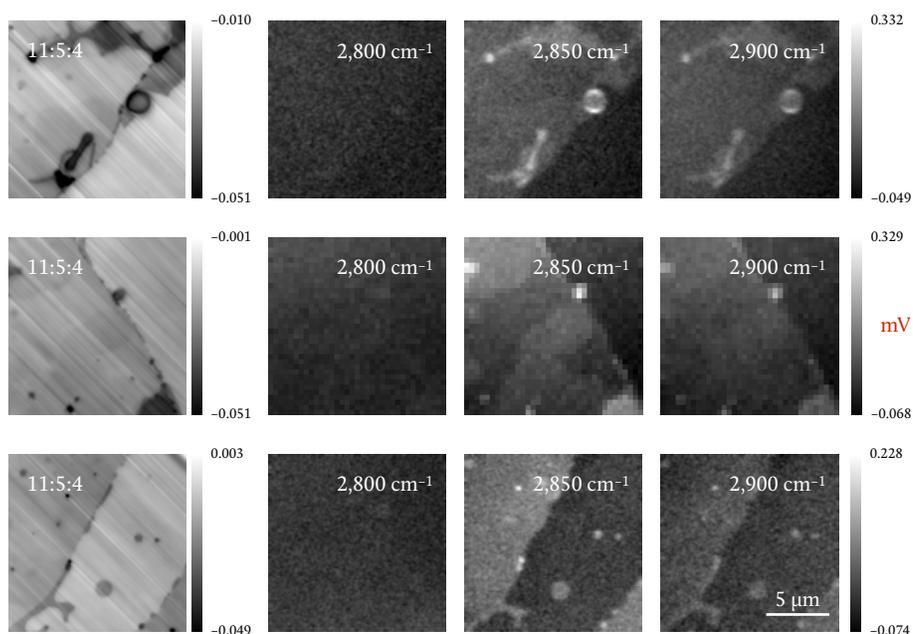


Figure 5.15: SRS images of the second 11:5:4 sample after SVD denoising and after a linear function of the y coordinate of the position was fitted to the PBS-only regions and subtracted from the images. The images are shown in the same order and at the same wave numbers as in figure 5.11; the composition (DOPC:sm:ch) of the sample is 11:5:4. As before, quantitative DIC phase images of the same fields of view (left) are shown. The scale is in radians for the quantitative DIC images and in mV for the SRS images.

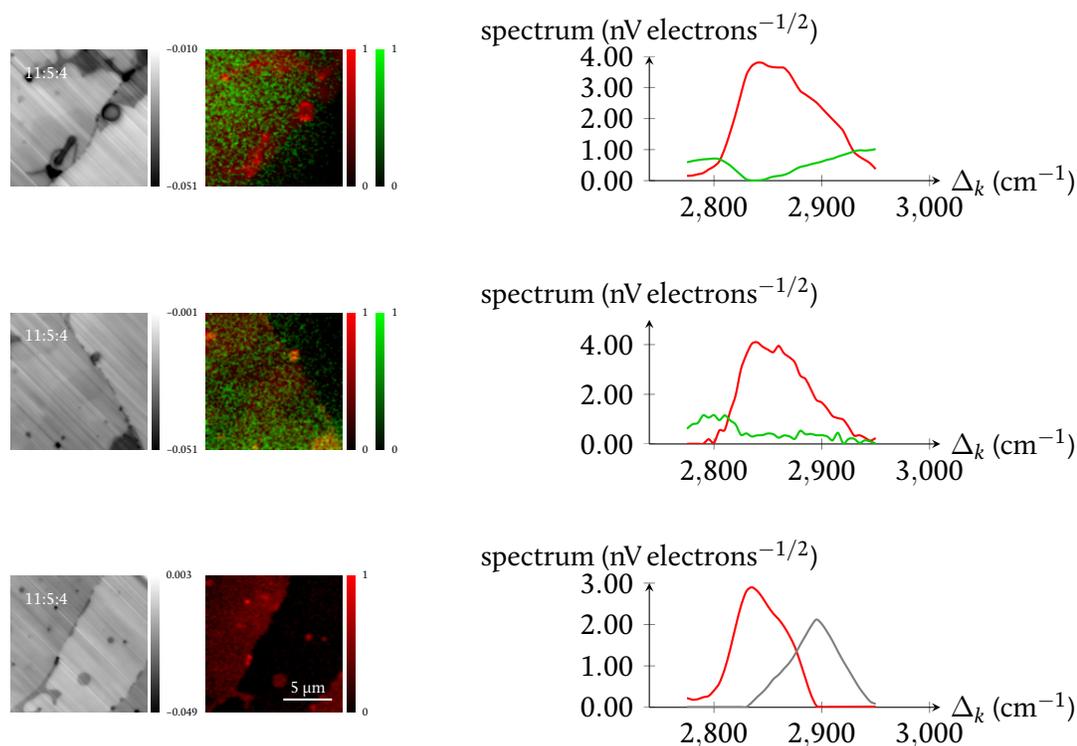


Figure 5.16: Quantitative DIC images (left), FSC³ images (centre) and FSC³ spectra of the regions shown in figure 5.15. The composition (DOPC:sm:ch) of the sample is 11:5:4. The top, central and bottom rows show different fields of view of the sample. The FSC³ was performed after the images were denoised via SVD and after a linear function in y was fitted to the PBS-only regions and subtracted from the images. The scale bars are in radians for the quantitative DIC images and in concentration for the colour components of the FSC³ images.

(DOPC:sm:ch = 2:2:1)— will be shown. These images were taken in order to allow the requirements established in the preceding paragraph to be satisfied and increase the chances of FSC³ being able to separate the two components.

5.2. Analysis of the 2:2:1 sample

5.2.1. Pre-factorisation analysis

SRS images of unilamellar and multilamellar sphingomyelin+cholesterol (sm+ch) and DOPC supported bilayers and unilamellar ternary (DOPC:sm:ch = 2:2:1) supported bilayers were taken. As with the samples discussed in section 5.1, CARS images were taken simultaneously and quantitative DIC images of the same regions of the same samples were taken immediately before the SRS and CARS images.

All images acquired were denoised by SVD. Gradients from thermal drift and from

spatial variations of the pump-Stokes overlap, as well as the contribution of PBS, were removed from the SRS images of the unilamellar samples by linear fitting and subtraction of the PBS-only regions; the fitted function was sometimes a function of y only and sometimes a function of both x and y (see appendix G). Noise and gradients also existed in the images from the multilamellar samples, but they were not significant because the large amount of lipid in the focal volume gave a signal strong enough to continue with the analysis as described below without background correction. Additionally, the PBS-only regions in these images were too few and too small for the fitting to be accurate.

5.2.2. Spectrum retrieval

In the case of the sm+ch and DOPC samples, the lipid SRS spectrum was obtained by simply averaging the spectrum of the regions containing a lipid bilayer in the denoised, gradient-corrected SRS images, avoiding regions with multiple bilayers (in the case of the unilamellar samples) or changes in the number of bilayers (in the case of the multilamellar samples), as well as vesicles and debris. This was possible because these samples exhibited a single lipid phase. For the multilamellar samples, the region with the largest number of bilayers in each image was selected in order to maximise the SRS signal (and thus the signal-to-noise ratio).

Each image from the ternary sample was subjected to unguided FSC³ with 2 and 3 components. For each image, the result showing the clearest segregation of lipid domains was selected and the other result was discarded; the selected result was usually the one of the 2-component FSC³ (see appendix G).

5.2.3. Post-factorisation corrections to the spectra

The spectra obtained as described above were scaled using the value of $\text{Im}(\chi_3)$ for a single lipid bilayer, obtained previously using CARS imaging.³ This allowed the determination of $\text{Im}(\chi_3)$ as a function of wave number for each lipid component.

The subtraction of the linear function introduces a small error in the retrieved spectra. If the thickness of a lipid bilayer is d , then the thickness of the PBS layer above the lipid is thinner by an amount d than the thickness of the PBS layer in the PBS-only regions. Since the SRS spectrum at the PBS-only regions was used to calculate the linear function subtracted from the entire image to correct for the gradient, the spectrum at the lipid regions is not truly the spectrum of the lipid, but the spectrum of the lipid minus the spectrum of a layer of PBS of thickness d . Therefore, to retrieve the lipid spectra, this had to be compensated. The spontaneous-Raman spectrum of a layer of PBS with thickness d (equal to 4.96 nm for the LO domains and 4.08 nm for the LD domains⁷) was added to the FSC³ SRS spectra in order to compensate for the initial background subtraction. Spontaneous Raman microscopy was used to obtain the spectrum of PBS because it was less noisy than the SRS spectrum. However, this meant the scale was different, so the Raman spectrum of PBS was scaled to the corresponding SRS signal amplitude. This was then rescaled to the $\text{Im}(\chi_3)$ of a layer of PBS of thickness d by multiplying it by the derivative of the SRS

signal of PBS with respect to the distance z between the sample and the focal plane, which itself was normalised to the SRS signal of PBS at $z = 0$, all at the central wave number of the SRS scan (figure 5.17).

The resulting SRS spectrum, then, is given by

$$\begin{aligned} \text{Im}\left(\chi_3^{(\text{lipid})}\right) = & \text{Im}\left(\chi_3^{(\text{PBS})}\right)\Big|_{\Delta_k^{(\text{centre})}} \frac{S_{\text{FSC}^3}}{2 S_{\text{FSC}^3}\Big|_{\Delta_k^{(\text{centre})}}} \sqrt{\frac{\langle S_{\text{CARS}} \rangle}{\langle S_{\text{CARS}} \rangle\Big|_{\Delta_k^{(\text{centre})}}}} \\ & + \frac{S_{\text{Raman}}^{(\text{PBS})}}{S_{\text{Raman}}^{(\text{PBS})}\Big|_{\Delta_k^{(\text{centre})}}} \frac{S_{\text{FSC}^3}^{(\text{PBS})}\Big|_{\Delta_k^{(\text{centre})}}}{2 S_{\text{FSC}^3}^{(\text{PBS})}\Big|_{\Delta_k^{(\text{centre}), z=0}}} \left(\frac{\partial S_{\text{FSC}^3}^{(\text{PBS})}\Big|_{\Delta_k^{(\text{centre})}}}{\partial z} \right)\Big|_{z=0} d, \end{aligned}$$

where S_{FSC^3} is the SRS FSC^3 signal prior to scaling, S_{CARS} is the CARS signal, S_{Raman} is the spontaneous Raman signal and $\Delta_k^{(\text{centre})}$ is the central wave number of the SRS scan.

5.2.4. Results from LO and DOPC samples

Figures 5.18 and 5.19 show the unilamellar and multilamellar spectra of the sm+ch and DOPC samples, respectively. In the case of the unilamellar samples, the average spectra shown (continuous curves) are the averages of the individual spatially averaged spectra of the images of each sample.

In the case of the multilamellar sm+ch sample, since each image was taken at a different field of view and the number of bilayers was different in each of them, the spectra were scaled down to single-bilayer strength before being averaged. In order to do this, the peak of the SRS spectrum of each image of the multilamellar sample was divided by the average of the peaks of the spatially averaged SRS spectra of the unilamellar images

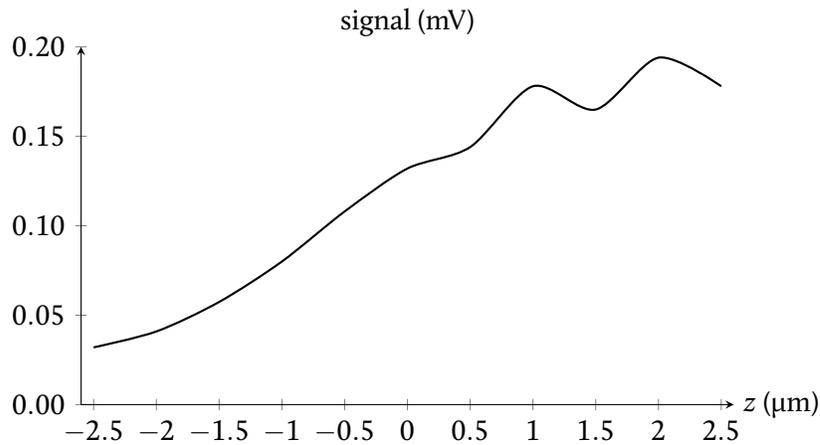


Figure 5.17: z dependence of the strength of the SRS signal in a PBS-only region.

(this average of the unilamellar spectra is shown in red on the top graph of figure 5.18; the individual spectra are shown as black dashed curves) and the resulting number, rounded to the nearest integer, was taken to be the number of bilayers.

In the case of the multilamellar DOPC sample, the spatially averaged spectrum of the most uniform image was kept and the other image was discarded, since the image's lack of uniformity greatly affected the averaged spectrum. This lack of uniformity came from the DOPC being locally heated by the lasers to the extent that bubbles appeared in the lipid, altering the signal significantly. These bubbles, which can be seen in figure 5.19, were avoided in the spatial averaging of the spectrum. The spectrum was scaled down to single-bilayer strength in the same way as the spectra of the multilamellar sm+ch sample.

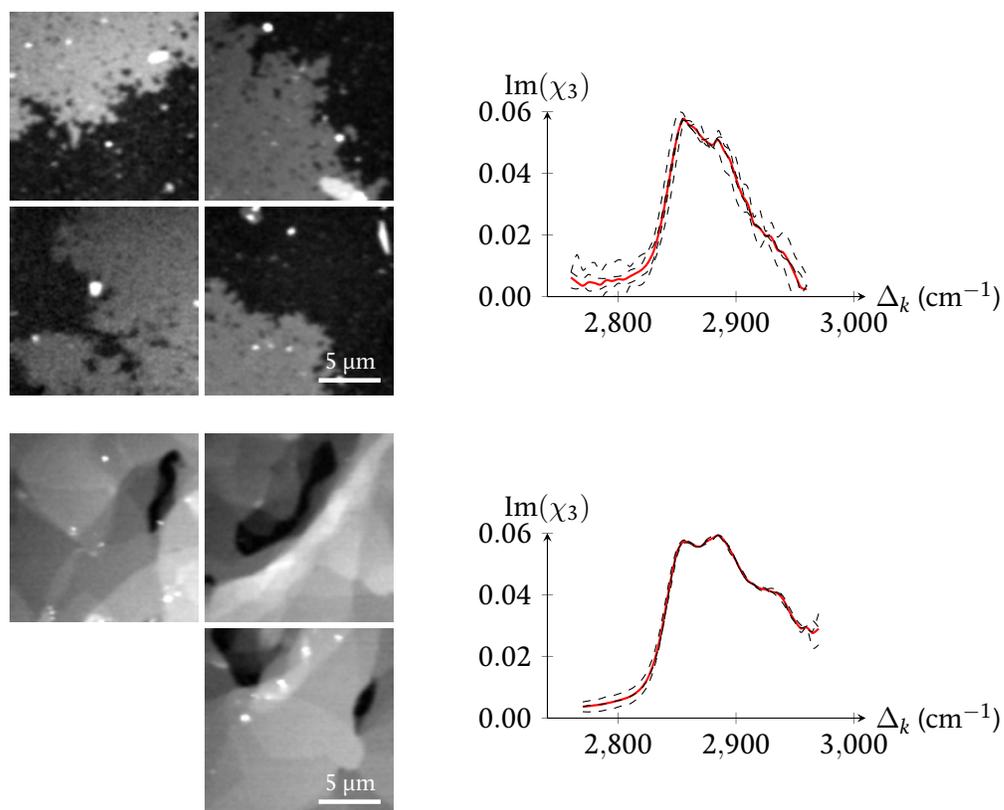


Figure 5.18: SRS images and spectra of the unilamellar (top) and multilamellar (bottom) sm+ch samples. Dark regions in the images are PBS-only regions, and light regions are lipid regions. Black dashed curves are the individual spectra, and continuous red curves are the average spectra. The multilamellar spectra have been scaled to single-bilayer strength. The grey ranges are, from top to bottom and left to right and in millivolts, $(-0.070, 1.290)$, $(-0.170, 3.720)$, $(-0.106, 0.777)$, $(-0.185, 1.190)$, $(-0.299, 7.346)$, $(-0.213, 3.900)$ and $(-4.863, 11.625)$. All images are shown at the centre wave number.

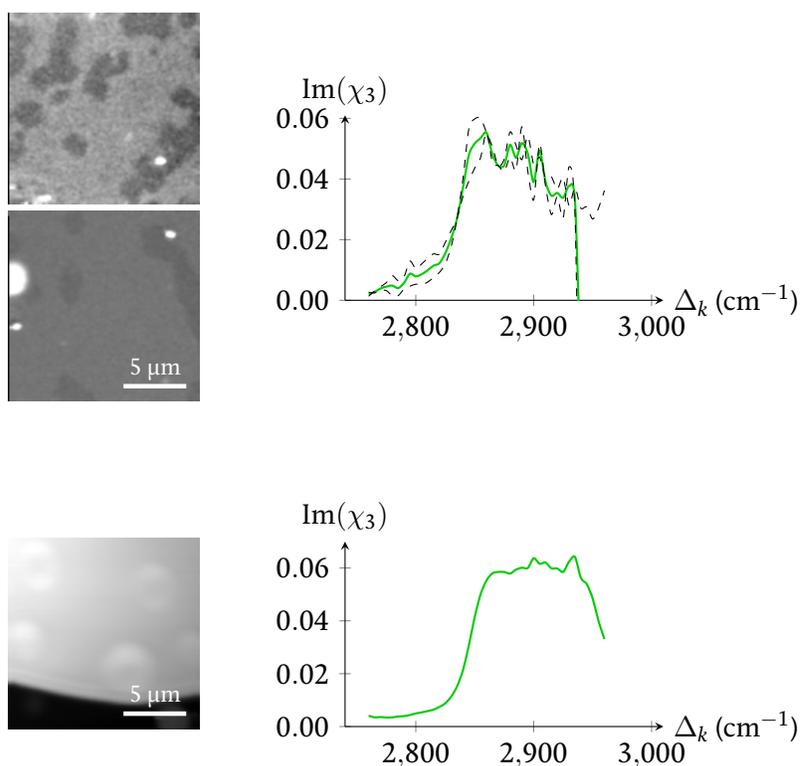


Figure 5.19: SRS image(s) and spectra of the unilamellar (top) and multilamellar (bottom) DOPC samples. As in figure 5.18, dark regions in the images are PBS-only regions and light regions are lipid regions. Black dashed curves are the individual spectra of the unilamellar sample images, and the continuous green curve in the top graph is the average spectrum; the graph on the bottom shows the spectrum of the only image which was considered for analysis, avoiding the lipid bubbles (visible as lighter circular regions), which formed due to local heating of the sample by the imaging lasers. The multilamellar spectrum has been scaled to single-bilayer strength. The grey ranges are, from top to bottom and in millivolts, $(-0.178, 4.322)$, $(-5.203, 10.470)$ and $(-1.908, 31.178)$. All images are shown at the centre wave number.

Although the multilamellar spectra have some contribution from the PBS (as can be seen by the gradual increase of the base signal towards higher wave numbers when compared with the unilamellar spectra), these provide a good measure of the $\text{Im}(\chi_3)$ of a sphingomyelin+cholesterol bilayer and a DOPC bilayer.

Concerning the sm+ch case, it is important to note that the composition (sm:ch) of the LO domains of a ternary mixture is not exactly 7:3. However, given that the relative concentration of sphingomyelin to cholesterol is only 26% higher in the 7:3 mixture than in the LO regions of a ternary mixture (see section 3.1.1), this remains a reasonably accurate measure of the $\text{Im}(\chi_3)$ of this type of lipid domain, both qualitatively and quantitatively.

5.2.5. Results from unguided analysis of the ternary sample

Of the nine images acquired from the ternary sample, only eight featured domains of both kinds; the remaining one had only a large patch of LD lipid. Of those eight, two had only

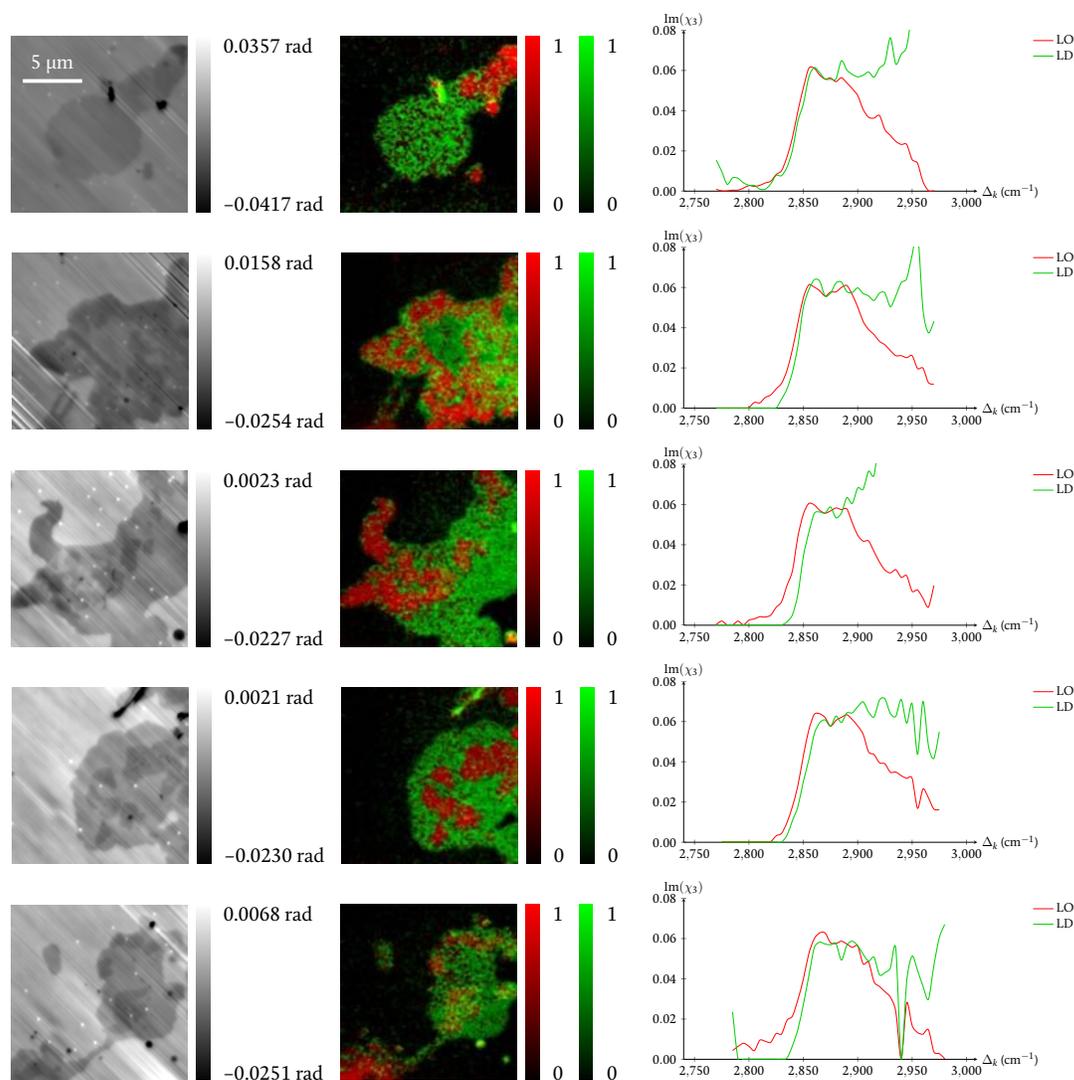


Figure 5.20: Quantitative DIC (left) and FSC³ (centre) images of different regions of the 2:2:1 sample. In the SRS images, the red regions are the LO domains and the green regions are the LD domains. The SRS spectra (right), obtained by unguided FSC³, follow the same colour code.

unilamellar LD domains coexisting with multilamellar regions which may have consisted of a combination of LO bilayers and LD bilayers; this made analysis difficult. Of the six remaining images, one had very strong gradients and not a large enough PBS-only region for the background fit to be accurate enough to remove the gradient; as with the second field of view shown in figure 5.16, the residual gradient caused the FSC³ algorithm to separate the image into components which did not correspond to the spatial distribution of the domains. Therefore, in the following discussion of the results of the ternary sample we will only consider the remaining five images.

Figure 5.20 shows the FSC³ images next to quantitative DIC images of the same fields of view. The figure also shows the spectra of the two domain types in each image, obtained by unguided FSC³. Since the LO domains were expected to be purely sphingomyelin and cholesterol and the LD domains were expected to be mostly DOPC (see section 3.1.1), the colours were chosen to match those used for the sm+ch and DOPC samples shown in figures 5.18 and 5.19.

The domain boundaries are visible in both the quantitative DIC images and the SRS images. Although there is some discrepancy in the exact shapes of the spectra from image to image, the general features of the LO and LD spectra are consistent between images; as in the case of figure 5.18, the LO spectrum has two peaks between 2,850 cm⁻¹ and 2,900 cm⁻¹, which correspond to CH₂ stretching,⁸ and then falls gradually as the wave number increases, while, as in the case of the multilamellar spectrum in figure 5.19, the LD spectrum features multiple peaks from 2,850 cm⁻¹ to around 2,950 cm⁻¹, corresponding to both CH₂ stretching and CH stretching,⁸ before falling abruptly; DOPC has twice as many double carbon bonds than sphingomyelin, so it contains twice as many CH bonds. The identification of which spectrum corresponds to the LO regions and which one corresponds to the LD regions was, in this case, based on a comparison with the sm+ch and DOPC spectra and is consistent with the determination of the phase of each domain by quantitative DIC.

Importantly, the domains are distinguishable without guidance even though the sample was unilamellar (which meant that the signal was rather low) and even though the spectra are fairly similar to each other (especially at the lower end of the wave number range shown). This means that, even for a mixture of unknown composition, as long as the spectra are different enough from each other in at least part of the spectral range and as long as enough of the image consists of PBS-only regions to perform the background correction with, the phases will be distinguishable with our technique (as long as the considerations discussed in section 5.1.10 are taken into account).

5.2.6. Results from self-guided analysis of the ternary sample

The LO and LD regions of the images from the ternary sample were determined with the aid of the quantitative DIC images, after which the LO and LD spectra of each image were obtained by averaging over the regions which were in the corresponding phase. These spectra were used as guides for guided FSC³ of the images. The guidance was local, meaning that the LO and LD spectra from each image were used as guides for that image

only.

Figure 5.21 shows the FSC³ and quantitative DIC images next to the spectra obtained by guided FSC³.

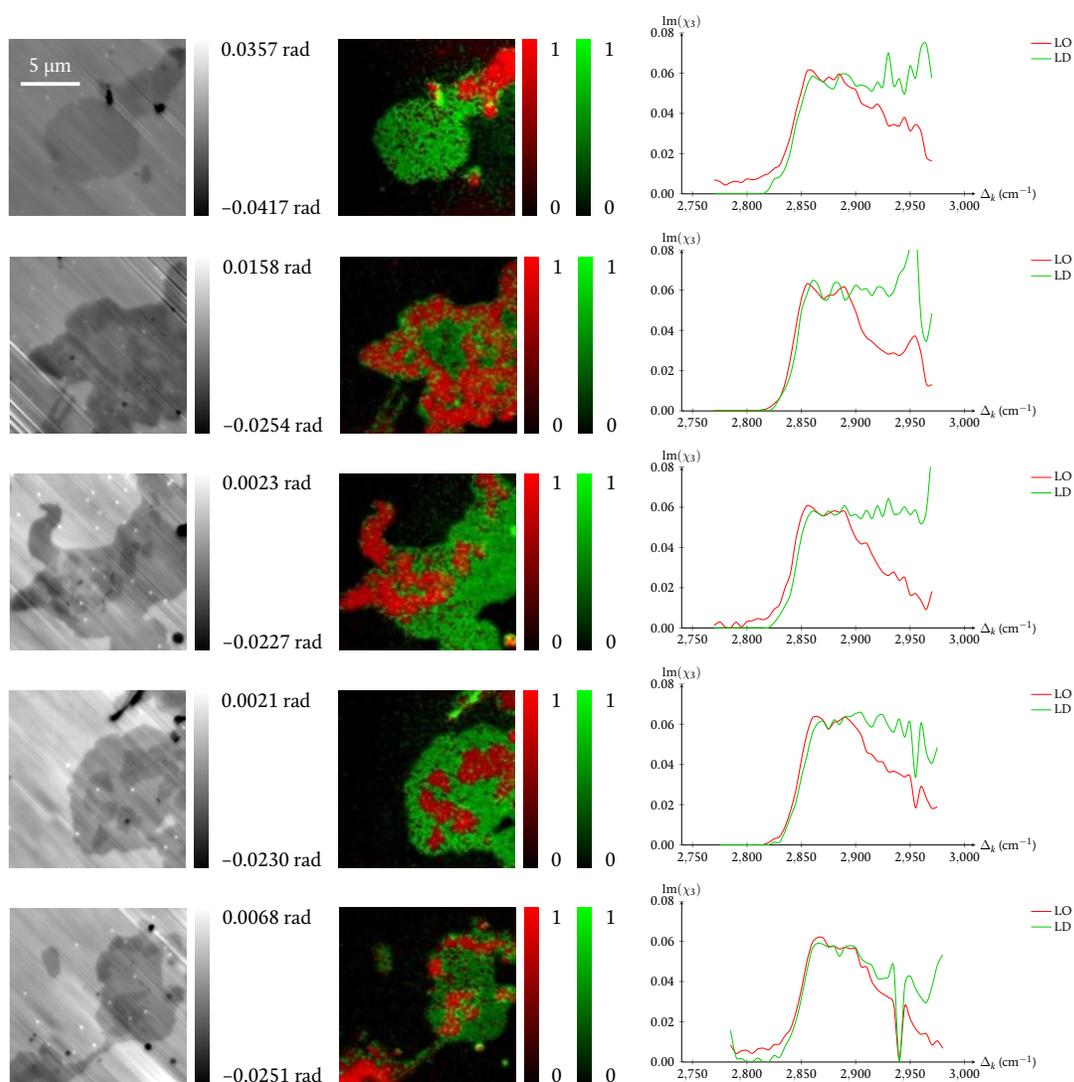


Figure 5.21: Quantitative DIC (left) and FSC³ (centre) images of different regions of the 2:2:1 sample. Colour coding is as in figure 5.20. In this case, the SRS spectra are from guided FSC³ and are likely more representative of the actual spectra of the domains. The guide spectra used for each image came from that same image.

The separation into phases in the FSC³ images is somewhat clearer than in the unguided case, especially for the second and fifth fields of view; the boundaries between the domains are sharper. Additionally, the domains themselves are visibly less noisy; this can be appreciated especially in the images of the first, third and fifth fields of view. These results were expected, as the algorithm no longer had to guess how to factorise the spectrum and partition the image.

The spectra, although still not quite identical to each other, are more consistent from image to image. In particular, the LD spectrum is flatter in the range between 2,850 cm⁻¹ and 2,950 cm⁻¹ except in the image of the fifth field of view. They are consistent with the spectra of sphingomyelin⁹ and DOPC^{9,10} layers reported in the literature; the sm+ch spectrum shows two strong peaks between 2,850 and 2,900 cm⁻¹ before decreasing gradually, as well as a small shoulder around 2,950 cm⁻¹, while the spectrum of DOPC has multiple peaks of approximately the same amplitude between 2,850 and 2,950 cm⁻¹.

It is important to recall that the guide spectra used for each image came from that same image. The fact that it is possible to obtain such consistent results with self-guidance removes the need for additional measurements of samples with the same composition as the different domains, as well as the need for prior knowledge of the sample composition; as long as the domains are distinguishable by the strength of the signal at some wave number and their spectra are different enough from each other, it should be possible to observe domain segregation using only SRS, and it should be possible to retrieve the SRS spectra of the domains if one has the spectrum of the surrounding medium and knows the thickness of each domain type.

5.2.7. Domain composition studies

Following the FSC³ analysis of the images of the ternary sample, we tried to discern the amount of each type of lipid present in each of the domains using the average of the scaled-down spectra of the multilamellar sm+ch sample (shown in red at the bottom of figure 5.18, and henceforth $S_{\text{sm+ch}}$) and the scaled-down spectrum of the multilamellar DOPC sample (shown at the bottom of figure 5.19, and henceforth S_{DOPC}). Each of the two spectra from each image was modelled as a linear combination of the aforementioned multilamellar spectra, $S = \alpha S_{\text{DOPC}} + \beta S_{\text{sm+ch}}$, and then the coefficients of the linear combination (α and β) were normalised to their sum. Since the spectra resulting from guided FSC³ are considered to be more representative of the true spectra of the domains, this analysis was only done on said spectra. The linear combination was unconstrained; however, as with the FSC³ and the background gradient fitting, wave numbers where the signal displayed a strong, unexpected change that was clearly the result of a fluctuation in the laser intensity, such as 2,940 cm⁻¹ in the fourth image of the ternary sample, were removed from the linear-combination fitting.

A composition of DOPC:(sm+ch) = $-0.016:1.016 \pm 0.130$ was obtained for the LO domains and a composition of 1.002: -0.002 ± 0.085 was obtained for the LD domains; uncertainties are standard errors.

Recalling the expected partitioning of lipids (DOPC:sm:ch = 0:17:10 for the LO do-

mains and all the DOPC with the remaining cholesterol for the LD domains) and the composition of the ternary sample used here (2:2:1), we would expect the composition to be close to DOPC:(sm+ch) = 0:1 for the LO domains and 1:0 to the LD domains. However, a few things complicate the matter. First, there was insufficient cholesterol in the sample for the LO domains to have a composition DOPC:sm:ch = 0:17:10 and the LD domains to have no sphingomyelin; either the LD domains had a small amount of sphingomyelin, the LO domains had less cholesterol than in a 0:17:10 ratio, or both. Second, the sm+ch sample had a composition 0:7:3 = 0:23:10, which is close, but not equal, to 0:17:10; therefore, whatever the composition of the LO domains (which was likely 0:2:1 = 0:20:10), the sm+ch spectrum was not exactly the same as the spectrum of the LO domains, and it is possible (if the LD domains had some sphingomyelin) that the DOPC spectrum was not exactly the same as the LD spectrum either. Third, any local variations in the relative abundance of the three types of lipid would have caused local variations in the LO and/or LD compositions and thus in the LO and/or LD spectra, which might account for the discrepancies observed in the spectra of the different images. Fourth, the sm+ch and DOPC spectra have no significant differences at wave numbers below about $2,900\text{ cm}^{-1}$, so only the highest few wavenumbers of the scan range contribute to the determination of the domain compositions.

Despite all of these complications, the compositions of the LO and LD domains are within error of DOPC:(sm+ch) = 0:1 and 1:0, respectively. However, individual linear combinations sometimes had either a negative α or a negative β (and, correspondingly, the positive coefficient was greater than 1), which is not physically meaningful. This is probably due to a combination of the factors described in the previous paragraph and the presence of residual noise even after denoising and background correction.

As a sanity check, the same fitting was done on the individual sm+ch and DOPC spectra. The average compositions were $-0.973:1.973 \pm 0.148$ for the unilamellar sm+ch spectra, $0.000:1.000 \pm 0.045$ for the multilamellar sm+ch spectra, $-0.205:1.205 \pm 0.035$ for the unilamellar DOPC spectra, and $1.000:0.000$ for the multilamellar DOPC spectrum. The unilamellar spectra clearly have too much noise at high wave numbers for accurate fitting. The multilamellar samples are, naturally, extremely close to the ideal concentrations of 0:1 and 1:0, since they were used for the fitting in the first place and they have much less noise relative to the signal.

5.3. Conclusions

Single lipid bilayers are extremely thin (only 4–5 nm thick) and thus produce a very small SRS signal (a relative modulation of the order of 10^{-6} with our setup). Here, we have described a data analysis method for SRS images of single lipid bilayers so that the individual lipid domains are visible and can be separated into different spectral components without the need for guidance. While unlabelled multilamellar vesicles¹¹ and cells^{12,13} and deuterium-labelled cells^{14,15} and organelles¹⁶ have been imaged using SRS, this is, to our knowledge, the first time this has been achieved with unlabelled single bilayers. SRS pro-

vides a much stronger signal than does spontaneous Raman scattering, as does CARS. At the same time, the analysis of SRS images is much more straightforward than that of CARS images, requiring no PBS/glass correction, and SRS affords a higher contrast than CARS. Finally, unlike techniques like surface-enhanced Raman scattering, it is noninvasive. For these reasons, SRS was selected over the aforementioned alternatives.

Our technique consists of denoising the data via singular-value decomposition, removing background gradients due to spatial pulse-overlap inhomogeneities and focus drift, factorising the hyperspectral images into individual spectra with different spatial distributions, correcting for the spectral pulse-overlap inhomogeneities by using CARS images acquired simultaneously with the SRS images, and finally correcting for the PBS spectrum subtraction incurred in the background-correction step. The acquisition of CARS data does not require any additional optics beyond a dichroic beam splitter, a filter and a detector; does not require additional alignment, since it employs the same beams used for SRS; and does not require additional imaging time, since the CARS and SRS acquisitions are simultaneous. Background correction does require PBS-only regions to be present, large enough and distributed widely enough in the sample to give an accurate background gradient fitting.

Further to discrimination between domains, we were able to obtain the SRS spectra of said domains by guiding the spectral factorisation. These spectra are consistent with those reported in the literature. The guidance was local, meaning only the data from an image was used to guide the factorisation of that image; therefore, no additional measurements on reference samples are required, and no knowledge of the domain composition is required either. When we compared the spectra retrieved from the locally guided factorisation of the ternary sample with the spectra of such reference samples, we obtained results consistent with previous knowledge about the segregation of sphingomyelin, DOPC and cholesterol into LO and LD domains when mixed at room temperature in such a way that there is between 2 and 2.5 times as much sphingomyelin as cholesterol, namely that these two lipids form LO domains with little or no DOPC and that the DOPC forms LD domains by itself (see chapter 3).

One of the main benefits of this technique is that it is quantitative and, unlike surface-enhanced Raman scattering or fluorescence, noninvasive. This means that the lipid domain segregation behaviour we have seen is likely to be more representative of that which occurs in physiological conditions than the behaviour which would be observed with these other techniques, which require inserting particles into the sample. Furthermore, this proves lipid phases can be observed with a noninvasive optical technique in single lipid bilayers, which is in line with our main goal of imaging a phase transition (or absence thereof) in live neurons.

Nevertheless, there remain challenges to be addressed. Domain discrimination requires hyperspectral measurements, as well as spatially distributed measurements for PBS subtraction. Since the pixel dwell time is 1 ms, which is already too slow to resolve an action potential, this renders the technique inappropriate for action potential imaging in its current state. Given the extremely low signal-to-noise ratio of unprocessed images,

it is unlikely that the dwell time can be reduced further. Even if cells were kept at room temperature (if mammalian) or close to 4 °C (if lobster giant axons) so action potentials lasted about 5 ms instead of 1 ms, the signal for at least two wave numbers would need to be measured for lipid and for the surrounding medium, requiring 4 ms, still well above the maximum acquisition time required to Nyquist-sample the action potential.

Heating of the cell (and the damage this would cause) is unlikely to be a problem. We observed heating and bubble formation in the multilamellar DOPC samples but not in any unilamellar sample. Since the portion of an axon which would lie inside the focal volume is a unilamellar object surrounded on both sides by substances which are non-resonant at the studied wave numbers (much like the unilamellar supported lipid bilayers studied in this work), the heating caused by the laser is likely negligible. However, a systematic study of how much the laser power can be increased (which would reduce the required exposure time) is required in order to determine the extent to which the sampling rate can be increased without the cell being damaged.

5.4. References

1. Langbein W, Regan D, Pope I & Borri P (2018): *Heterodyne dual-polarization epideTECTED CARS microscopy for chemical and topographic imaging of interfaces*, APL Photonics 3, 092402
2. Masia F, Glen A, Stephens P, Borri P & Langbein W (2013): *Quantitative chemical imaging and unsupervised analysis using hyperspectral coherent anti-Stokes Raman scattering microscopy*, Analytical Chemistry 85, 10820–10828
3. Nahmad-Rohen A, Regan D, McPhee C, Masia F, Borri P & Langbein W (2019): *Determining the composition of domains in single lipid bilayers by coherent Raman scattering*, in preparation
4. Krafft C, Dietzek B, Popp J & Schmitt M (2012): *Raman and coherent anti-Stokes Raman scattering microspectroscopy for biomedical applications*, Journal of Biomedical Optics 17, 040801
5. Min W, Freudiger C W, Lu S & Xie X S (2011): *Coherent nonlinear optical imaging: beyond fluorescence microscopy*, Annual Review of Physical Chemistry 62, 507–530
6. Karuma A, Masia F, Borri P & Langbein W (2016): *Hyperspectral volumetric coherent anti-Stokes Raman scattering microscopy: quantitative volume determination and NaCl as non-resonant standard*, Journal of Raman Spectroscopy 47, 1167–1173
7. Regan D, Williams J, Borri P & Langbein W (2019): *Lipid bilayer thickness measured by quantitative DIC reveals phase transitions and effects of substrate hydrophilicity*, Langmuir 35, 13805–13814

8. Bazylewski P, Divigalpitiya R & Fanchini G (2017): *In situ Raman spectroscopy distinguishes between reversible and irreversible thiol modifications in L-cysteine*, Royal Society of Chemistry Advances 7, 2964–2970
9. Ando J, Kinoshita M, Cui J, Yamakoshi H, Dodo K, Fujita K, Murata M & Sodeoka M (2015): *Sphingomyelin distribution in lipid rafts of artificial monolayer membranes visualized by Raman microscopy*, Proceedings of the National Academy of Sciences of the United States of America 112, 4558–4563
10. Mari M, Mouras R, Downes A & Elfick A (2011): *Exploring the interactions between peptides and lipid bilayers using coherent anti-Stokes Raman scattering and two-photon fluorescence*, Proceedings of SPIE 8086, 80861G
11. Hofer M, Balla N K & Brasselet S (2017): *High-speed polarization-resolved coherent Raman scattering imaging*, Optica 4, 795–801
12. Yue S et al (2014): *Cholesteryl ester accumulation induced by PTEN loss and PI3K/AKT activation underlies human prostate cancer aggressiveness*, Cell Metabolism 19, 393–406
13. Zhang X, Roeffaers M B J, Basu S, Daniele J R, Fu D, Freudiger C W, Holtom G R & Xie X S (2012): *Label-free live cell imaging of nucleic acids using stimulated Raman scattering (SRS) microscopy*, ChemPhysChem 13, 1054–1059
14. Li J & Cheng J-X (2014): *Direct visualization of de novo lipogenesis in single living cells*, Scientific Reports 4, 6807
15. Hu F, Wei L, Zheng C, Shen Y & Min W (2014): *Live-cell vibrational imaging of choline metabolites by stimulated Raman scattering coupled with isotope-based metabolic labeling*, Analyst 139, 2312–2317
16. Shen Y, Zhao Z, Zhang L, Shi L, Shahriar S, Chan R B, di Paolo G & Min W (2017): *Metabolic activity induces membrane phase separation in endoplasmic reticulum*, Proceedings of the National Academy of Sciences of the United States of America 114, 13394–13399

Chapter 6

Conclusions

Our initial objective was to attempt to observe action potentials with a noninvasive, real-time optical measurement technique and ultimately determine whether action potentials involve a phase transition, as proposed by T R Heimburg and A D Jackson in 2005.

To do so, we first summarised the Heimburg-Jackson model and the traditionally accepted Hodgkin-Huxley model for action potential propagation in chapter 1.

The Hodgkin-Huxley model views the cell membrane as a capacitor and the ion channel proteins embedded in it as variable resistors in an RC circuit. It postulates that the voltage across the membrane changes due to an exchange of potassium and sodium ions between the cell exterior and interior through the ion channels, which are activated (and, in the case of the sodium channels, also deactivated) by “gating particles”. While these particles seem more a mathematical device than a biophysical mechanism and have never been detected, the model does accurately predict the time course of the action potential, in part thanks to it having a large number of parameters which could be adjusted to fit experimental data.

The Heimburg-Jackson model predicts that a phase transition in the cell membrane, from liquid-disordered to solid-ordered, is responsible for the action potential. Using a dispersion equation, it shows how this phase transition can propagate along the axon in a solitonic fashion without the need for ion exchange or gating particles. This isn't to say that ion exchange doesn't occur, merely that the mechanism responsible for the action potential is a phase transition and not the ions themselves.

This prediction implies important changes to the cell membrane during activity, and it can be experimentally tested in order to discern how complete the model is. It currently seems to be more complete than the Hodgkin-Huxley model given that it explains many of the phenomena concomitant with the action potential better than its counterpart. Armed with the mathematical description of the two models, we studied almost a century's worth of experiments involving action potentials and the aforementioned phenomena; this is the content of chapter 2.

These phenomena include an increase in the thickness of, and pressure in, the axon. While the Hodgkin-Huxley model would require such changes to be caused by elec-

trostriction or by the movement of water molecules attached to the ions, calculations assuming this is the case do not arrive at the observed values or even at the correct orders of magnitude. In contrast, the Heimburg-Jackson model explains these changes as a direct result of the membrane phase transition; studies have shown that lipid bilayers in the ordered phases are thicker than bilayers in the liquid-disordered phase by an amount consistent with the aforementioned observations in axons.

Other phenomena are opacity, scattering and birefringence changes in the axon, which further studies have shown are due to the membrane only and not the cytoplasm. This is consistent with both models, which focus on changes in the membrane properties. That said, the Hodgkin-Huxley model does not provide an explanation that is more than guesswork; it provides no mechanism by which the membrane might undergo the structural changes required to cause such optical effects. The Heimburg-Jackson model, on the other hand, predicts precisely the type of structural change which would cause these effects.

Release of heat from the axon has been reported together with the action potential. Both models offer plausible explanations for this —heat dissipation by the resistors in the Hodgkin-Huxley model and the latent heat of the phase transition in the Heimburg-Jackson model— but fall short of the observed release by many orders of magnitude; furthermore, while the qualitative nature of the heat release predicted by each model is different (irreversible in the Hodgkin-Huxley model and fully reversible in the Heimburg-Jackson model), the qualitative nature of the experimentally detected heat falls precisely in between (only partially reversible).

Finally, the Heimburg-Jackson model, being a purely thermodynamic model, is fully consistent with the Meyer-Overton rule, the finding that the potency of an anaesthetic is proportional to its solubility in lipid. In contrast, the Hodgkin-Huxley model requires anaesthetics to possess a degree of chemical specificity for ion channel proteins which is not always observed (for example, xenon is chemically inert, yet it functions as an anaesthetic).

From this study came two important results. The first is the conclusion that current knowledge of the physical mechanism behind action potential propagation is too nebulous to determine the extent to which each model accurately describes this mechanism. The second is a set of conditions which an imaging technique must have if it is to be non-invasive, be sensitive enough to resolve a phase transition in a single 4-nm lipid bilayer, and have a temporal resolution better than 1 ms.

With these conclusions in mind, electrophysiology techniques were learnt with the hope of being able to conduct an experiment which could clarify the nature of the action potential, an optical technique called interferometric reflectometry was theoretically developed and experimentally tested, and the data analysis procedure for stimulated Raman scattering (SRS) microscopy data was improved to make the technique sensitive to different thermodynamic phases in single lipid bilayers. The optical work was done on supported lipid bilayers, which are an appropriate model to study phase transitions in the cell membrane without the complications which arise from working with live cells, but the

aim was to ultimately move to live-cell imaging, as stated before.

The reflectometry theory, described in chapter 4, allows the independent determination of the thickness and refractive index of a thin layer from the layer's reflection coefficient, which includes multiple reflections inside the layer and is thus a complex quantity. It also takes into account the possibility of multiple layers with different thicknesses and refractive indices being stacked on each other and the finite numerical aperture of real microscope optics.

Such thin layers as a lipid bilayer produce a very small signal, however, meaning any reflectometry measurements on single lipid bilayers must have very little noise in order to have a signal-to-noise ratio high enough to accurately determine the bilayers' properties. Two possible strategies to reduce noise were presented: phase filtering and referencing. However, each brings its own problems. Phase filtering reduces noise at the cost of loss of information and blurring of the edges of the structures in the sample, and the amount of filtering which could be done without sacrificing sample information was insufficient to reduce to the noise enough for accurate measurements. Referencing, on the other hand, causes an image to contain information from two orthogonally polarised light beams incident on the sample at different positions, which means that the information corresponding to each polarisation must be separated. Three algorithms —one based on minimising the cross-correlation of the images, one based on combining the images and then minimising via gradient descent the error of a reconstruction of the original images using the combination, and one which used a genetic algorithm in place of gradient descent— were employed to attempt to separate the polarisations, but ultimately they were unable to achieve this; the reconstruction algorithms were successful only with simulated data, but the noise present in experimental data was too much for separation to be achieved, while the correlation algorithm was inadequate due to the data being complex.

Ultimately, the height of signal steps at bilayer edges was used to attempt to calculate the thickness and refractive index of the sample — not as an entire image, but only at the edges of the lipid bilayers. This yielded results which were 2–3 times too large in the thickness and slightly too small in the refractive index given what is known about the properties of a lipid bilayer, indicating that, as before, the noise was too strong for accurate measurements. Nevertheless, that the measured thicknesses were of the correct order of magnitude when the signal was so small is an indication that, while more work must still be done, interferometric reflectometry is close to the required sensitivity. Not only this, it was shown that, while the noise affects quantitative measurements, the technique is sensitive enough to distinguish between liquid-ordered and liquid-disordered domains — even with the current level of noise, it is capable of reaching sub-nanometre resolution in the axial direction, though it presently has insufficient accuracy.

SRS, together with our data analysis procedure, was more successful. This procedure, as well as its evolution throughout this work, is described in chapter 5. By using regions without lipid to correct for background gradients and singular-value decomposition to reduce the noise in the data, we were able to make the technique sensitive enough to distinguish lipid domains from each other without any external guidance. By simultaneously ac-

quiring coherent anti-Stokes Raman scattering images of the same field of view, it was possible to correct for the uneven pulse overlap as a function of wave number and ultimately obtain spectra consistent with the scientific literature. By comparing the domain spectra obtained with locally guided spectral factorisation with the spectra of pure DOPC and a mixture of sphingomyelin and cholesterol, it was determined that the lipids were fully segregated into DOPC liquid-disordered domains and sphingomyelin+cholesterol liquid-ordered domains, which is also consistent with the literature. Local guidance means that no external references other than PBS (or water) are necessary to retrieve the correct domain spectra.

Nevertheless, our SRS technique presently lacks the temporal resolution required for action potential observation. The pixel acquisition time employed is precisely the duration of an action potential, so, even if action potentials were slowed fivefold by cell cooling and only two spectral points were obtained, the temporal resolution would be insufficient. Because the signal from a single lipid bilayer is so low, decreasing the acquisition time sufficiently to Nyquist-sample action potentials while preserving the laser power would reduce the signal-to-noise ratio to a point where different thermodynamic phases would be undistinguishable, and increasing the laser power would risk damage to the cells. However, it must be said that a rigorous study of how much the acquisition time can be reduced without losing sensitivity to phase differences and how much the power can be increased without damaging the sample was not performed.

Two cell types were studied electrophysiologically: lobster nerves —which it was hoped would eventually lead to lobster giant axons, themselves selected for the ease of handling due to their large size— and mouse hippocampal neurons — selected because they are mammalian cells (which makes them more relevant to us as humans) and because they display spontaneous activity. Ultimately, this last trait, spontaneous neural activity, determined that mouse neurons were to be used going forwards.

Unfortunately, there was insufficient time remaining in the project to conduct all the experiments we wanted to conduct. Our electrophysiology experiments were limited to the preliminary stage of learning how to reliably use the equipment and which cell culture conditions resulted in active cells in the case of mouse hippocampal neurons. The optical and electrophysiological measurement techniques were not combined, although that was the next planned stage.

Future work will consist of attempting to increase the accuracy of our interferometric reflectometry setup so it can be used with single bilayers and, ultimately, combining electrophysiology measurements with single-point reflectometry or two-point SRS measurements on cultured mouse neurons to try to shed more light on the mechanisms underlying neural activity.

Appendix A

Additional reflectometry theory

A.I. Remarks on some mathematical properties of the reflection coefficient

Substituting $s = 0$ in equation 4.3, we see immediately that the zeros of s occur at $n_s = \sqrt{n_w n_g}$.

Substituting $n_s = \sqrt{n_w n_g}$ and $s = 0$ in equation 4.1, rearranging terms and writing $e^{i\psi}$ as $\cos(\psi) + i \sin(\psi)$, we obtain

$$\cos(2kd\sqrt{n_w n_g}) + i \sin(2kd\sqrt{n_w n_g}) = -\frac{\sqrt{n_w n_g} - n_w}{\sqrt{n_w n_g} + n_w} \frac{n_g + \sqrt{n_w n_g}}{n_g - \sqrt{n_w n_g}} = -1,$$

whereby

$$d = \frac{(2N - 1)\pi}{2k\sqrt{n_w n_g}}$$

with $N \in \mathbb{N}$. This means that the zeros occur at the value of d which is exactly at the centre of each period of s .

Thus, $d = (2N - 1)\lambda/4\sqrt{n_w n_g}$ and $n_s = \sqrt{n_w n_g}$ are the conditions the layer must have in order to constitute a perfect antireflective coating for light of wavelength λ . Conversely, having $d = N\lambda/2n_s$ for any value of n_s is, in terms of reflectivity, equivalent to not having a layer at all.

Let us now recall that there is a value of n_s below which $\text{Re}(s) > 0$. To find this value, we first multiply and divide s in equation 4.1 by the conjugate of the denominator of the right-hand-term to obtain

$$s = \frac{2(n_g^2 - n_s^2)(n_s^2 + n_w^2) + 2(n_s^2 - n_w^2)((n_g^2 + n_s^2)\cos(\psi) + 2in_g n_s \sin(\psi))}{(n_g - n_s)^2(n_s - n_w)^2 + (n_g + n_s)^2(n_s + n_w)^2 + 2(n_g^2 - n_s^2)(n_s^2 - n_w^2)\cos(\psi)},$$

where $\psi = 2kdn_s$.

The denominator is always positive because it is the product of a non-zero complex number and its conjugate, so the values of n_s for which the real part of the numerator is negative are exactly those values for which $\text{Re}(s) < 0$. We thus have the condition

$$\cos(2kdn_s) < -\frac{(n_g^2 - n_s^2)(n_s^2 + n_w^2)}{(n_g^2 + n_s^2)(n_s^2 - n_w^2)},$$

which requires

$$\frac{(n_g^2 - n_s^2)(n_s^2 + n_w^2)}{(n_g^2 + n_s^2)(n_s^2 - n_w^2)} < 1.$$

Now, this is a monotonically decreasing function of n_s in the range $n_w < n_s < n_g$, as shown by the fact that

$$\frac{\partial}{\partial n_s} \left(\frac{(n_g^2 - n_s^2)(n_s^2 + n_w^2)}{(n_g^2 + n_s^2)(n_s^2 - n_w^2)} \right) = -4n_s \frac{(n_g^2 - n_w^2)(n_w^2 n_g^2 + n_s^4)}{(n_g^2 + n_s^2)^2 (n_s^2 - n_w^2)^2} < 0.$$

Therefore, the value of n_s for which

$$\frac{(n_g^2 - n_s^2)(n_s^2 + n_w^2)}{(n_g^2 + n_s^2)(n_s^2 - n_w^2)} = 1$$

is the value of n_s below which $\text{Re}(s)$ is necessarily non-negative. If we write $n_s = \sqrt{n_w n_g}$, we obtain

$$\frac{(n_g^2 - n_s^2)(n_s^2 + n_w^2)}{(n_g^2 + n_s^2)(n_s^2 - n_w^2)} = \frac{n_w n_g (n_g^2 - n_w^2)}{n_w n_g (n_g^2 - n_w^2)} = 1,$$

so $\sqrt{n_w n_g}$ is the value we seek (figure A.1), as already suggested by figure 4.2.

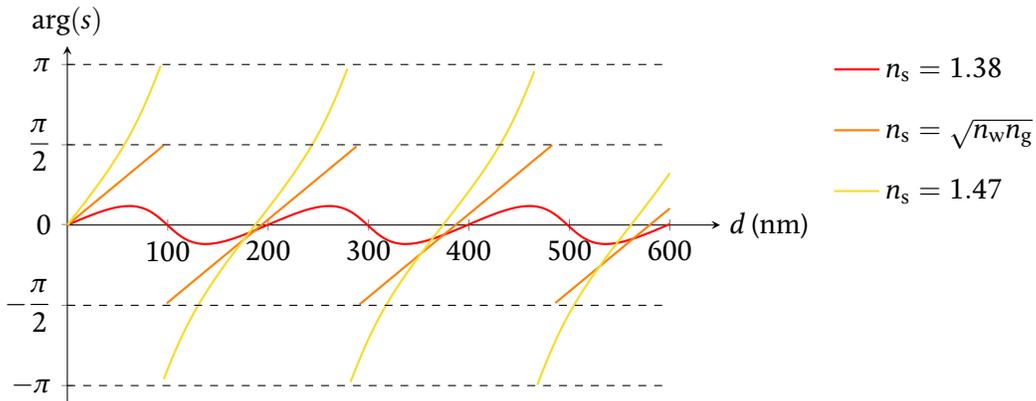


Figure A.1: The effect of n_s on the range of $\arg(s)$. If $n_s < \sqrt{n_w n_g}$ (red curve), $\arg(s)$ can only take small values; if $n_s = \sqrt{n_w n_g}$ (orange curve), $\arg(s)$ can take all values between $-\pi/2$ and $\pi/2$; if $n_s > \sqrt{n_w n_g}$ (yellow curve), $\arg(s)$ can take any value. Here, n_w , n_g and λ are as in previous figures.

A.2. Reflection by multiple layers

The model described in section 4.1.1 can be generalised to an arbitrary number of layers, as already mentioned in chapter 4.

A.2.1. Two layers

If the sample consists of two layers, each with its own thickness and refractive index (figure A.2), the situation becomes more complicated. Let the angle of incidence be denoted by θ , as before, and let θ_1 and θ_2 be the transmission angles in the first and second layers, respectively, given by Snell's law. Let us further denote the properties of the materials by d_1 , n_1 , d_2 and n_2 . We now have the following possibilities:

The light may be reflected at the interface between the glass and the first material and never enter the sample, as before. This contributes a term $s_0 = r_{g1}$.

The transmitted light may enter the first material and then be reflected any number $\ell + 1$ of times at the interface between the two materials and ℓ times at the interface between the first material and the glass before finally being transmitted back through the glass. This contributes a term similar to the one in the 1-layer case:

$$s_1 = t_{g1} r_{12} t_{1g} e^{2ik \frac{d_1 n_1}{\cos(\theta_1)}} \sum_{\ell=0}^{\infty} r_{1g}^{\ell} r_{12}^{\ell} e^{2\ell ik \frac{d_1 n_1}{\cos(\theta_1)}}.$$

Finally, the light may be transmitted through the interface between the two materials and enter the second layer (after any number of reflections within the first layer). This light can be either transmitted into the water, in which case it does not contribute to the reflection coefficient, or reflected back into the second layer and then transmitted back through the first layer and into the glass (after any number of reflections within either material or within both materials). Since we have different kinds of reflections, there are multiple ways of ordering them, so a multiplicity factor must be included in the two-layer term. If we have j reflections within the first bilayer, ℓ reflections within both bilayers and m reflections within the second bilayer, there are $(j + \ell + m)!/j!\ell!m!$ different ways of

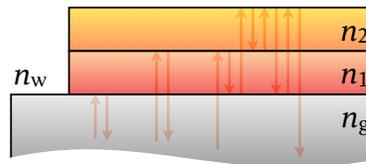


Figure A.2: Reflection from a thin sample consisting of two layers of different materials.

ordering them. The term contributed by this case is then

$$\begin{aligned}
 s_2 &= t_{g1}t_{12}r_{2w}t_{21}t_{1g}e^{2ik\left(\frac{d_1n_1}{\cos(\theta_1)} + \frac{d_2n_2}{\cos(\theta_2)}\right)} \times \\
 &\times \sum_{j=0}^{\infty} \sum_{\ell=0}^{\infty} \sum_{m=0}^{\infty} \frac{(j+\ell+m)!}{j!\ell!m!} r_{1g}^j r_{12}^j e^{2jik\frac{d_1n_1}{\cos(\theta_1)}} \times \\
 &\times r_{1g}^{\ell} t_{12}^{\ell} r_{2w}^{\ell} t_{21}^{\ell} e^{2lik\left(\frac{d_1n_1}{\cos(\theta_1)} + \frac{d_2n_2}{\cos(\theta_2)}\right)} r_{21}^m r_{2w}^m e^{2mik\frac{d_2n_2}{\cos(\theta_2)}}.
 \end{aligned}$$

The total reflection coefficient is, of course, $s = s_0 + s_1 + s_2$, which simplifies to

$$\begin{aligned}
 s &= r_{g1} + \frac{t_{g1}r_{12}t_{1g}e^{2ik\frac{d_1n_1}{\cos(\theta_1)}}}{1 - r_{1g}r_{12}e^{2ik\frac{d_1n_1}{\cos(\theta_1)}}} \\
 &+ \frac{t_{g1}t_{12}r_{2w}t_{21}t_{1g}e^{2ik\left(\frac{d_1n_1}{\cos(\theta_1)} + \frac{d_2n_2}{\cos(\theta_2)}\right)}}{1 - r_{1g}r_{12}e^{2ik\frac{d_1n_1}{\cos(\theta_1)}} - r_{1g}t_{12}r_{2w}t_{21}e^{2ik\left(\frac{d_1n_1}{\cos(\theta_1)} + \frac{d_2n_2}{\cos(\theta_2)}\right)} - r_{21}r_{2w}e^{2ik\frac{d_2n_2}{\cos(\theta_2)}}}. \quad (\text{A.1})
 \end{aligned}$$

As a sanity check, we may set $n_2 = n_1$ (which will turn the Fresnel coefficients for the interface between the two layers into $r_{12} = r_{21} = 0$ and $t_{12} = t_{21} = 1$). This yields

$$s = r_{g1} + \frac{t_{g1}r_{1w}t_{1g}e^{2ik\frac{(d_1+d_2)n_1}{\cos(\theta_1)}}}{1 - r_{1g}r_{1w}e^{2ik\frac{(d_1+d_2)n_1}{\cos(\theta_1)}}},$$

which (except for the cosines, which arise from non-normal incidence) is identical to equation 4.1 with $d = d_1 + d_2$ and $n_s = n_1$, as is expected. Similarly, substituting $n_2 = n_w$ results in $r_{2w} = 0$ and thus

$$s = r_{g1} + \frac{t_{g1}r_{1w}t_{1g}e^{2ik\frac{d_1n_1}{\cos(\theta_1)}}}{1 - r_{1g}r_{1w}e^{2ik\frac{d_1n_1}{\cos(\theta_1)}}},$$

which is again equation 4.1 with $d = d_1$ and $n_s = n_1$. Notably, substituting $d_2 = 0$ does not reduce equation A.1 to the 1-layer case; this is because r_{1w} (which would be the new reflection coefficient at the top of the first layer) is not necessarily equal to $t_{12}r_{2w}t_{21}$.

In the single-reflection approximation, equation A.1 becomes

$$s \approx r_{g1} + t_{g1}r_{12}t_{1g}e^{2ik\frac{d_1n_1}{\cos(\theta_1)}} + t_{g1}t_{12}r_{2w}t_{21}t_{1g}e^{2ik\left(\frac{d_1n_1}{\cos(\theta_1)} + \frac{d_2n_2}{\cos(\theta_2)}\right)}.$$

A.2.2. An arbitrary number of layers

For N layers with thicknesses d_1, \dots, d_N and refractive indices n_1, \dots, n_N , we may separate the problem into $N + 1$ partial reflection coefficients s_0, \dots, s_N , each taking one more layer than the previous one, as we did for the 1-layer case and the 2-layer case. We may obtain a generalisable expression if we assign the index 0 to the glass slide and the index $N + 1$ to the water.

We first define the operator Ξ , which we will use to denote nested sums, as follows:

$$\Xi_{\ell, m, M}^{q, Q} \equiv \sum_{\ell_{q, q} = m}^M \sum_{\ell_{q, q+1} = m}^M \cdots \sum_{\ell_{q, Q} = m}^M \sum_{\ell_{q+1, q+1} = m}^M \cdots \sum_{\ell_{q+1, Q} = m}^M \cdots \sum_{\ell_{Q, Q} = m}^M .$$

Here, the first index of ℓ runs from q to Q and the second one runs from the first one's value to Q , so there are $(Q - q + 1)(Q - q + 2)/2$ sums. For example, if $m = 0$, $M = \infty$, $q = 1$ and $Q = 3$, then, for any function f ,

$$\Xi_{\ell, 0, \infty}^{1, 3} f(\{\ell\}) = \sum_{\ell_{1,1}=0}^{\infty} \sum_{\ell_{1,2}=0}^{\infty} \sum_{\ell_{1,3}=0}^{\infty} \sum_{\ell_{2,2}=0}^{\infty} \sum_{\ell_{2,3}=0}^{\infty} \sum_{\ell_{3,3}=0}^{\infty} f(\ell_{1,1}, \ell_{1,2}, \ell_{1,3}, \ell_{2,2}, \ell_{2,3}, \ell_{3,3}).$$

Now that we have defined the nested-sum operator, we may calculate the reflection coefficient of a collection of N parallel layers.

The j -th partial reflection coefficient is the combination of all possible reflections within each of the first j layers, each pair of adjacent layers (with the corresponding transmission coefficients) within the first j layers, and so on, taking all sets of ℓ adjacent layers with $1 \leq \ell \leq j$ and remembering to write the multiplicity factor for each combination of reflections:

$$s_j = r_{j,j+1} \left(\prod_{\ell=1}^j t_{\ell-1, \ell} t_{\ell, \ell-1} \right) e^{2ik \sum_{\ell=1}^j \frac{d_{\ell} n_{\ell}}{\cos(\theta_{\ell})}} \times$$

$$\times \Xi_{m, 0, \infty}^{1, j} \left(\frac{\left(\sum_{p=1}^j \sum_{q=p}^j m_{p, q} \right)!}{\prod_{p=1}^j \prod_{q=p}^j m_{p, q}!} \prod_{p=1}^j \prod_{q=p}^j r_{p, p-1}^{m_{p, q}} r_{q, q+1}^{m_{p, q}} e^{2m_{p, q} ik \sum_{\ell=p}^q \frac{d_{\ell} n_{\ell}}{\cos(\theta_{\ell})}} \times \right.$$

$$\left. \times \prod_{\ell=p+1}^q t_{\ell-1, \ell}^{m_{p, q}} t_{\ell, \ell-1}^{m_{p, q}} \right).$$

The total reflection coefficient is thus

$$\begin{aligned}
 s = & \sum_{j=0}^N r_{j,j+1} \left(\prod_{\ell=1}^j t_{\ell-1,\ell} t_{\ell,\ell-1} \right) e^{2ik \sum_{\ell=1}^j \frac{d_{\ell} n_{\ell}}{\cos(\theta_{\ell})}} \times \\
 & \times \sum_{m,0,\infty}^{1,j} \left(\frac{\left(\sum_{p=1}^j \sum_{q=p}^j m_{p,q} \right)!}{\prod_{p=1}^j \prod_{q=p}^j m_{p,q}!} \prod_{p=1}^j \prod_{q=p}^j r_{p,p-1}^{m_{p,q}} r_{q,q+1}^{m_{p,q}} e^{2m_{p,q} ik \sum_{\ell=p}^q \frac{d_{\ell} n_{\ell}}{\cos(\theta_{\ell})}} \times \right. \\
 & \left. \times \prod_{\ell=p+1}^q t_{\ell-1,\ell}^{m_{p,q}} t_{\ell,\ell-1}^{m_{p,q}} \right) \quad (A.2)
 \end{aligned}$$

with the understanding that, if $q > Q$, then

$$\begin{aligned}
 \sum_{\ell=q}^Q f &= 0, \\
 \prod_{\ell=q}^Q f &= 1, \\
 \sum_{\ell,m,M}^{q,Q} f &= 1
 \end{aligned}$$

for any function f .

Substitution of $N = 0$, $N = 1$ or $N = 2$ into equation A.2 turns s into the glass-water reflection coefficient (r_{gw}), the 1-layer reflection coefficient (equation 4.1) or the 2-layer reflection coefficient (equation A.1), respectively.

In the single-reflection approximation, equation A.2 becomes

$$s \approx \sum_{j=0}^N r_{j,j+1} \left(\prod_{\ell=1}^j t_{\ell-1,\ell} t_{\ell,\ell-1} \right) e^{2ik \sum_{\ell=1}^j \frac{d_{\ell} n_{\ell}}{\cos(\theta_{\ell})}}.$$

Again, setting $N = 0$, $N = 1$ or $N = 2$ turns this into r_{gw} or the 1- or 2-layer single-reflection approximations, respectively.

Appendix B

Correlation algorithm for interferometric reflectometry

B.I. The code

Reproduced below is the code for the correlation programme. It was written and tested on Mathematica 10.

```
(* clear all variables *)
Clear[Evaluate[Context[]<>"*"]];

(* set export directory to notebook directory *)
SetDirectory[NotebookDirectory[]];

(* read data from files *)
(* assumes data is contained in six files: real part of D_j in files
   called Djx, imaginary part in Djy, amplitude in Djr, where j=1 or
   j=2 *)
D1xf=SystemDialogInput["FileOpen"];
D1x=ReadList[D1xf,Number,RecordLists->True];
D1yf=SystemDialogInput["FileOpen",DirectoryName@D1xf];
D1y=ReadList[D1yf,Number,RecordLists->True];
D1rf=SystemDialogInput["FileOpen",DirectoryName@D1yf];
D1r=ReadList[D1rf,Number,RecordLists->True];
D2xf=SystemDialogInput["FileOpen",DirectoryName@D1rf];
D2x=ReadList[D2xf,Number,RecordLists->True];
D2yf=SystemDialogInput["FileOpen",DirectoryName@D2xf];
D2y=ReadList[D2yf,Number,RecordLists->True];
D2rf=SystemDialogInput["FileOpen",DirectoryName@D2yf];
D2r=ReadList[D2rf,Number,RecordLists->True];
```

```

(* rescale Re & Im data using amplitude *)
  x=Sqrt[D1r^2/(D1x^2+D1y^2)]D1x;
  y=Sqrt[D1r^2/(D1x^2+D1y^2)]D1y;
  D1x=x;
  D1y=y;
  Clear[x,y];
  D1=D1x+I*D1y;
  x=Sqrt[D2r^2/(D2x^2+D2y^2)]D2x;
  y=Sqrt[D2r^2/(D2x^2+D2y^2)]D2y;
  D2x=x;
  D2y=y;
  Clear[x,y];
  D2=D2x+I*D2y;

(* calculate data amplitude & phase *)
  D1r=Abs[D1];
  D1p=Arg[D1];
  D2r=Abs[D2];
  D2p=Arg[D2];

(* determine image dimensions *)
  rows=Length[D1];
  cols=Length[Transpose[D1]];

(* apply high-pass filter to images *)
  sigma=(rows+cols)/10; (* this value was found to be appropriate for the
                          tested data *)
  D1hpf=InverseFourier[Fourier[D1]
    *ParallelTable[1-Exp[-(x^2+y^2)/(2sigma^2)],
      {y,1,rows},{x,1,cols}]
    *ParallelTable[1-Exp[-((x-cols)^2+y^2)/(2sigma^2)],
      {y,1,rows},{x,1,cols}]
    *ParallelTable[1-Exp[-(x^2+(y-rows)^2)/(2sigma^2)],
      {y,1,rows},{x,1,cols}]
    *ParallelTable[1-Exp[-((x-cols)^2+(y-cols)^2)
      /(2sigma^2)],
      {y,1,rows},{x,1,cols}]]];
  D2hpf=InverseFourier[Fourier[D2]
    *ParallelTable[1-Exp[-(x^2+y^2)/(2sigma^2)],
      {y,1,rows},{x,1,cols}]
    *ParallelTable[1-Exp[-((x-cols)^2+y^2)/(2sigma^2)],
      {y,1,rows},{x,1,cols}]

```

```

*ParallelTable[1-Exp[-(x^2+(y-rows)^2)/(2sigma^2)],
              {y,1,rows},{x,1,cols}]
*ParallelTable[1-Exp[-((x-cols)^2+(y-cols)^2)
              /(2sigma^2)],
              {y,1,rows},{x,1,cols}]];

(* correlation function *)
(* calculates normalised correlation of d1 & d2 at a displacement
   (cn,rn) *)
corr[d1_,d2_,rn_,cn_] := Flatten[Conjugate[d1[[Max[1,1+rn];;
                                             Min[rows,rows+rn],
                                             Max[1,1+cn];;
                                             Min[cols,cols+cn]]]
                                -Mean[Flatten[d1]]]
    .Flatten[d2[[Max[1,1-rn];;
                Min[rows,rows-rn],
                Max[1,1-cn];;
                Min[cols,cols-cn]]]
              -Mean[Flatten[d2]]]
  /((rows-Abs[rn])*(cols-Abs[cn])
   *Variance[Flatten[d1]]*Variance[Flatten[d2]]);

(* partition parameters *)
wstart=0.5;
wend=2;
astart=0;
aend=2Pi;
steps=100;
wstep=(wend-wstart)/steps;
astep=(aend-astart)/steps;

(* partition the (w1,a1,w2,a2,w3,a3) space & calculate the correlation at
   each point in the partition *)
partition=ParallelTable[w1=wstart+wstep*x1;
                       a1=astart+astep*y1;
                       c1=w1*Exp[I*a1];
                       Clear[w1,a1];
                       w2=wstart+wstep*x2;
                       a2=astart+astep*y2;
                       c2=w2*Exp[I*a2];
                       Clear[w2,a2];
                       w3=wstart+wstep*x3;

```

```

a3=astart+astep*y3;
c3=w3*Exp[I*a3];
Clear[w3,a3];
Abs[corr[D1+c1*D2,c3(c2*D1+D2),0,0]],
{x1,0,steps},{y1,0,steps},{x2,0,steps},
{y2,0,steps},{x3,0,steps},{y3,0,steps}];

(* interpolate calculated values *)
index=Table[{j1,j2,j3,j4,j5,j6},
{j1,1,steps+1},{j2,1,steps+1},{j3,1,steps+1},
{j4,1,steps+1},{j5,1,steps+1},{j6,1,steps+1}];
intcorr=Interpolation[Flatten[Table[{index[[j1,j2,j3,j4,j5,j6]],
partition[[j1,j2,j3,j4,j5,j6]],
{j1,1,steps+1},{j2,1,steps+1},
{j3,1,steps+1},{j4,1,steps+1},
{j5,1,steps+1},{j6,1,steps+1}],3],
Method->"Spline",InterpolationOrder->3];
min=FindMinimum[{intcorr[w1,a1,w2,a2,w3,a3],
wstart<w1<wend&&wstart<w2<wend&&wstart<w3<wend
&&0<=a1<=2Pi&&0<=a2<=2Pi&&0<=a3<=2Pi},
{w1,0.5},{a1,0},{w2,0.5},{a2,0},{w3,0.5},{a3,0}];

(* perform linear combination with optimal parameters *)
w1=w1/.min[[2]];
a1=a1/.min[[2]];
c1=w1*Exp[I*a1];
w2=w2/.min[[2]];
a2=a2/.min[[2]];
c2=w2*Exp[I*a2];
w3=w3/.min[[2]];
a3=a3/.min[[2]];
c3=w3*Exp[I*a3];
A1=D1+c1*D2;
A2=c3(c2*D1+D2);

(* display separated images *)
Print[{ListDensityPlot[Abs[A1],
ColorFunction->"Rainbow",
FrameLabel->{Row[{x," (px)"}],Row[{y," (py)"}]}},
PlotLabel->Style[Row[{"|", "A"Subscript[1], "|"}],
Large],
PlotRange->Full,PlotRangePadding->0,

```

```

LabelStyle->FontFamily->"Calluna",
PlotLegends->Placed[Automatic,Right],
ImageSize->Small,AspectRatio->rows/cols],
ListDensityPlot[Arg[A1],
ColorFunction->"Rainbow",
FrameLabel->{Row[{x," (px)"}],Row[{y," (px)"}]}},
PlotLabel->Style[Row[{"arg(", "A"Subscript[1],
                    ")}"],
                Large],
PlotRange->Full,PlotRangePadding->0,
LabelStyle->FontFamily->"Calluna",
PlotLegends->Placed[Automatic,Right],
ImageSize->Small,AspectRatio->rows/cols],
ListDensityPlot[Abs[A2],
ColorFunction->"Rainbow",
FrameLabel->{Row[{x," (px)"}],Row[{y," (px)"}]}},
PlotLabel->Style[Row[{"|", "A"Subscript[2], "|"}],
                Large],
PlotRange->Full,PlotRangePadding->0,
LabelStyle->FontFamily->"Calluna",
PlotLegends->Placed[Automatic,Right],
ImageSize->Small,AspectRatio->rows/cols],
ListDensityPlot[Arg[A2],
ColorFunction->"Rainbow",
FrameLabel->{Row[{x," (px)"}],Row[{y," (px)"}]}},
PlotLabel->Style[Row[{"arg(", "A"Subscript[2],
                    ")}"],
                Large],
PlotRange->Full,PlotRangePadding->0,
LabelStyle->FontFamily->"Calluna",
PlotLegends->Placed[Automatic,Right],
ImageSize->Small,AspectRatio->rows/cols]};

```

B.2. Results with simulated data

Figure B.1 shows the amplitude and phase of the reflection coefficient s of a simulated sample consisting of a single 4-nm layer of material with refractive index 1.44 for each of the two orthogonal linear polarisations; the spatial shift, or shear, simulates the effect of a Wollaston prism on the incident beam. The simulated data is noiseless, and the values were calculated assuming normal incidence (equation 4.1) and the already familiar parameters used in previous examples ($n_w = 1.333$, $n_g = 1.518$, $\lambda = 550$ nm). Figure B.2 shows what would be detected in this case assuming perfect detectors and assuming the beam is

circularly polarised before reaching the Wollaston prism; at each position \vec{r} , the simulated data shown (D_1 and D_2) was calculated as

$$D_1(\vec{r}) = \frac{s_H(\vec{r})}{\sqrt{2}} + i \frac{s_V(\vec{r})}{\sqrt{2}},$$

$$D_2(\vec{r}) = i \frac{s_H(\vec{r})}{\sqrt{2}} + \frac{s_V(\vec{r})}{\sqrt{2}},$$

where $s_j(\vec{r})$ is the reflection coefficient of the sample at \vec{r} (about $0.6466 + 0.003595i$ if there is part of the layer at \vec{r} and $r_{\text{gw}} = 0.06489$ otherwise) for polarisation j and it is assumed that the images on the top row of figure B.1 correspond to the \vec{H} polarisation and those on the bottom row correspond to the \vec{V} polarisation. By construction, the linear-combination parameters we expect to find using the correlation method are $c_1 = c_2 = \pm i$ and $c_3 = 1$,

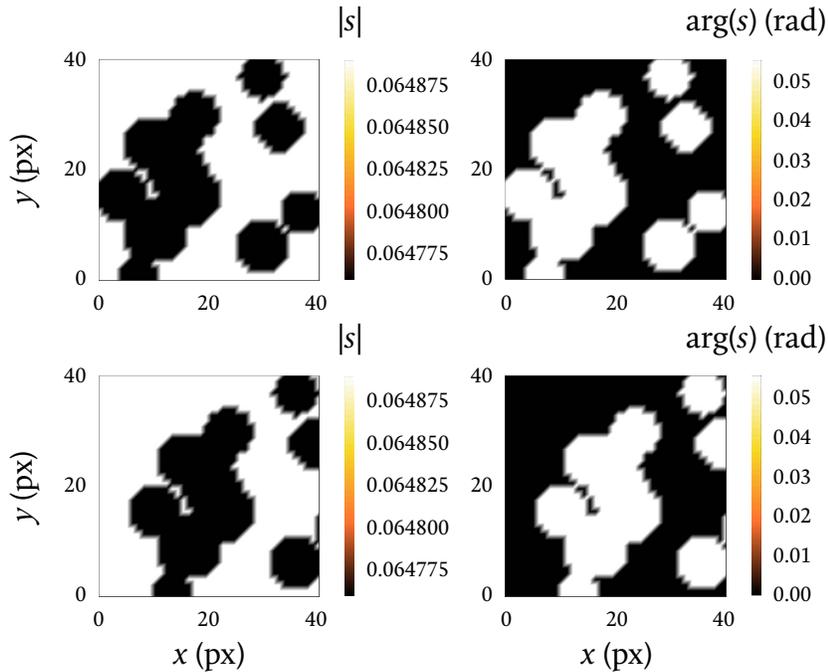


Figure B.1: Simulated amplitude (left) and phase (right) of the reflection coefficient of 4-nm sample with refractive index 1.44 as would be seen by the co- (top) and cross-polarised (bottom) components of a probe beam after passing through a Wollaston prism. The shear induced by the hypothetical Wollaston prism is 6 pixels in the horizontal direction. No noise has been added to this particular simulation.

since

$$\begin{aligned} D_1 - iD_2 &= \frac{s_H + is_V}{\sqrt{2}} - i \frac{is_H + s_V}{\sqrt{2}} = \sqrt{2}s_H, \\ -iD_1 + D_2 &= -i \frac{s_H + is_V}{\sqrt{2}} + \frac{is_H + s_V}{\sqrt{2}} = \sqrt{2}s_V \end{aligned}$$

and

$$\begin{aligned} D_1 + iD_2 &= \frac{s_H + is_V}{\sqrt{2}} + i \frac{is_H + s_V}{\sqrt{2}} = \sqrt{2}is_V, \\ iD_1 + D_2 &= i \frac{s_H + is_V}{\sqrt{2}} + \frac{is_H + s_V}{\sqrt{2}} = \sqrt{2}is_H \end{aligned}$$

(recall that obtaining a multiple of the original images is acceptable).

In this case, the high-pass-filtered data (figure B.3) was calculated as

$$(D_j)_f = \mathcal{F}^{-1} \left[1 - e^{-\frac{\xi^2 + \eta^2}{2\sigma^2}} \mathcal{F}[D_j] \right]$$

with $\sigma = 0.02 \text{ px}^{-1}$.

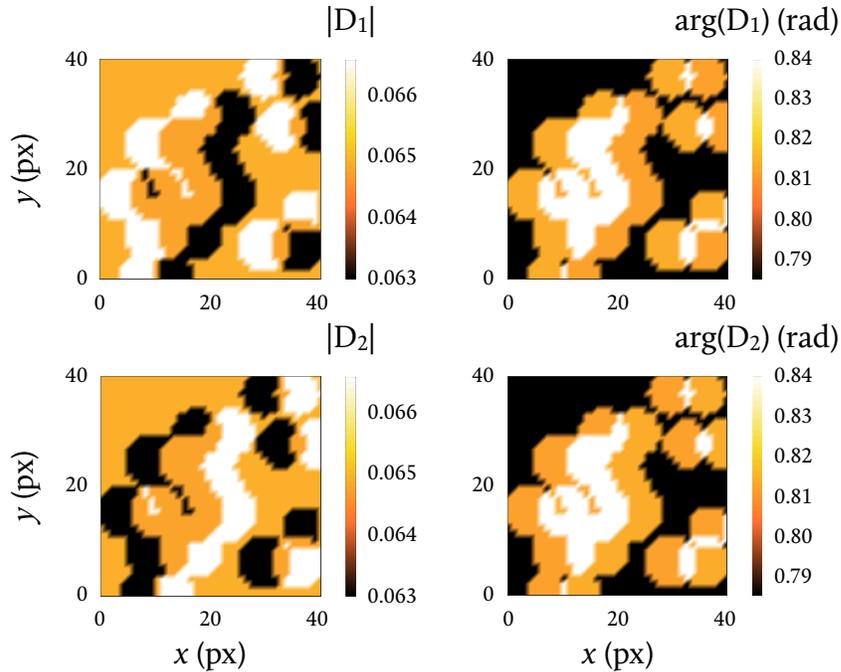


Figure B.2: Simulated amplitude (left) and phase (right) of the signal that would be detected by the interferometric reflectometry setup, denoted by D_j and calculated as linear combinations of the data shown on figure B.1.

Figure B.4 shows the correlation of the simulated $(D_1)_f$ and $(D_2)_f$, given by equation 4.6, as a function of the displacement $\vec{r} = (r_x, r_y)$ for $r_x, r_y \in [-15 \text{ px}, 15 \text{ px}]$; the projections are the same as those in section 4.2.2. A small peak is visible at $\vec{r} = (0, 0)$, and very strong peaks are visible at $(\pm 6 \text{ px}, 0)$, which corresponds to the shear.

It should be noted at this point that maximising either of the lateral correlation peaks would be an equally valid approach. We have chosen to minimise the correlation peak at $(0, 0)$ because, although the calculation of the correlation is somewhat more computationally expensive (because the overlap area is the entire image and is thus larger), it is simpler to code and it does not require finding the shear (which for experimental data is not necessarily parallel to one of the scan axes and may even be a fractional number of pixels) and shifting one of the images.

Figure B.5 shows projections of the correlation of the $(A_j)_f$ (calculated using equations 4.5 with the high-pass-filtered data) setting $c_3 = 1$. The correlation as a function of any two variables in the 4-dimensional space that results from fixing c_3 is qualitatively similar to that of experimental data (figure 4.17), but the range (the difference between the correlation's maximum and minimum values) is reduced in the experimental data due to noise; here it is much greater.

Since the data is complex, the magnitude of the product of two images with high spatial correspondence of the structures they show is not necessarily greater than that of two

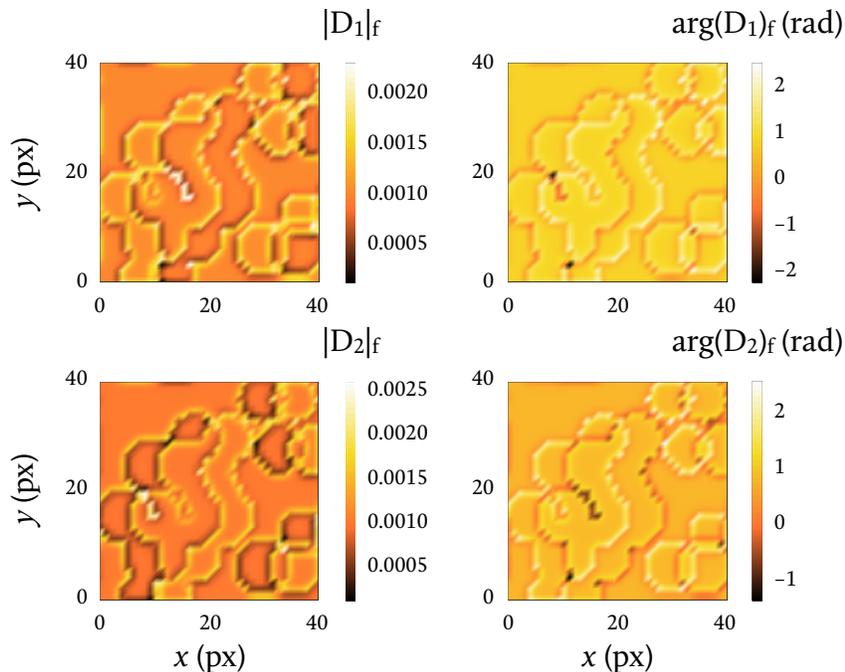


Figure B.3: Amplitude (left) and phase (right) of the simulated data from figure B.2 after a high-pass filter has been applied to the data.

images with less spatial correspondence, as was seen in the correlation graphs for simulated data in section 4.2.2. Additionally, noise in the data can translate into noise in the correlation.

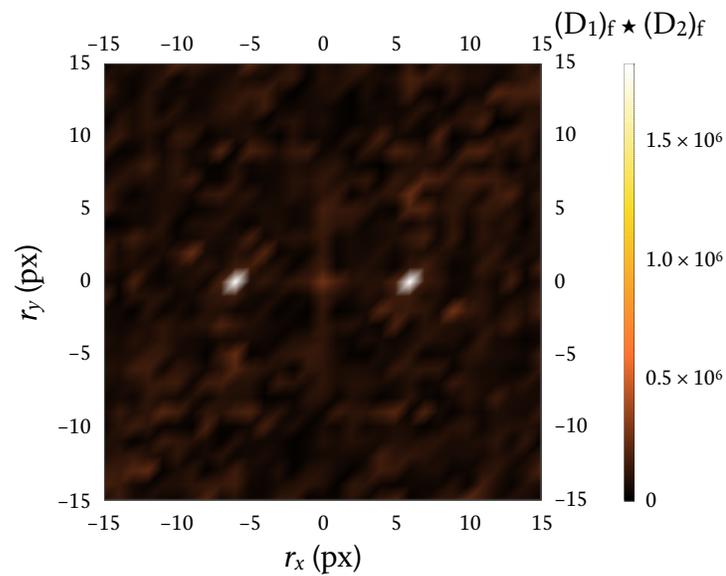


Figure B.4: Correlation of $(D_1)_f$ and $(D_2)_f$ from figure B.3 for different values of the displacement \vec{r} . There are two peaks at positions symmetric about the origin; said positions correspond to the shear introduced by the hypothetical Wollaston prism in the simulation. There is also a much weaker peak at the origin.

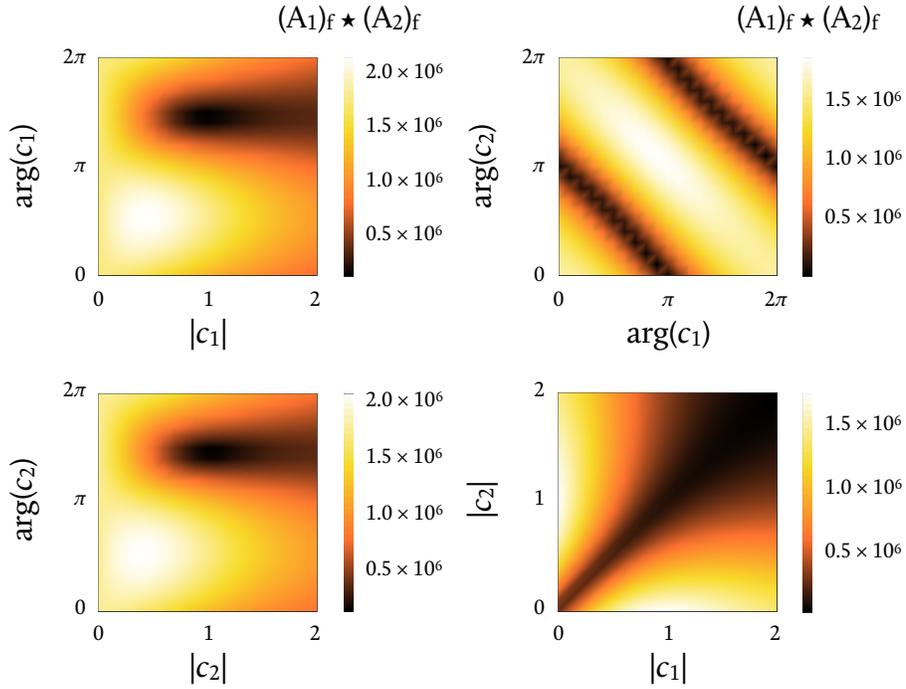


Figure B.5: Several 2-dimensional projections of the correlation at zero displacement of the simulated $(A_1)_f$ and $(A_2)_f$ for $c_3 = 1$. In $(|c_1|, \arg(c_1), |c_2|, \arg(c_2), |c_3|, \arg(c_3))$ space, the projections are as follows:

$$\begin{aligned} & \left(|c_1|, \arg(c_1), 1, \frac{3\pi}{2}, 1, 0 \right) \quad \left(1, \arg(c_1), 1, \arg(c_2), 1, 0 \right) \\ & \left(1, \frac{3\pi}{2}, |c_2|, \arg(c_2), 1, 0 \right) \quad \left(|c_1|, \frac{3\pi}{2}, |c_2|, \frac{3\pi}{2}, 1, 0 \right) \end{aligned}$$

The projections look qualitatively similar to the same projections for experimental data, but the correlation range is enlarged here. Notably, there is no minimum at $(1, 3\pi/2, 1, 3\pi/2, 1, 0)$, which corresponds to $(c_1, c_2, c_3) = (-i, -i, 1)$, one of the points at which the polarisations would be perfectly separated.

Appendix C

Gradient descent algorithm for interferometric reflectometry

C.I. The code

Reproduced below is the code for the comparison programme using a gradient descent algorithm. It was written and tested on MATLAB 2017b and tested on MATLAB 2018b. The code includes a section for the generation of test data because it was tested more extensively with simulated data than with experimental data.

```
% global variables
    global rows cols shearx sheary sigmar sigmap D1 D2 delta

% build test data
    rows=40;
    cols=40;
    s=0.06466+0.003595i; % one lipid bilayer
    w=0.06489; % no layer
    I1=w*ones(rows,cols);
    I2=I1;
    shearx=6;
    sheary=0;
    num=randi([1,3])+randi([1,3])+randi([1,3])+randi([1,3])+randi([1,3]);
    sizeavg=randi([1,3])+randi([1,3]);
    size=random('Normal',sizeavg,0.5,1,num);
    posx=randi([1,cols],1,num);
    posy=randi([1,rows],1,num);
    for m=1:num
        for j=1:rows
            for l=1:cols
                if abs((j-posy(m)+sheary/2)^2+(l-posx(m)+shearx/2)^2) ...
```

```

        <=size(m)^2
        I1(j,1)=s;
    end
    if abs((j-posy(m)-sheary/2)^2+(1-posx(m)-shearx/2)^2) ...
        <=size(m)^2
        I2(j,1)=s;
    end
end
end
end
end

% add noise to data
amp=2000*imag(s); % 2,000 approximately simulates experimentally
                % observed noise level

sigma=imag(s);
cutoff=1/10; % this value has been found to be appropriate for the
            % tested data
noise=amp*random('Normal',0,sigma,1,rows*cols);
noise=fft(noise);
for m=1:rows*cols
    noise(m)=exp(-(m/(2*rows*cols*cutoff)).^2))*noise(m);
end
noise=abs(iff(noise));
noise=transpose(reshape(noise,cols,rows));
I1=abs(I1).*exp(1i*(atan2(imag(I1),real(I1))+noise));
I2=abs(I2).*exp(1i*(atan2(imag(I2),real(I2))+noise));

% build "measurements" (D1 & D2) as linear combinations of I1 & I2
a1=1/sqrt(2);
a2=1i/sqrt(2);
a3=1i/sqrt(2);
a4=1/sqrt(2);
D1=a1*I1+a2*I2;
D2=a3*I1+a4*I2;
figD=figure('pos',[1,535,560,420]);
subplot(2,2,1)
    image(abs(D1),'CDataMapping','scaled'), ...
    colormap(jet),colorbar,title(colorbar,''), ...
    set(gcf,'name','data','numbertitle','off'), ...
    xlabel('\it{x} (px)'),ylabel('\it{y} (px)')
    pbaspect([1,1,1])
    title('|D_{1}|')
```

```

subplot(2,2,2)
    image(atan2(imag(D1),real(D1)),'CDataMapping','scaled'), ...
    colormap(jet),colorbar,title(colorbar,''), ...
    set(gcf,'name','data','numbertitle','off'), ...
    xlabel('\it{x} (px)'),ylabel('\it{y} (px)')
    pbaspect([1,1,1])
    title('arg(D_{1})')
subplot(2,2,3)
    image(abs(D2),'CDataMapping','scaled'), ...
    colormap(jet),colorbar,title(colorbar,''), ...
    set(gcf,'name','data','numbertitle','off'), ...
    xlabel('\it{x} (px)'),ylabel('\it{y} (px)')
    pbaspect([1,1,1])
    title('|D_{2}|')
subplot(2,2,4)
    image(atan2(imag(D2),real(D2)),'CDataMapping','scaled'), ...
    colormap(jet),colorbar,title(colorbar,''), ...
    set(gcf,'name','data','numbertitle','off'), ...
    xlabel('\it{x} (px)'),ylabel('\it{y} (px)')
    pbaspect([1,1,1])
    title('arg(D_{2})')

% parameters
sigmar=1;
sigmap=1;
delta=0.001;
params(1,:)=[random('Uniform',0.8,1.2), ...
             random('Uniform',pi/2-0.2,pi/2+0.2), ...
             random('Uniform',0.8,1.2), ...
             random('Uniform',pi/2-0.2,pi/2+0.2), ...
             random('Uniform',0.8,1.2), ...
             random('Uniform',-0.2,0.2)];

gamma=1*10^0;
alpha=1;

% iteratively determine c1, c2 & c3
errorsiter=0;
disp(' ');
tic;
disp('programme start');
disp(' ');
p=1;

```

```

while alpha>10^-9
    params=[params; [0,0,0,0,0,0]];
    errorsiter(p)=seperror(params(p,1),params(p,2),params(p,3), ...
        params(p,4),params(p,5),params(p,6));
    if p>=2
        gamma=[gamma,0];
        gamma(p)=min(gamma(p-1), ...
            abs((params(p,:) - params(p-1,:)) ...
                *transpose(sepgrad(params(p,1),params(p,2), ...
                    params(p,3),params(p,4), ...
                    params(p,5),params(p,6)) ...
                    -sepgrad(params(p-1,1), ...
                        params(p-1,2), ...
                        params(p-1,3), ...
                        params(p-1,4), ...
                        params(p-1,5), ...
                        params(p-1,6))))
            / (norm(sepgrad(params(p,1),params(p,2), ...
                params(p,3),params(p,4), ...
                params(p,5),params(p,6)) ...
                -sepgrad(params(p-1,1),params(p-1,2), ...
                    params(p-1,3),params(p-1,4), ...
                    params(p-1,5),params(p-1,6)))));
    end
    params(p+1,:)=params(p,:)-alpha*gamma(p) ...
        *sepgrad(params(p,1),params(p,2), ...
            params(p,3),params(p,4), ...
            params(p,5),params(p,6));
    while seperror(params(p+1,1),params(p+1,2),params(p+1,3), ...
        params(p+1,4),params(p+1,5),params(p+1,6)) ...
        >seperror(params(p,1),params(p,2),params(p,3), ...
            params(p,4),params(p,5),params(p,6))
        alpha=alpha/sqrt(2);
        params(p+1,:)=params(p,:)-alpha*gamma(p) ...
            *sepgrad(params(p,1),params(p,2),params(p,3), ...
                params(p,4),params(p,5),params(p,6));
    end
    alpha=alpha*sqrt(2);
    X=['iteration ',num2str(p-1),': error=',num2str(errorsiter(p))];
    disp(X);
    p=p+1;
end

```

```

% display results
disp(' ');
X=['after ',num2str(p-2),' iterations, the error has gone from ', ...
   num2str(errorsiter(1)),' to ',num2str(errorsiter(p-1))];
disp(X);
disp(' ');
disp('initial parameters:');
disp(params(1,:));
disp('final parameters:');
disp(params(p,:));
x=0:p-2;
c1=params(p,1)*exp(1i*params(p,2));
c2=params(p,3)*exp(1i*params(p,4));
c3=params(p,5)*exp(1i*params(p,6));
A1=D1+c1.*D2;
A2=c3.*(c2.*D1+D2);
figA=figure('pos',[1,1,891,420]);
subplot(2,3,1)
plot(x,log10(errorsiter(:)),'LineWidth',1,'Color',[1,0,0]), ...
set(gcf,'name','results','numbertitle','off'), ...
xlabel('iteration'),ylabel('log_{10}(error)')
title('error evolution')
subplot(2,3,4)
plot([x,0],params(1:p,1),'LineWidth',1,'Color',[1,0,0]), ...
set(gcf,'name','results','numbertitle','off'), ...
xlabel('iteration'),ylabel('')
hold on
plot([x,0],params(1:p,2),'LineWidth',1,'Color',[1,.6,0]), ...
set(gcf,'name','results','numbertitle','off'), ...
xlabel('iteration'),ylabel('')
plot([x,0],params(1:p,3),'LineWidth',1,'Color',[.9,.8,0]), ...
set(gcf,'name','results','numbertitle','off'), ...
xlabel('iteration'),ylabel('')
plot([x,0],params(1:p,4),'LineWidth',1,'Color',[0,.8,0]), ...
set(gcf,'name','results','numbertitle','off'), ...
xlabel('iteration'),ylabel('')
plot([x,0],params(1:p,5),'LineWidth',1,'Color',[.2,.2,1]), ...
set(gcf,'name','results','numbertitle','off'), ...
xlabel('iteration'),ylabel('')
plot([x,0],params(1:p,6),'LineWidth',1,'Color',[.7,.3,.7]), ...
set(gcf,'name','results','numbertitle','off'), ...

```

```

xlabel('iteration'),ylabel('')
hold off
legend('\it{w}_{1}','\theta_{1}','\it{w}_{2}', ...
       '\theta_{2}','\it{w}_{3}','\theta_{3}')
title('parameter evolution')
subplot(2,3,2)
image(abs(A1),'CDataMapping','scaled'), ...
colormap(jet),colorbar,title(colorbar,''), ...
set(gcf,'name','results','numbertitle','off'), ...
xlabel('\it{x} (px)'),ylabel('\it{y} (px)')
pbaspect([1,1,1])
title('|A_{1}|')
subplot(2,3,3)
image(atan2(imag(A1),real(A1)),'CDataMapping','scaled'), ...
colormap(jet),colorbar,title(colorbar,''), ...
set(gcf,'name','results','numbertitle','off'), ...
xlabel('\it{x} (px)'),ylabel('\it{y} (px)')
pbaspect([1,1,1])
title('arg(A_{1})')
subplot(2,3,5)
image(abs(A2),'CDataMapping','scaled'), ...
colormap(jet),colorbar,title(colorbar,''), ...
set(gcf,'name','results','numbertitle','off'), ...
xlabel('\it{x} (px)'),ylabel('\it{y} (px)')
pbaspect([1,1,1])
title('|A_{2}|')
subplot(2,3,6)
image(atan2(imag(A2),real(A2)),'CDataMapping','scaled'), ...
colormap(jet),colorbar,title(colorbar,''), ...
set(gcf,'name','results','numbertitle','off'), ...
xlabel('\it{x} (px)'),ylabel('\it{y} (px)')
pbaspect([1,1,1])
title('arg(A_{2})')

toc;

% export data
dlmwrite('params.txt',params(1:p,:), 'delimiter',' ');
dlmwrite('error.txt',errorsiter,'delimiter',' ');

```

The code makes use of custom-written functions called `sepperror` and `seppgrad` for the calculation of the error and the gradient descent algorithm, respectively. These are reproduced below.

```

function sgrad=sepgrad(w1,q1,w2,q2,w3,q3)
    global rows cols shearx sheary sigmar sigmap D1 D2 delta
    sgrad=[seperror(w1+delta/2,q1,w2,q2,w3,q3) ...
           -seperror(w1-delta/2,q1,w2,q2,w3,q3), ...
           seperror(w1,q1+delta/2,w2,q2,w3,q3) ...
           -seperror(w1,q1-delta/2,w2,q2,w3,q3), ...
           seperror(w1,q1,w2+delta/2,q2,w3,q3) ...
           -seperror(w1,q1,w2-delta/2,q2,w3,q3), ...
           seperror(w1,q1,w2,q2+delta/2,w3,q3) ...
           -seperror(w1,q1,w2,q2-delta/2,w3,q3), ...
           seperror(w1,q1,w2,q2,w3+delta/2,q3) ...
           -seperror(w1,q1,w2,q2,w3-delta/2,q3), ...
           seperror(w1,q1,w2,q2,w3,q3+delta/2) ...
           -seperror(w1,q1,w2,q2,w3,q3-delta/2)];
end

function serror=seperror(w1,q1,w2,q2,w3,q3)
    global rows cols shearx sheary sigmar sigmap D1 D2 delta
    c1=w1*exp(1i*q1);
    c2=w2*exp(1i*q2);
    c3=w3*exp(1i*q3);
    A1=D1+c1.*D2;
    A2=c3.*(c2.*D1+D2);
    A=zeros(rows,cols);
    A(1+abs(sheary)/2:rows-abs(sheary)/2, ...
       1+abs(shearx)/2:cols-abs(shearx)/2) ...
       =(A1(1+abs(sheary)/2+sheary/2:rows-abs(sheary)/2+sheary/2, ...
          1+abs(shearx)/2+shearx/2:cols-abs(shearx)/2+shearx/2) ...
          +A2(1+abs(sheary)/2-sheary/2:rows-abs(sheary)/2-sheary/2, ...
              1+abs(shearx)/2-shearx/2:cols-abs(shearx)/2-shearx/2))/2;
    D1rec=zeros(rows,cols);
    D2rec=D1rec;
    D1rec(1+abs(sheary):rows-abs(sheary), ...
           1+abs(shearx):cols-abs(shearx)) ...
           =(c3*A(1+abs(sheary)-sheary/2:rows-abs(sheary)-sheary/2, ...
                  1+abs(shearx)-shearx/2:cols-abs(shearx)-shearx/2) ...
              -c1*A(1+abs(sheary)+sheary/2:rows-abs(sheary)+sheary/2, ...
                    1+abs(shearx)+shearx/2:cols-abs(shearx)+shearx/2)) ...
           /((1-c1*c2)*c3);
    D2rec(1+abs(sheary):rows-abs(sheary), ...
           1+abs(shearx):cols-abs(shearx)) ...
           =(A(1+abs(sheary)+sheary/2:rows-abs(sheary)+sheary/2, ...

```

```

    1+abs (shearx)+shearx/2:cols-abs (shearx)+shearx/2) ...
    -c2*c3*A(1+abs (sheary)-sheary/2:rows-abs (sheary)-sheary/2, ...
        1+abs (shearx)-shearx/2 ...
        :cols-abs (shearx)-shearx/2)/((1-c1*c2)*c3);
D1crop=D1(1+abs (sheary):rows-abs (sheary), ...
    1+abs (shearx):cols-abs (shearx));
D1reccrop=D1rec(1+abs (sheary):rows-abs (sheary), ...
    1+abs (shearx):cols-abs (shearx));
D2crop=D2(1+abs (sheary):rows-abs (sheary), ...
    1+abs (shearx):cols-abs (shearx));
D2reccrop=D2rec(1+abs (sheary):rows-abs (sheary), ...
    1+abs (shearx):cols-abs (shearx));
error1=log(D1crop./D1reccrop);
error2=log(D2crop./D2reccrop);
serror=sum(sum((real(error1)).^2+(real(error2)).^2)/sigmar
    +(imag(error1)).^2+(imag(error2)).^2)/sigmap);
end

```

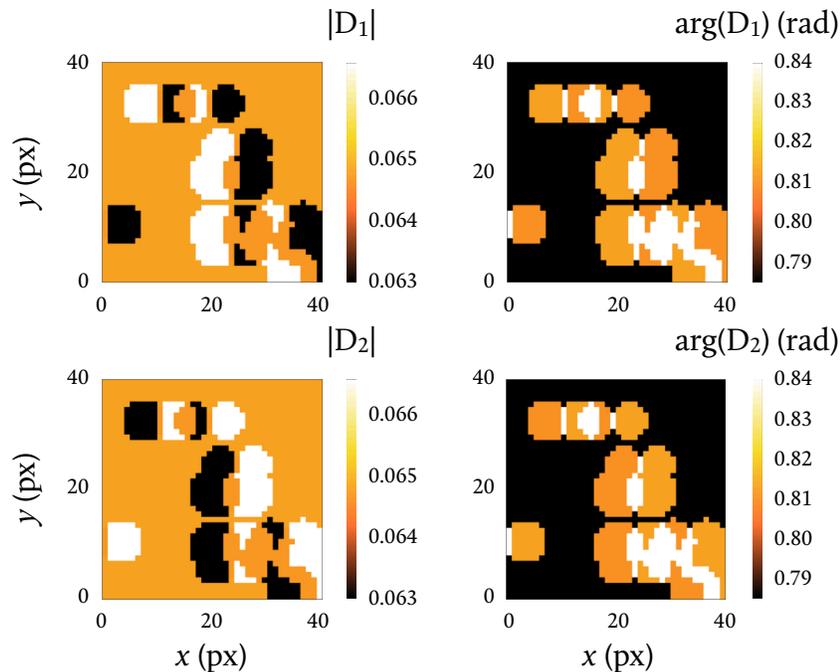


Figure C.1: Simulated amplitude (left) and phase (right) of the co- (top) and cross-polarised (bottom) components of the reflectometry signal that would be produced by another sample with the same characteristics as in appendix B. The shear is horizontal and 6 pixels, as before.

C.2. Results with simulated data

Figure C.1 shows the simulated signal that would be generated by a sample with the same characteristics as the one in the correlation example. As before, the polarisations are mixed, simulating what the Wollaston prism would do; again the combination is

$$\begin{aligned} D_1(\vec{r}) &= \frac{s_H(\vec{r})}{\sqrt{2}} + i \frac{s_V(\vec{r})}{\sqrt{2}}, \\ D_2(\vec{r}) &= i \frac{s_H(\vec{r})}{\sqrt{2}} + \frac{s_V(\vec{r})}{\sqrt{2}}, \end{aligned}$$

so the coefficients we seek are $c_1 = c_2 = \pm i$ and $c_3 = 1$.

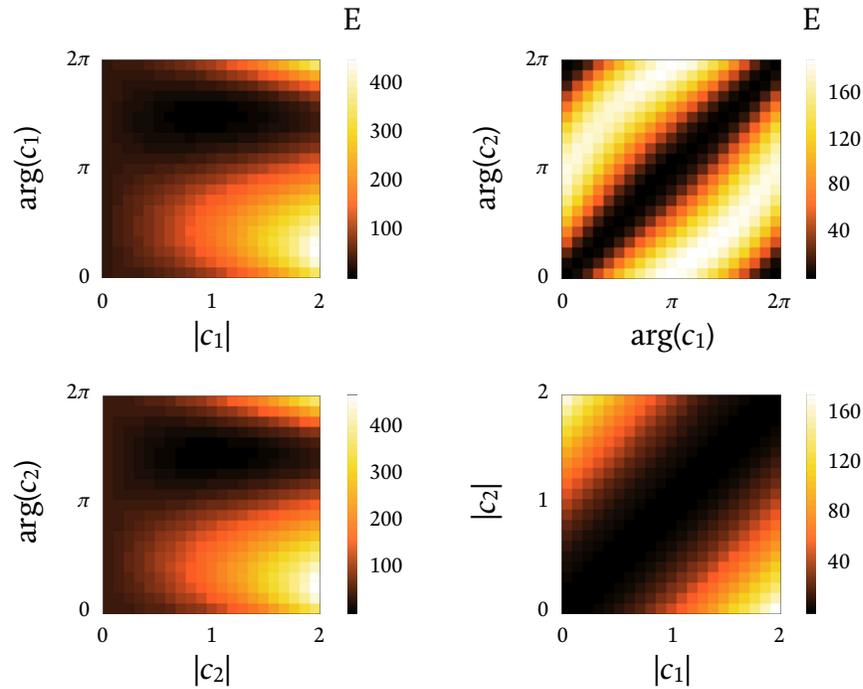


Figure C.2: Several 2-dimensional projections of the correlation at zero displacement of the simulated $(A_1)_f$ and $(A_2)_f$ for $c_3 = 1$. In $(|c_1|, \arg(c_1), |c_2|, \arg(c_2), |c_3|, \arg(c_3))$ space, the projections are as follows:

$$\begin{aligned} \left(|c_1|, \arg(c_1), 1, \frac{3\pi}{2}, 1, 0 \right) & \quad (1, \arg(c_1), 1, \arg(c_2), 1, 0) \\ \left(1, \frac{3\pi}{2}, |c_2|, \arg(c_2), 1, 0 \right) & \quad \left(|c_1|, \frac{3\pi}{2}, |c_2|, \frac{3\pi}{2}, 1, 0 \right) \end{aligned}$$

This time, there does appear to be a minimum at $(1, 3\pi/2, 1, 3\pi/2, 1, 0)$, which corresponds to $(c_1, c_2, c_3) = (-i, -i, 1)$, the point at which the polarisations would be perfectly separated.

Figure C.2 shows several projections of the error E as a function of the c_j . The projections are the same as those shown in figure B.5. It is interesting to note that the structures displayed by E are similar to those displayed by $(A_1)_f \star (A_2)_f$. This time, however, there

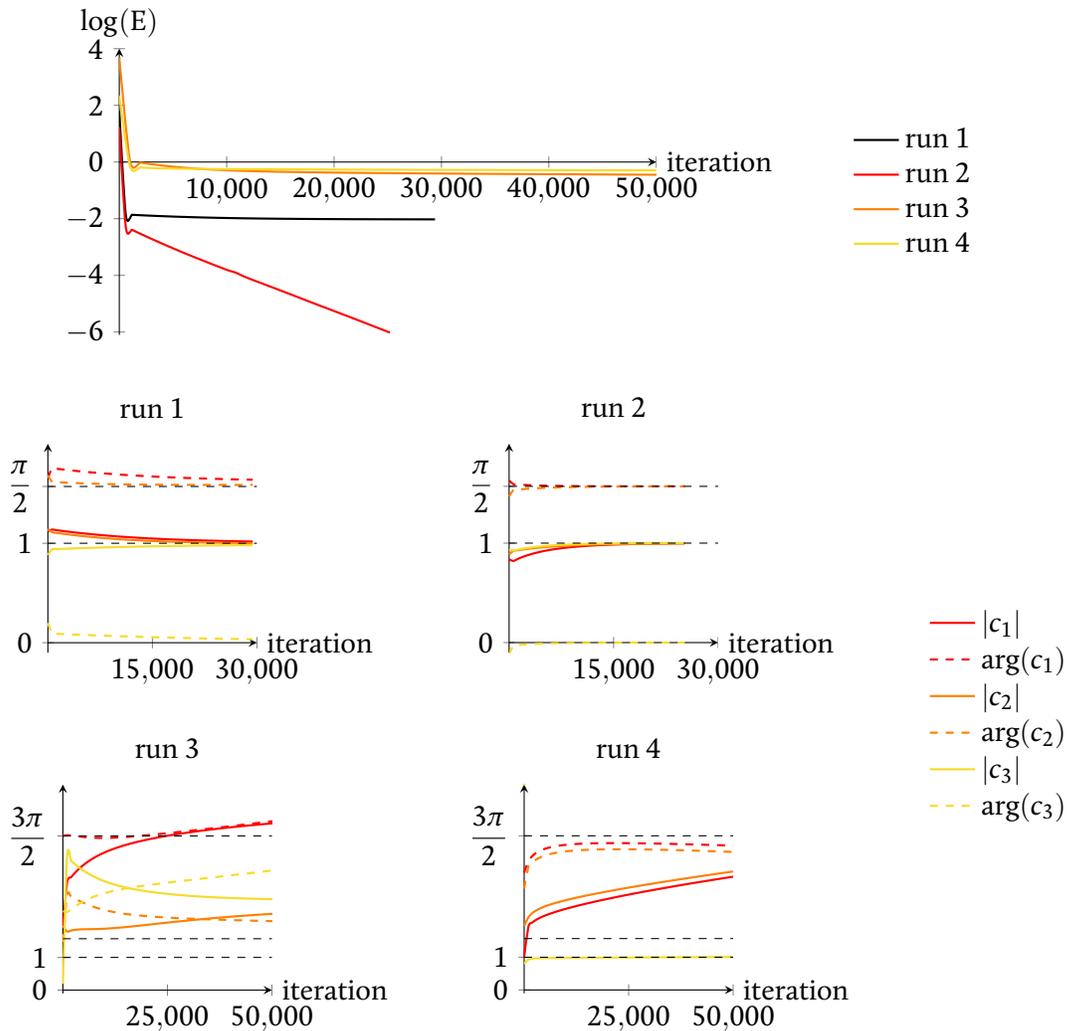


Figure C.3: Evolution of the error E (top) and the parameters (centre and bottom) in four different runs with the same simulated data. The starting points were random; in runs 1 and 2 the starting $|c_j|$ were in the range $[0.8, 1.2]$ and the starting $\arg(c_j)$ were in the range $[\pi/2 - 0.2, \pi/2 + 0.2]$; in runs 3 and 4, the starting $|c_j|$ were in the range $[0, 2]$ and the starting $\arg(c_j)$ could take any value between 0 and 2π . In runs 1 and 2, especially the latter, the algorithm reached a point close enough to the global minimum at $(c_1, c_2, c_3) = (i, i, 1)$ for the error to be extremely small and for separation to be nearly perfect (see figures C.4 and C.5). In runs 3 and 4, it failed to reach a low-error point after 50,000 iterations. The black dashed lines in the centre and bottom graphs mark the values 1, $\pi/2$ and $3\pi/2$.

does appear to be a minimum at $(-i, -i, 1)$, although it is impossible to be certain because E is a function of 6 variables and cannot be visualised in 2 dimensions.

Figure C.3 shows the results of four runs of the gradient descent algorithm on the same data, the one shown in figure C.1. The first two runs had random starting points $(c_1, c_2, c_3)_0$ such that $|c_j| \in [0.8, 1.2]$ and $\arg(c_j) \in [\pi/2 - 0.2, \pi/2 + 0.2]$ for all j ; in these two runs, the algorithm reached the perfect-separation point $(i, i, 1)$ within acceptable error after some 25,000 iterations. The third and fourth runs had random starting points in a much larger region: $|c_j| \in [0, 2]$ and $\arg(c_j) \in [0, 2\pi]$ for all j ; in both of them, the algorithm failed to converge after 50,000 iterations, although, left to run longer, it would have eventually found a minimum, but likely not the global minimum, judging by the fact that $|c_1|$ and $|c_2|$ were moving in the wrong direction in both runs. Letting the algorithm run for 400,000 iterations using the extended region for the possible starting point did not result in convergence; in one instance, for example, the difference between the amplitude of the residual copies of the structures and the amplitude of the no-layer region was about 15% of the no-layer amplitude.

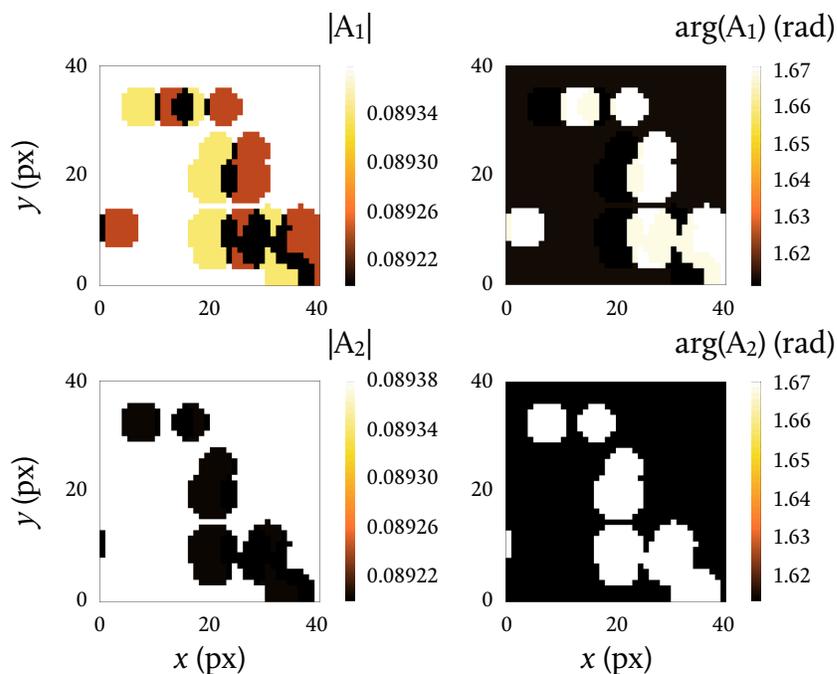


Figure C.4: Final result of the first run shown in figure C.3: amplitude (left) and phase (right) of the A_j calculated using the final parameters of the run with equations 4.5. $|A_1|$ (as well as $\arg(A_1)$ to a lesser extent) shows imperfect separation of the polarisations; however, a closer look reveals the separation is good enough: the amplitude of the residual copy of the wrong polarisation is about 0.09340, while that of the no-layer regions is about 0.09344, causing an error of 3.860×10^{-5} , or about 0.04%.

Figure C.4 shows the amplitude and phase of the A_j calculated using the final parameters of the first run. Although the polarisations are separated in A_2 and almost separated in the phase of A_1 , they appear not to be very well separated in the amplitude of A_1 . However, the residual copies of the features of the wrong polarisation have an amplitude of about 0.09340, which is only 0.04% lower than the no-layer amplitude of 0.09344; the error introduced by the imperfect separation is thus only 0.04%. Figure C.5, which shows the A_j calculated with the final parameters of the second run, shows an even better separation of the polarisations; the residual duplicates of the features have an amplitude of about 0.09165, only 0.009% higher than the no-layer amplitude, making for an error of less than 0.01% in any subsequent analysis.

Unfortunately, in the case of our experimental data it was impossible to separate the polarisations using this algorithm, since a good region in (c_1, c_2, c_3) space for the starting point could not be determined. This suggests that the gradient descent analysis is, in general, unable to separate the polarisations so the noise can be reduced by referencing. As suggested by the results with simulated data presented here, this is likely because the algorithm only converges at the global minimum if it starts close enough to it to avoid local minima.

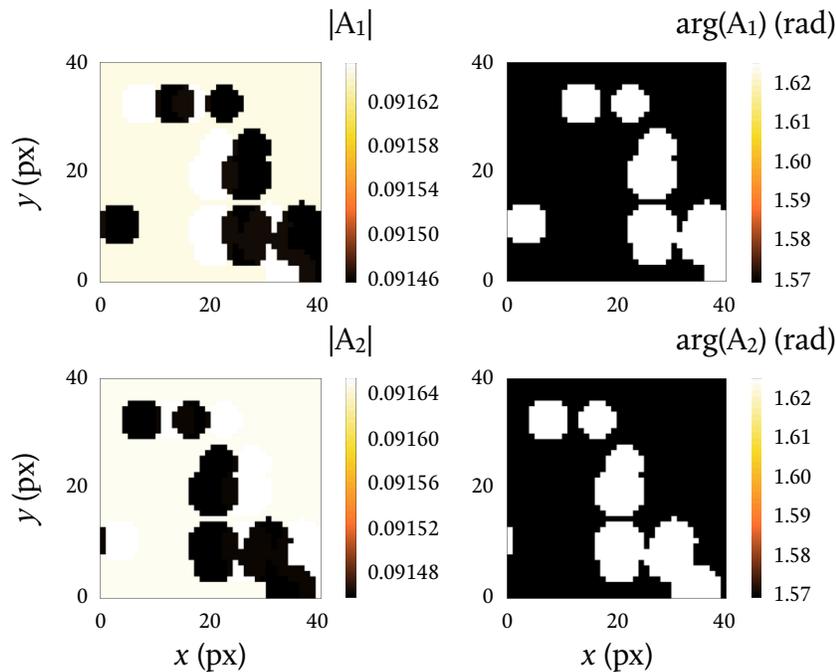


Figure C.5: Final result of the second run shown in figure C.3: amplitude (left) and phase (right) of the A_j . As suggested by the final value of E , the separation is even better than that achieved in the first run. This time, the wrong polarisation in $|A_1|$ has an amplitude only 7.949×10^{-6} , or about 0.009%, larger than the no-layer amplitude.

Appendix D

Genetic algorithm for interferometric reflectometry

D.I. The code

Reproduced below is the code for the comparison programme using a genetic algorithm. It was written and tested on MATLAB 2018b. As in the case of the gradient descent code, it includes a section for the generation of test data because it was tested more extensively with simulated data than with experimental data.

```
% global variables
    global rows cols shearx sheary sigmar sigmap D1 D2

% initialise random number generator with random seed
    rng('shuffle');

% build test data
    disp('Building data...');
    tic;
    rows=40;
    cols=40;
    s=0.06466+0.003595i; % one 4-nm lipid bilayer
    w=0.06489; % no layer
    I1=w*ones(rows,cols);
    I2=I1;
    shearx=8;
    sheary=0;
    num=randi([1,3])+randi([1,3])+randi([1,3])+randi([1,3])+randi([1,3]);
    sizeavg=randi([1,3])+randi([1,2]);
    size=random('Normal',sizeavg,0.5,1,num);
    posx=randi([1,cols],1,num);
```

```

posy=randi([1,rows],1,num);
for m=1:num
    for j=1:rows
        for l=1:cols
            if abs((j-posy(m)+sheary/2)^2+(l-posx(m)+shearx/2)^2) ...
                <=size(m)^2
                I1(j,l)=I1(j,l)+s-w;
            end
            if abs((j-posy(m)-sheary/2)^2+(l-posx(m)-shearx/2)^2) ...
                <=size(m)^2
                I2(j,l)=I2(j,l)+s-w;
            end
        end
    end
end
disp('    Data built.');
```

% add noise to data

```

amp=2000*imag(s); % 2,000 approximately simulates experimentally
                % observed noise level

sigma=imag(s);
cutoff=1/10; % this value has been found to be appropriate for the
            % tested data

noise=amp*random('Normal',0,sigma,1,rows*cols);
noise=fft(noise);
for m=1:rows*cols
    noise(m)=exp(-(m/(2*rows*cols*cutoff)).^2))*noise(m);
end
noise=abs(ifft(noise));
noise=transpose(reshape(noise,cols,rows));
I1=abs(I1).*exp(1i*(atan2(imag(I1),real(I1))+noise));
I2=abs(I2).*exp(1i*(atan2(imag(I2),real(I2))+noise));
disp('    Noise added.');
```

% build "measurements" (D1 & D2) as linear combinations of I1 & I2

```

a1=1/sqrt(2);
a2=1i/sqrt(2);
a3=1i/sqrt(2);
a4=1/sqrt(2);
D1=a1*I1+a2*I2;
D2=a3*I1+a4*I2;
figD=figure('pos',[1,535,560,420]);
```

```

subplot(2,2,1)
    image(abs(D1),'CDataMapping','scaled'), ...
    colormap(jet),colorbar,title(colorbar,''), ...
    set(gcf,'name','data','numbertitle','off'), ...
    xlabel('\it{x}} (px)'),ylabel('\it{y}} (px)')
    pbaspect([1,1,1])
    title('|D_{1}|')
subplot(2,2,2)
    image(atan2(imag(D1),real(D1)),'CDataMapping','scaled'), ...
    colormap(jet),colorbar,title(colorbar,''), ...
    set(gcf,'name','data','numbertitle','off'), ...
    xlabel('\it{x}} (px)'),ylabel('\it{y}} (px)')
    pbaspect([1,1,1])
    title('arg(D_{1})')
subplot(2,2,3)
    image(abs(D2),'CDataMapping','scaled'), ...
    colormap(jet),colorbar,title(colorbar,''), ...
    set(gcf,'name','data','numbertitle','off'), ...
    xlabel('\it{x}} (px)'),ylabel('\it{y}} (px)')
    pbaspect([1,1,1])
    title('|D_{2}|')
subplot(2,2,4)
    image(atan2(imag(D2),real(D2)),'CDataMapping','scaled'), ...
    colormap(jet),colorbar,title(colorbar,''), ...
    set(gcf,'name','data','numbertitle','off'), ...
    xlabel('\it{x}} (px)'),ylabel('\it{y}} (px)')
    pbaspect([1,1,1])
    title('arg(D_{2})')

disp('    Experiment simulated. ');
toc;
disp(' ');

% parameters
sigmar=1;
sigmap=1;
Nchrom=250;
Nxover=6;
Nrepr=1;
mutamp=10^0;
alpha=1/5;
Ngenstop=100000;

```

```

% partition & error matrix
disp('Partitioning space...');
tic;
wsteps=0;
wstart=1;
wend=1;
if wsteps==0
    wstep=0;
else
    wstep=(wend-wstart)/wsteps;
end
qsteps=9;
qstart=0;
qend=2*pi;
if qsteps==0
    qstep=0;
else
    qstep=(qend-qstart)/qsteps;
end
errorgrid=zeros(wsteps+1,qsteps,wsteps+1,qsteps,wsteps+1,qsteps);
for u1=1:wsteps+1
    for v1=1:qsteps
        for u2=1:wsteps+1
            for v2=1:qsteps
                for u3=1:wsteps+1
                    for v3=1:qsteps
                        errorgrid(u1,v1,u2,v2,u3,v3) ...
                            =sepperrorwq(wstart+(u1-1)*wstep, ...
                                qstart+(v1-1)*qstep,
                                wstart+(u2-1)*wstep,
                                qstart+(v2-1)*qstep,
                                wstart+(u3-1)*wstep,
                                qstart+(v3-1)*qstep);
                    end
                end
            end
        end
    end
end
X=['    Building preliminary error matrix (' , ...
    num2str(100*((u1-1)+v1/qsteps)/(wsteps+1)), '% complete) .'];
disp(X);
end
end

```



```

        <=errorgrid(u1,v1,u2,v2,u3,v3-1)) ...
    %% (v3==qsteps) || ...
    (errorgrid(u1,v1,u2,v2,u3,v3) ...
        <=errorgrid(u1,v1,u2,v2,u3,v3+1))
    locmin=transpose([ ...
        transpose(locmin), ...
        transpose([wstart+(u1-1)*wstep, ...
            qstart+(v1-1)*qstep, ...
            wstart+(u2-1)*wstep, ...
            qstart+(v2-1)*qstep, ...
            wstart+(u3-1)*wstep, ...
            qstart+(v3-1)*qstep])]);
    locminerr=[locminerr, ...
        errorgrid(u1,v1,u2,v2,u3,v3)];
    end
    end
    end
    end
    X=[' Finding local minima (' , ...
        num2str(100*((u1-1)+v1/qsteps)/(wsteps+1)), '% complete).'];
    disp(X);
    end
end
Nlocmin=length(locmin);
locminord=zeros([Nlocmin,1]); % vector for ordered local minimum
                                % positions
for j=1:Nlocmin % order local minima by lowest error
    for l=1:Nlocmin
        if locminerr(l)==min(locminerr)
            locminord(j)=l;
            locminerr(l)=max(locminerr)+1; % remove local minimum from
                % local minimum ordering
                % pool
            break;
        end
    end
end
end
toc;
disp(' ');

% minimise error

```



```

    pchrom(j,2)=pchromwq(j,1)*sin(pchromwq(j,2));
    pchrom(j,3)=pchromwq(j,3)*cos(pchromwq(j,4));
    pchrom(j,4)=pchromwq(j,3)*sin(pchromwq(j,4));
    pchrom(j,5)=pchromwq(j,5)*cos(pchromwq(j,6));
    pchrom(j,6)=pchromwq(j,5)*sin(pchromwq(j,6));
end
fitnessp=zeros([Nchrom,1]); % parent fitness vector
for j=1:Nchrom % determine parent fitness
    fitnessp(j)=-log10(seperrorxy(pchrom(j,1),pchrom(j,2), ...
                                pchrom(j,3),pchrom(j,4), ...
                                pchrom(j,5),pchrom(j,6)));
end
cchrom=pchrom; % child vector
fitnessc=fitnessp; % child fitness vector
reprobl=zeros([Nchrom,1]); % parent reproduction probability vector
reprob2=reprobl; % temporary vector for second parent
remove=ones([Nrepr*Noffspr,1]); % indices of non-surviving parents
parents=ones([Nrepr,2]); % parents in reproduction events
iter=1;
iterlast=1; % last iteration with improvement
error0=10^-max(fitnessp);
errors=0;
mutations=0;
X=['    Initial error = ',num2str(error0)];
disp(X);
X=['    Target error = ',num2str(seperrorxy(0,1,0,1,1,0)),'.'];
disp(X);
stop=0;
while stop==0
    if iter>1
        errors=[errors,0];
        mutations=[mutations,mutations(iter-1)];
    end
    cchrom=pchrom; % clone parent generation to child generation
    fitnessc=fitnessp; % inherit fitnesses
    remprob=(max(fitnessp)-fitnessc)./sum(max(fitnessp)-fitnessc);
    % calculate parent removal probabilities
    for j=1:Nrepr*Noffspr % determine removed parent in each
        % reproduction event
        u=random('Uniform',0,sum(remprob));
        for l=1:Nchrom
            if u>sum(remprob(1:l))

```

```

        remove(j)=1+1;
    end
end
remprob(remove(j))=0;
end
repprob1=(fitnessp-min(fitnessp))./sum(fitnessp-min(fitnessp));
% calculate parent reproduction probabilities
for j=1:Nrepr % determine first parent in each reproduction event
    for l=1:Nchrom
        u=random('Uniform',0,1);
        if u>sum(repprob1(1:l))
            parents(j,1)=1+1;
        end
    end
end
repprob2=repprob1; % second parent reproduction probabilities
for j=1:Nrepr % determine second parent in each reproduction event
    repprob2(parents(j,1))=0; % remove each reproduction event's
        % first parent from the reproduction
        % pool for that event
    u=random('Uniform',0,sum(repprob2));
    for l=1:Nchrom
        if u>sum(repprob2(1:l))
            parents(j,2)=1+1;
        end
    end
    repprob2=repprob1; % reintroduce first parent into reproduction
        % pool for future reproduction events
end
for j=1:Nrepr % generate children & calculate their fitnesses
    xover=randperm(6,6); % determine crossover traits
    for l=1:Noffspr
        for m=1:Nxover % cross parent traits over
            beta=random('Uniform',0,1);
            cchrom(remove((j-1)*Noffspr+1),xover(m)) ...
                =beta*pchrom(parents(j,1),xover(m)) ...
                +(1-beta)*pchrom(parents(j,2),xover(m));
        end
        for m=Nxover+1:6 % inherit non-crossover traits from
            % parents
            cchrom(remove((j-1)*Noffspr+1),xover(m)) ...
                =pchrom(parents(j,1.5+0.5*(-1)^ceil(1/2^(6-m))), ...

```

```

                                xover(m) ;
    end
    fitnessc(remove((j-1)*Noffspr+1)) ...
        =seperorxy(cchrom(remove((j-1)*Noffspr+1),1), ...
                    cchrom(remove((j-1)*Noffspr+1),2), ...
                    cchrom(remove((j-1)*Noffspr+1),3), ...
                    cchrom(remove((j-1)*Noffspr+1),4), ...
                    cchrom(remove((j-1)*Noffspr+1),5), ...
                    cchrom(remove((j-1)*Noffspr+1),6));
    end
end
for j=1:Nrepr*Noffspr % introduce mutations & calculate fitnesses
    % of new children
    mutationwq=[random('Uniform',0,5),random('Uniform',0,2*pi), ...
                random('Uniform',0,5),random('Uniform',0,2*pi), ...
                random('Uniform',0,5),random('Uniform',0,2*pi)];
    mutation=[mutationwq(1)*cos(mutationwq(2)), ...
              mutationwq(1)*sin(mutationwq(2)), ...
              mutationwq(3)*cos(mutationwq(4)), ...
              mutationwq(3)*sin(mutationwq(4)), ...
              mutationwq(5)*cos(mutationwq(6)), ...
              mutationwq(5)*sin(mutationwq(6))];
    dist=norm(mutation-cchrom(remove(j),:));
    mutprob=mutamp*exp(-alpha*dist);
    mutparam=random('Uniform',0,1);
    if mutparam<mutprob
        cchrom(remove(j),:)=mutation;
        if iter>1
            mutations(iter)=mutations(iter-1)+1;
        else
            mutations(iter)=1;
        end
    end
end
fitnessc(remove(j))=-log10(seperorxy(cchrom(remove(j),1), ...
                                     cchrom(remove(j),2), ...
                                     cchrom(remove(j),3), ...
                                     cchrom(remove(j),4), ...
                                     cchrom(remove(j),5), ...
                                     cchrom(remove(j),6)));
end
pchrom=cchrom; % turn child generation into new parent generation
fitnessp=fitnessc; % back-inherit fitnesses

```

```

errors(iter)=10^-max(fitnessc); % determine best error of new
                                % generation
if iter>1 % increase mutation likelihood when error becomes small
    if errors(iter)<1&&errors(iter-1)>=1
        alpha=alpha/2;
    end
end
if iter==1||iter==iterlast+10000|errors(iter)<errlast
    X=['    Iteration ',num2str(iter),' error = ', ...
        num2str(errors(iter)),'.'];
    disp(X);
    iterlast=iter;
    errlast=errors(iter);
end
if iter>Ngenstop % stop run if target number of generations without
                % improvement has been reached
    if errors(iter)==errors(iter-Ngenstop)
        stop=1;
    end
end
for l=1:Nchrom
    if fitnessc(l)==max(fitnessc)
        bestchrom=cchrom(l,:);
    end
end
iter=iter+1;
end
errors=[error0,errors];
X=['    After ',num2str(iter-1), ...
    ' iterations, the error has gone from ',num2str(errors(1)), ...
    ' to ',num2str(errors(iter)),'.'];
disp(X);
X=['    Target error = ',num2str(seperrorxy(0,1,0,1,1,0)),'.'];
disp(X);
X=['    Parameters obtained: (c1,c2,c3) = (' ,num2str(bestchrom(1))];
if bestchrom(2)>=0
    X=[X,'+'];
end
X=[X,num2str(bestchrom(2)),'i',num2str(bestchrom(3))];
if bestchrom(4)>=0
    X=[X,'+'];
end

```

```

X=[X,num2str(bestchrom(4)),'i',num2str(bestchrom(5))];
if bestchrom(6)>=0
    X=[X,+'];
end
X=[X,num2str(bestchrom(6)),'i.'];
disp(X);

% display results
xe=0:iter-1; % x-axis range for error graph
xm=1:iter-1; % x-axis range for mutation graph
c1=bestchrom(1)+1i*bestchrom(2);
c2=bestchrom(3)+1i*bestchrom(4);
c3=bestchrom(5)+1i*bestchrom(6);
A1=D1+c1.*D2;
A2=c3.*(c2.*D1+D2);
figA=figure('pos',[1,1,891,420]);
    subplot(2,3,1)
        plot(xe,log10(errors(1:iter)),'LineWidth',1,'Color',[1,0,0]), ...
            set(gcf,'name','Results','numbertitle','off'), ...
            xlabel('iteration'),ylabel('log_{10}(error)')
            title('Error evolution')
    subplot(2,3,4)
        plot(xm,mutations(1:iter-1),
            'LineWidth',1,'Color',[0,0.8,0]), ...
            set(gcf,'name','Results','numbertitle','off'), ...
            xlabel('iteration'),ylabel('total mutations')
            title('Mutations')
    subplot(2,3,2)
        image(abs(A1),'CDataMapping','scaled'), ...
        colormap(jet),colorbar,title(colorbar,''), ...
        set(gcf,'name','Results','numbertitle','off'), ...
        xlabel('\it{x} (px)'),ylabel('\it{y} (px)')
        pbaspect([1,1,1])
        title('|A_{1}|')
    subplot(2,3,3)
        image(atan2(imag(A1),real(A1)),'CDataMapping','scaled'), ...
        colormap(jet),colorbar,title(colorbar,''), ...
        set(gcf,'name','Results','numbertitle','off'), ...
        xlabel('\it{x} (px)'),ylabel('\it{y} (px)')
        pbaspect([1,1,1])
        title('arg(A_{1})')
    subplot(2,3,5)

```

```

        image(abs(A2), 'CDataMapping', 'scaled'), ...
        colormap(jet), colorbar, title(colorbar, ''), ...
        set(gcf, 'name', 'Results', 'numbertitle', 'off'), ...
        xlabel('\it{x} (px)'), ylabel('\it{y} (px)')
        pbaspect([1,1,1])
        title('|A_{2}|')
    subplot(2,3,6)
        image(atan2(imag(A2), real(A2)), 'CDataMapping', 'scaled'), ...
        colormap(jet), colorbar, title(colorbar, ''), ...
        set(gcf, 'name', 'Results', 'numbertitle', 'off'), ...
        xlabel('\it{x} (px)'), ylabel('\it{y} (px)')
        pbaspect([1,1,1])
        title('arg(A_{2})')

    toc;
    disp(' ');

% export data
    disp('Exporting results...');
    tic;
    dlmwrite('error.txt', errors, 'delimiter', ' ');
    disp('    Error progression exported as error.txt. ');
    dlmwrite('mutations.txt', mutations, 'delimiter', ' ');
    disp('    Mutation counter exported as mutations.txt. ');
    toc;
    disp(' ');

```

The code makes use of a custom-written function called `seperorxy` for the calculation of the error. It is very similar to the `seperor` function from the gradient descent algorithm but takes input in cartesian coordinates instead of polar coordinates. It is reproduced below.

```

function seporxy=seperorxy(x1,y1,x2,y2,x3,y3)
    global rows cols shearx sheary sigmar sigmap D1 D2
    c1=x1+1i*y1;
    c2=x2+1i*y2;
    c3=x3+1i*y3;
    A1=D1+c1.*D2;
    A2=c3.*(c2.*D1+D2);
    A=zeros(rows,cols);
    A(1+abs(sheary)/2:rows-abs(sheary)/2, ...
    1+abs(shearx)/2:cols-abs(shearx)/2) ...
    =(A1(1+abs(sheary)/2+sheary/2:rows-abs(sheary)/2+sheary/2, ...
    1+abs(shearx)/2+shearx/2:cols-abs(shearx)/2+shearx/2) ...

```

```

        +A2 (1+abs (sheary) /2-sheary/2:rows-abs (sheary) /2-sheary/2, ...
            1+abs (shearx) /2-shearx/2:cols-abs (shearx) /2-shearx/2) /2;
D1rec=zeros (rows,cols) ;
D2rec=D1rec;
D1rec (1+abs (sheary) :rows-abs (sheary) , ...
        1+abs (shearx) :cols-abs (shearx) ) ...
    =(c3*A (1+abs (sheary) -sheary/2:rows-abs (sheary) -sheary/2, ...
        1+abs (shearx) -shearx/2:cols-abs (shearx) -shearx/2) ...
        -c1*A (1+abs (sheary) +sheary/2:rows-abs (sheary) +sheary/2, ...
        1+abs (shearx) +shearx/2:cols-abs (shearx) +shearx/2) ) ...
        /((1-c1*c2) *c3) ;
D2rec (1+abs (sheary) :rows-abs (sheary) , ...
        1+abs (shearx) :cols-abs (shearx) ) ...
    =(A (1+abs (sheary) +sheary/2:rows-abs (sheary) +sheary/2, ...
        1+abs (shearx) +shearx/2:cols-abs (shearx) +shearx/2) ...
        -c2*c3*A (1+abs (sheary) -sheary/2:rows-abs (sheary) -sheary/2, ...
        1+abs (shearx) -shearx/2 ...
        :cols-abs (shearx) -shearx/2) ) ...
        /((1-c1*c2) *c3) ;
D1crop=D1 (1+abs (sheary) :rows-abs (sheary) , ...
            1+abs (shearx) :cols-abs (shearx) ) ;
D1reccrop=D1rec (1+abs (sheary) :rows-abs (sheary) , ...
            1+abs (shearx) :cols-abs (shearx) ) ;
D2crop=D2 (1+abs (sheary) :rows-abs (sheary) , ...
            1+abs (shearx) :cols-abs (shearx) ) ;
D2reccrop=D2rec (1+abs (sheary) :rows-abs (sheary) , ...
            1+abs (shearx) :cols-abs (shearx) ) ;
error1=log (D1crop ./D1reccrop) ;
error2=log (D2crop ./D2reccrop) ;
serrorxy=sum (sum ((real (error1) ) .^2+(real (error2) ) .^2) /sigmar
                +((imag (error1) ) .^2+(imag (error2) ) .^2) /sigmap) ) ;
end

```

An alternative function to calculate the error, `sepperrorwq`, was also used in order to see whether the results of the algorithm depended on the coordinate system used. This function is identical to `sepperror` from the gradient descent algorithm, but it is reproduced below for completeness.

```

function sepperrorwq=sepperrorwq(w1,q1,w2,q2,w3,q3)
    global rows cols shearx sheary sigmar sigmap D1 D2
    c1=w1*exp (1i*q1) ;
    c2=w2*exp (1i*q2) ;
    c3=w3*exp (1i*q3) ;

```

```

A1=D1+c1.*D2;
A2=c3.*(c2.*D1+D2);
A=zeros(rows,cols);
A(1+abs(sheary)/2:rows-abs(sheary)/2, ...
  1+abs(shearx)/2:cols-abs(shearx)/2) ...
  =(A1(1+abs(sheary)/2+sheary/2:rows-abs(sheary)/2+sheary/2, ...
    1+abs(shearx)/2+shearx/2:cols-abs(shearx)/2+shearx/2) ...
    +A2(1+abs(sheary)/2-sheary/2:rows-abs(sheary)/2-sheary/2, ...
      1+abs(shearx)/2-shearx/2:cols-abs(shearx)/2-shearx/2))/2;
D1rec=zeros(rows,cols);
D2rec=D1rec;
D1rec(1+abs(sheary):rows-abs(sheary), ...
  1+abs(shearx):cols-abs(shearx)) ...
  =(c3*A(1+abs(sheary)-sheary/2:rows-abs(sheary)-sheary/2, ...
    1+abs(shearx)-shearx/2:cols-abs(shearx)-shearx/2) ...
    -c1*A(1+abs(sheary)+sheary/2:rows-abs(sheary)+sheary/2, ...
      1+abs(shearx)+shearx/2:cols-abs(shearx)+shearx/2)) ...
  /((1-c1*c2)*c3);
D2rec(1+abs(sheary):rows-abs(sheary), ...
  1+abs(shearx):cols-abs(shearx)) ...
  =(A(1+abs(sheary)+sheary/2:rows-abs(sheary)+sheary/2, ...
    1+abs(shearx)+shearx/2:cols-abs(shearx)+shearx/2) ...
    -c2*c3*A(1+abs(sheary)-sheary/2:rows-abs(sheary)-sheary/2, ...
      1+abs(shearx)-shearx/2 ...
      :cols-abs(shearx)-shearx/2))/((1-c1*c2)*c3);
D1crop=D1(1+abs(sheary):rows-abs(sheary), ...
  1+abs(shearx):cols-abs(shearx));
D1reccrop=D1rec(1+abs(sheary):rows-abs(sheary), ...
  1+abs(shearx):cols-abs(shearx));
D2crop=D2(1+abs(sheary):rows-abs(sheary), ...
  1+abs(shearx):cols-abs(shearx));
D2reccrop=D2rec(1+abs(sheary):rows-abs(sheary), ...
  1+abs(shearx):cols-abs(shearx));
error1=log(D1crop./D1reccrop);
error2=log(D2crop./D2reccrop);
serrorwq=sum(sum(((real(error1)).^2+(real(error2)).^2)/sigmar
  +((imag(error1)).^2+(imag(error2)).^2)/sigmap));
end

```

D.2. Results with experimental data

Figure D.1 shows data similar to that shown previously, but with added phase noise with amplitude $a = 10,000 \text{ Im}(s)$ to more accurately simulate experimental data. The error E is similar to what is shown in figure 4.21; in the case of this particular data set, a total of 2,771 local minima were found with a $11 \times 11 \times 11 \times 11 \times 11$ partition of the $(|c_1|, \arg(c_1), |c_2|, \arg(c_2), |c_3|, \arg(c_3))$ space.

Figure D.2 shows the evolution of the error E over 7 different runs. A run typically lasted between 110,000 and 150,000 iterations. Of the runs shown in the figure, run 1 resulted in the lowest error. The point at which this error was achieved is $(0.139 + 0.967i, 0.143 + 0.942i, 0.984 + 0.006i)$, which is very close to $(i, i, 1)$; the cartesian distance between the two points is only about 0.211. Figure D.3 shows the A_j calculated using the aforementioned point with equations 4.5. The images do not appear to be properly separated, but the amplitude of the error (i.e. the amplitude of the residual copies of the structures due to imperfect separation) is only about 2.5% for A_1 and 0.09% for A_2 , so separation was achieved within reasonable error.

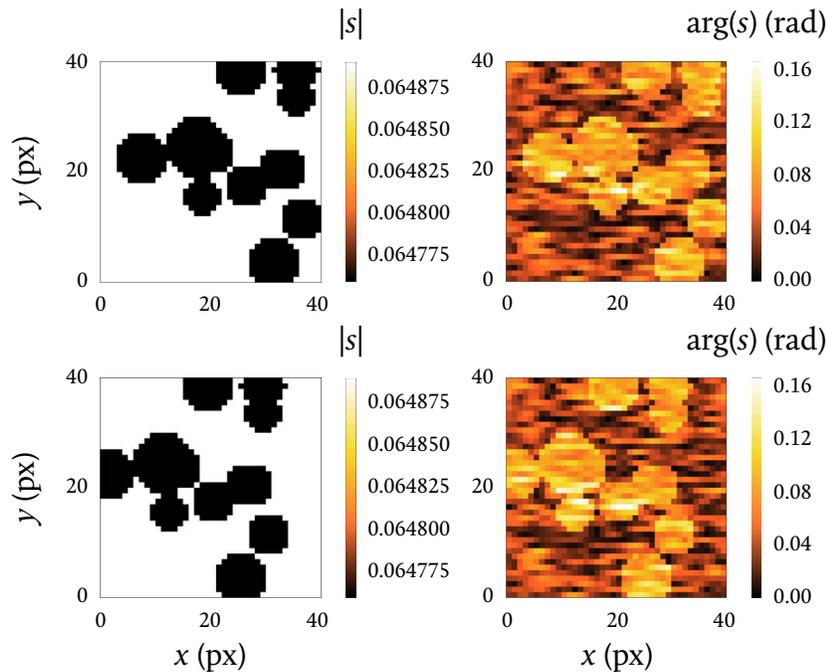


Figure D.1: Simulated amplitude (left) and phase (right) of the reflectometry signal that would be produced by another sample with the same characteristics as the previous simulated samples. As before, the co- (top) and cross-polarised (bottom) components have mixed polarisation information and the shear is 6 pixels in the horizontal direction. This time, however, the phase has noise of amplitude $a = 10,000 \text{ Im}(s)$.

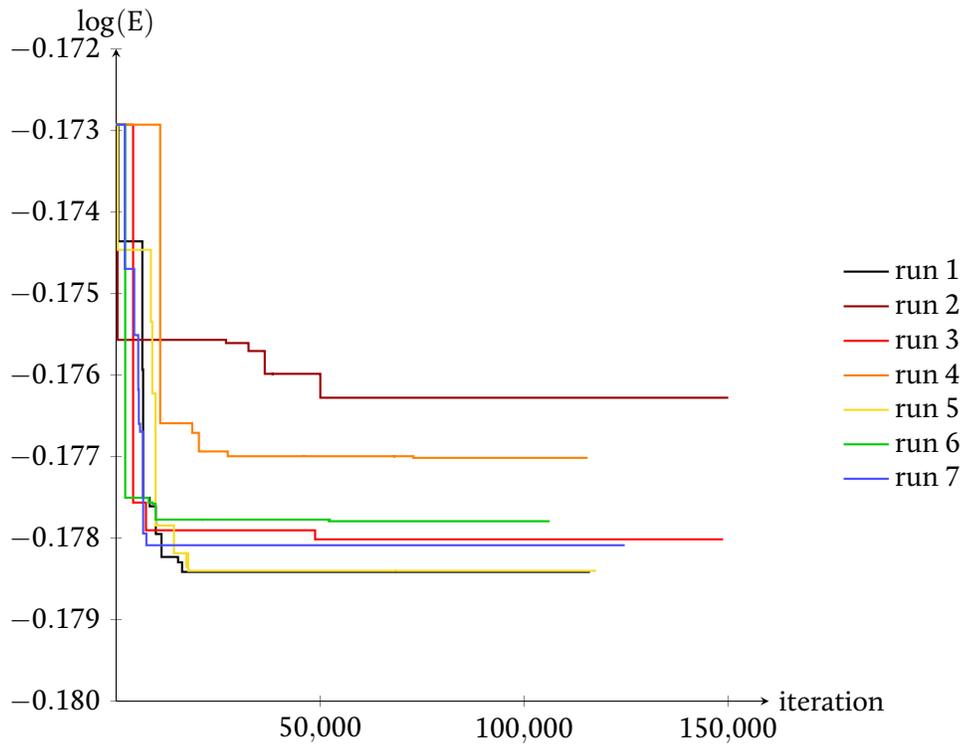


Figure D.2: Evolution of the error E in seven different runs with the same simulated data. The initial population, which consisted of the 250 local minima of E with the lowest value of E (which, notably, did not include the points $(c_1, c_2, c_3) = (\pm i, \pm i, 1)$, likely due to the noise in the data increasing the error of these points slightly more than for other points), was identical in every run, but the random nature of reproduction and mutation caused every run to be different.

Choosing an initial population entirely at random (with, as before, $|c_j| \in [0, 2]$ and $\arg(c_j) \in [0, 2\pi]$ for every j) resulted in runs which took longer (typically around 220,000 iterations, but sometimes considerably fewer and occasionally considerably more) and achieved errors between 0.5% and 2%.

It is not, at present, known why the genetic algorithm was unable to separate the polarisations in the case of experimental data. It is possible that the large amount of (complex) noise in the data caused the point of perfect separation to have a larger error than other points where the separation was imperfect; this would have resulted in the algorithm favouring these other points over the correct one, so even if the correct point had been reached in any generation it would have had a lower chance of surviving for many generations and would have eventually been lost. Without knowing which point corresponds to perfect separation, it is impossible to tell whether this was indeed the case. It is evident from the example shown in this appendix that this is not the case for noisy simulated data, or at least the effect of noise is not pronounced enough for the algorithm to favour points where separation is not achieved to within acceptable error.

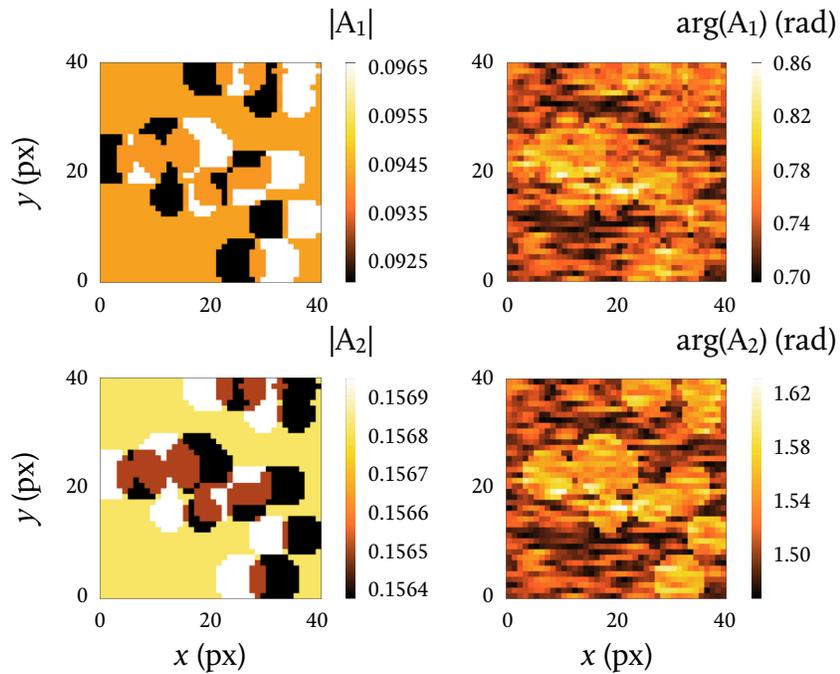


Figure D.3: Amplitude (left) and phase (right) of the linear combinations A_j of the simulated signal components using the c_j obtained from the first of the runs shown in figure D.2. The amplitude of the no-sample region is about 0.094251 for A_1 and about 0.15682 for A_2 , while that of the residual duplicates of the structures is about 0.096592 (an error of about 2.5%) and 0.15695 (an error of about 0.09%), respectively.

Appendix E

Line trace analysis for interferometric reflectometry

Reproduced below is the code for the reflectometry line trace analysis. It was written and tested on Mathematica 11. It is split into three parts: partitioning of part of the (d, n_s) space and calculation of s for every point in the partition, data & parameter reading, and line trace analysis.

```
(* clear all variables *)
Clear[Evaluate[Context[]<>"*"]];

(* set export directory to notebook directory *)
SetDirectory[NotebookDirectory[]];

(* calculation parameters *)
nw=1.333;
ng=1.518;
lambda=550;
k=2*Pi/lambda;
NA=1.27;
ff=1;
Deltad=0.1;
dmax=30;
Deltan=0.001;

(* build s(d,n) table *)
s=ParallelTable[NIntegrate[Exp[-(ff^2)*(Sin[theta]/(NA/ng))^2]
*Cos[theta]*Sin[theta]*(1/2)
*((ng*Cos[theta]
-n*Cos[ArcSin[(n/ng)*Sin[theta]])]
/(ng*Cos[theta]
```

$$\begin{aligned}
& +n*\text{Cos}[\text{ArcSin}[(n/ng)*\text{Sin}[\text{theta}]]]) \\
+ & ((2*ng*\text{Cos}[\text{theta}] \\
& / (ng*\text{Cos}[\text{theta}] \\
& +n*\text{Cos}[\text{ArcSin}[(n/ng)*\text{Sin}[\text{theta}]]])) \\
& * (2*n*\text{Cos}[\text{ArcSin}[(n/ng)*\text{Sin}[\text{theta}]]] \\
& / (n*\text{Cos}[\text{ArcSin}[(n/ng)*\text{Sin}[\text{theta}]]] \\
& +ng*\text{Cos}[\text{theta}])) \\
& * ((n*\text{Cos}[\text{ArcSin}[(n/ng)*\text{Sin}[\text{theta}]]] \\
& -nw*\text{Cos}[\text{ArcSin}[(nw/ng)*\text{Sin}[\text{theta}]]]) \\
& / (n*\text{Cos}[\text{ArcSin}[(n/ng)*\text{Sin}[\text{theta}]]] \\
& +nw*\text{Cos}[\text{ArcSin}[(nw/ng)*\text{Sin}[\text{theta}]]])) \\
& * \text{Exp}[2*I*d*n*k \\
& / \text{Cos}[\text{ArcSin}[(n/ng)*\text{Sin}[\text{theta}]]]]) \\
/ & (1 - ((n*\text{Cos}[\text{ArcSin}[(n/ng)*\text{Sin}[\text{theta}]]] \\
& -nw*\text{Cos}[\text{ArcSin}[(nw/ng)*\text{Sin}[\text{theta}]]]) \\
& * (n*\text{Cos}[\text{ArcSin}[(n/ng)*\text{Sin}[\text{theta}]]] \\
& +nw*\text{Cos}[\text{ArcSin}[(nw/ng) \\
& \quad \quad \quad * \text{Sin}[\text{theta}]]])) \\
& * ((n*\text{Cos}[\text{ArcSin}[(n/ng)*\text{Sin}[\text{theta}]]] \\
& -ng*\text{Cos}[\text{theta}]) \\
& / (n*\text{Cos}[\text{ArcSin}[(n/ng)*\text{Sin}[\text{theta}]]] \\
& +ng*\text{Cos}[\text{theta}])) \\
& * \text{Exp}[2*I*d*n*k \\
& / \text{Cos}[\text{ArcSin}[(n/ng)*\text{Sin}[\text{theta}]]]]) \\
+ & (n*\text{Cos}[\text{theta}] \\
& -ng*\text{Cos}[\text{ArcSin}[(n/ng)*\text{Sin}[\text{theta}]]]) \\
/ & (n*\text{Cos}[\text{theta}] \\
& +ng*\text{Cos}[\text{ArcSin}[(n/ng)*\text{Sin}[\text{theta}]]]) \\
+ & ((2*ng*\text{Cos}[\text{theta}] \\
& / (n*\text{Cos}[\text{theta}] \\
& +ng*\text{Cos}[\text{ArcSin}[(n/ng)*\text{Sin}[\text{theta}]]])) \\
& * (2*n*\text{Cos}[\text{ArcSin}[(n/ng)*\text{Sin}[\text{theta}]]] \\
& / (ng*\text{Cos}[\text{ArcSin}[(n/ng)*\text{Sin}[\text{theta}]]] \\
& +n*\text{Cos}[\text{theta}])) \\
& * ((nw*\text{Cos}[\text{ArcSin}[(n/ng)*\text{Sin}[\text{theta}]]] \\
& -n*\text{Cos}[\text{ArcSin}[(nw/ng)*\text{Sin}[\text{theta}]]]) \\
& / (nw*\text{Cos}[\text{ArcSin}[(n/ng)*\text{Sin}[\text{theta}]]] \\
& +n*\text{Cos}[\text{ArcSin}[(nw/ng)*\text{Sin}[\text{theta}]]])) \\
& * \text{Exp}[2*I*d*k*n \\
& / \text{Cos}[\text{ArcSin}[(n/ng)*\text{Sin}[\text{theta}]]]]) \\
/ & (1 - ((nw*\text{Cos}[\text{ArcSin}[(n/ng)*\text{Sin}[\text{theta}]]] \\
& -n*\text{Cos}[\text{ArcSin}[(nw/ng)*\text{Sin}[\text{theta}]]])
\end{aligned}$$

```

        / (nw*cos[ArcSin[(n/ng)*Sin[theta]]]
          +n*cos[ArcSin[(nw/ng)
                    *Sin[theta]]])
      * ((ng*cos[ArcSin[(n/ng)*Sin[theta]]]
        -n*cos[theta])
        / (ng*cos[ArcSin[(n/ng)*Sin[theta]]]
          +n*cos[theta]))
      *Exp[2*I*d*n*k
          /Cos[ArcSin[(n/ng)
                  *Sin[theta]]]),
      {theta,0,ArcSin[NA/ng]]
/NIntegrate[Exp[-(ff^2)*(Sin[theta]/(NA/ng))^2]
            *Cos[theta]*Sin[theta]*(1/2)
            *( (ng*cos[theta]
                -nw*cos[ArcSin[(ng/ng)*Sin[theta]]]
                / (ng*cos[theta]
                  +nw*cos[ArcSin[(nw/ng)*Sin[theta]]]
                  + (nw*cos[theta]
                    -ng*cos[ArcSin[(nw/ng)*Sin[theta]]]
                    / (nw*cos[theta]
                      +ng*cos[ArcSin[(nw/ng)*Sin[theta]]])
                  ),
            {theta,0,ArcSin[NA/ng]}],
      {n,nw,ng,Deltan},{d,0,dmax,Deltad}];

(* data files *)
(* the parameter files contain the coordinates, lengths and directions
   (+x or -x) of all line traces for a given image *)
(* file paths below are incomplete *)
data={".../Data/20180127_iRef/pf/003_B2x.dat",
      ".../Data/20180127_iRef/pf/003_B2y.dat",
      ".../Data/20180127_iRef/pf/003_params_v15.dat"},
{".../Data/20180127_iRef/pf/011_B2x.dat",
      ".../Data/20180127_iRef/pf/011_B2y.dat",
      ".../Data/20180127_iRef/pf/011_params_v15.dat"},
{".../Data/20180127_iRef/pf/019_B2x.dat",
      ".../Data/20180127_iRef/pf/019_B2y.dat",
      ".../Data/20180127_iRef/pf/019_params_v15.dat"},
{".../Data/20180127_iRef/pf/027_B2x.dat",
      ".../Data/20180127_iRef/pf/027_B2y.dat",
      ".../Data/20180127_iRef/pf/027_params_v15.dat"},
{".../Data/20180710_iRef/pf/013_B2x.dat",
      ".../Data/20180710_iRef/pf/013_B2y.dat",

```

```

    ".../Data/20180710_iRef/pf/013_params_v15.dat"},
{".../Data/20180710_iRef/pf/028_B2x.dat",
 ".../Data/20180710_iRef/pf/028_B2y.dat",
 ".../Data/20180710_iRef/pf/028_params_v15.dat"},
{".../Data/20180710_iRef/pf/043_B2x.dat",
 ".../Data/20180710_iRef/pf/043_B2y.dat",
 ".../Data/20180710_iRef/pf/043_params_v15.dat"}}

(* fit steps and obtain d & n *)
For[dn=1, dn<=Length[data], dn++,
  D2xf=data[[dn,1]];
  D2yf=data[[dn,2]];
  D2x=ReadList[D2xf,Number,RecordLists->True];
  D2y=ReadList[D2yf,Number,RecordLists->True];
  D2=D2x+I*D2y;
  rangemin=10;
  rangemax=25;
  paramsf=data[[dn,3]]; (* line trace parameters *)
  params=ReadList[paramsf,Number,RecordLists->True]
  results=ConstantArray[0,{Length[params],4}]; (* result vector *)
  For[j=1, j<=Length[params], j++,
    Print["DATA ",dn," LINE ",j];
    x=params[[j,1]]+1;
    y=params[[j,2]]+1;
    linelength=params[[j,3]]; (* the sign of the line length
                                determines the direction of the
                                line trace *)
    xstart=x+Min[0,linelength];
    xend=x+Max[0,linelength];
    r=Abs[If[linelength>0,
             D2[[y,xstart;;xend]],
             Reverse[D2[[y,xstart;;xend]]]]];
    p=Arg[If[linelength>0,
             D2[[y,xstart;;xend]],
             Reverse[D2[[y,xstart;;xend]]]]];
    Clear[ar,ap,b,c,dr,dp,er,ep,fr,fp,g,rvalue,pvalue,range];
    (* clear fit parameters & values *)
    r0=Transpose[{Range[Abs[linelength]+1]-1,r}];
    p0=Transpose[{Range[Abs[linelength]+1]+Abs[linelength],p}];
    d0=Join[r0,p0];
    fit=ConstantArray[0,1+rangemax-rangemin];
    err=fit;

```

```

For[range=rangemin, range<=rangemax, range++,
  fit[[1+range-rangemin]]
  =FindFit[
    Join[d0[[1+Floor[Abs[linelength]/2,1]-range
      ;1+Floor[Abs[linelength]/2,1]+range,
      All]],
    d0[[2+Floor[Abs[linelength]/2,1]-range
      +Abs[linelength]
      ;2+Floor[Abs[linelength]/2,1]+range
      +Abs[linelength],
      All]]],
    {If[t<Abs[linelength]+0.5,
      (ar/2)*Tanh[(t-b)/c]+dr*t+er
      +(fr/2)*Sech[(t-g)/c],
      (ap/2)*Tanh[(t-b-Abs[linelength]-1)/c]+dp*t+ep
      +(fp/2)*Sech[(t-g-Abs[linelength]-1)/c]],
      {Abs[linelength]/3<b<2*Abs[linelength]/3,
        1/3<c<3}},
    {{ar,Max[r]-Min[r]},{b,Floor[Abs[linelength]/2,1]},
    {c,1},{dr},{er,Min[r[[1]],r[[-1]]]},
    {fr,Max[r]-Min[r]},{g,Floor[Abs[linelength]/2,1]},
    {ap,Max[p]-Min[p]},{dp},{ep,Min[p[[1]],p[[-1]]]},
    {fp,Max[p]-Min[p]}},
    t];
  {ar,b,c,dr,er,fr,g,ap,dp,ep,fp}
  ={ar,b,c,dr,er,fr,g,ap,dp,ep,fp}
  /.fit[[1+range-rangemin]];
dat=Join[d0[[1+Floor[Abs[linelength]/2,1]-range
  ;1+Floor[Abs[linelength]/2,1]+range,
  All]],
  d0[[2+Floor[Abs[linelength]/2,1]-range
  +Abs[linelength]
  ;2+Floor[Abs[linelength]/2,1]+range
  +Abs[linelength],
  All]]];
err[[1+range-rangemin]]
  =Sum[(1/(4*range+2))*(If[m<range+1.5,
    (ar/2)*Tanh[(dat[[m,1]]-b)/c]
    +dr*dat[[m,1]]+er
    +(fr/2)*Sech[(dat[[m,1]]-g)/c]
    -dat[[m,2]],
    (ap/2)*Tanh[(dat[[m,1]]-b-1

```

```

                                -Abs[linelength])
                                /c]
                                +dp*dat[[m,1]]+ep
                                +(fp/2)*Sech[(dat[[m,1]]-g-1
                                -Abs[linelength])
                                /c]
                                -dat[[m,2]]]^2,
                                {m,1,4*range+2}];
Clear[ar,b,c,dr,er,fr,g,ap,dp,ep,fp];];
{ar,b,c,dr,er,fr,g,ap,dp,ep,fp}
={ar,b,c,dr,er,fr,g,ap,dp,ep,fp}
/.fit[[Max[Flatten[Position[err,Min[err]],1]]]];
(* choose best fit range based on least-squares criterion *)
rvalue=1-Abs[ar]/(er+Abs[ar]/2);
pvalue=Abs[ap];
Deltar=1;
Deltap=1;
dvalue
=Deltad
*Flatten[
  Position[
    ((rvalue-Abs[s])/Deltar)^2
    +((pvalue-Arg[s])/Deltap)^2,
    Min[Min[((rvalue-Abs[s])/Deltar)^2
            +((pvalue-Arg[s])/Deltap)^2]]][[2]];
nvalue
=nw
+Deltan
*Flatten[
  Position[
    ((rvalue-Abs[s])/Deltar)^2
    +((pvalue-Arg[s])/Deltap)^2,
    Min[Min[((rvalue-Abs[s])/Deltar)^2
            +((pvalue-Arg[s])/Deltap)^2]]][[1]];
Print[Style["x",Italic]," = ",x];
Print[Style["y",Italic]," = ",y];
Print["line length = ",linelength," px"];
Print[Style["r",Italic]," = ",rvalue];
Print[Style["\["CurlyPhi]",Italic]," = ",pvalue," rad"];
Print[Style["d",Italic]," = ",dvalue," nm"];
Print[Style["n",Italic]," = ",nvalue];
rrange={Min[Min[r],FindMinimum[{(ar/2)*Tanh[(t-b)/c]+dr*t

```

```

+er+(fr/2)*Sech[(t-b)/c],
t>=0,t<=Abs[linelength]},
t][[1]]],
Max[Max[r],-FindMinimum[{-((ar/2)*Tanh[(t-b)/c]+dr*t
+er+(fr/2)*Sech[(t-b)/c]),
t>=0,t<=Abs[linelength]},
t][[1]]]];
prange={Min[Min[p],FindMinimum[{(ap/2)*Tanh[(t-b-1
-Abs[linelength])
/c]
+dr*t+er
+(fr/2)*Sech[(t-b-1
-Abs[linelength])
/c],
t>=0,t<=Abs[linelength]},
t][[1]]],
Max[Max[r],-FindMinimum[{-((ap/2)
*Tanh[(t-b-1-Abs[linelength])
/c]
+dr*t+er
+(fr/2)
*Sech[(t-b-1
-Abs[linelength])
/c]),
t>=0,t<=Abs[linelength]},
t][[1]]]];
Print[Show[{ListPlot[r0,
PlotStyle->Red,
PlotRange->rrange,
LabelStyle->FontFamily->"Calluna",
AxesLabel->{Row[{Style["x",Italic],
" (px)"}],
Row[{Style["r",Italic]}}],
PlotMarkers->Automatic],
Plot[(ar/2)*Tanh[(t-b)/c]+dr*t
+er+(fr/2)*Sech[(t-g)/c],
{t,1,1+Abs[linelength]},
PlotStyle->Red}]];
Print[Show[{ListPlot[p0,
PlotStyle->Lighter[Orange],
PlotRange->Full,
LabelStyle->FontFamily->"Calluna",

```

```

        AxesLabel->{Row[{Style["x",Italic],
                        " (px)}],
                    Row[{Style["\[CurlyPhi]",
                                Italic]}]},
        PlotMarkers->Automatic],
    Plot[(ap/2)*Tanh[(t-b-Abs[linelength]-1)/c]
          +dp*t+ep
          +(fp/2)*Sech[(t-g-Abs[linelength]-1)/c],
          {t,1,1+Abs[linelength]},
          PlotStyle->Lighter[Orange]]];
    Print[""];
    results[[j]]={rvalue,pvalue,dvalue,nvalue};
    (* place results in result vector *)
  ];
  Export["results_ "<>ToString[dn]<>".dat",results];
];

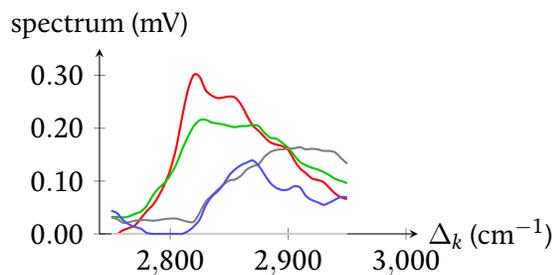
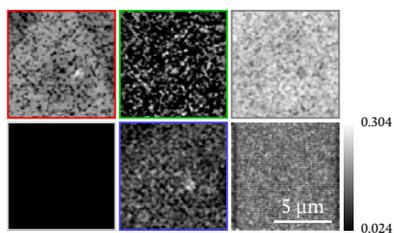
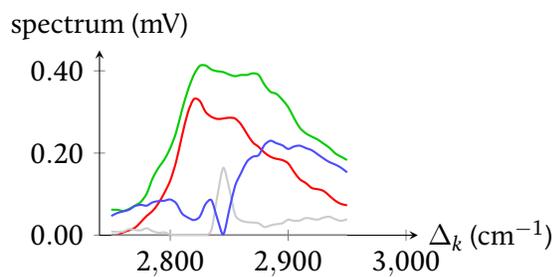
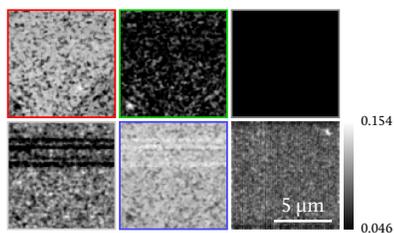
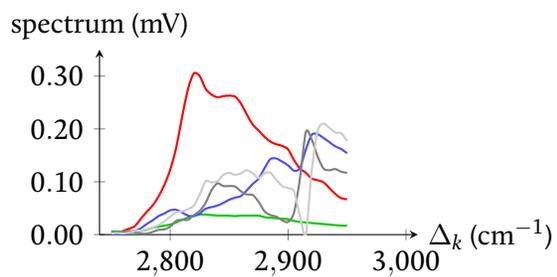
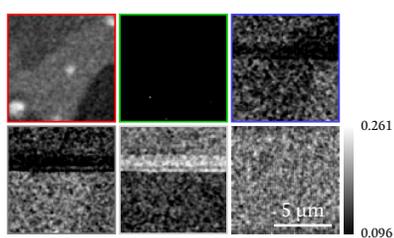
```

Appendix F

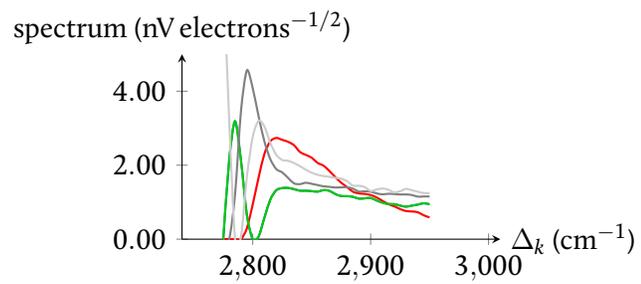
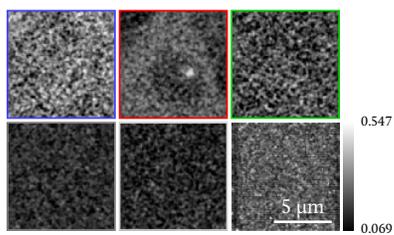
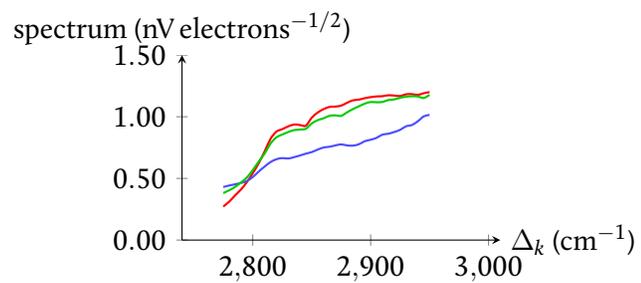
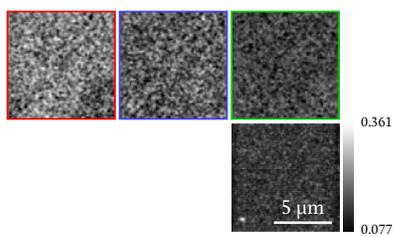
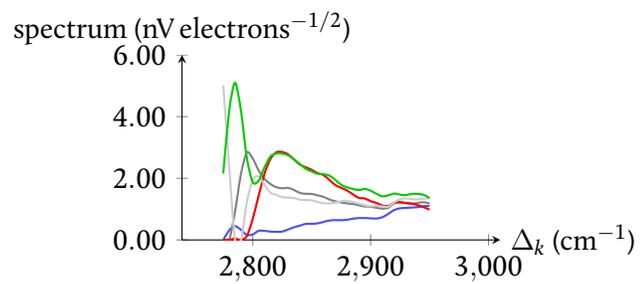
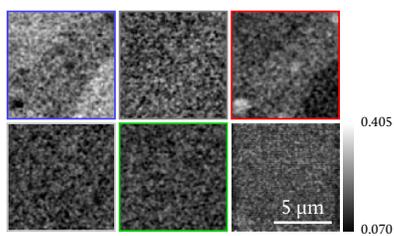
Full results of FSC³ of SRS images

Below are the full results of FSC³ of the images from chapter 5.

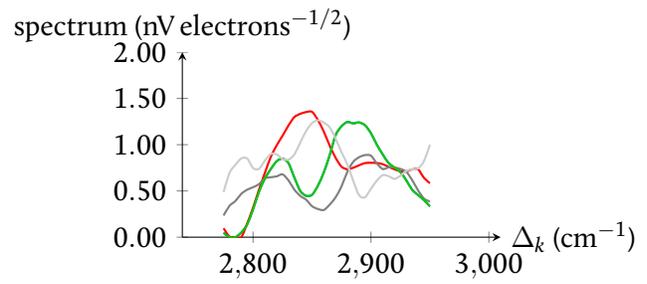
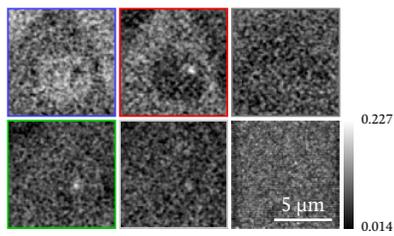
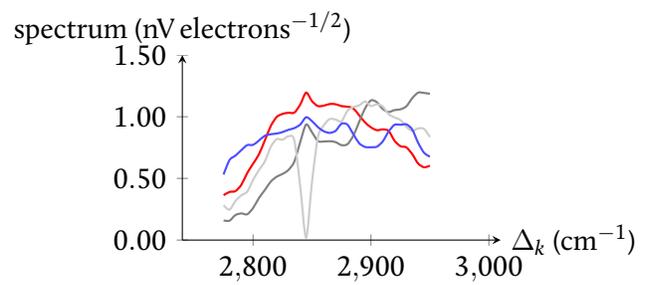
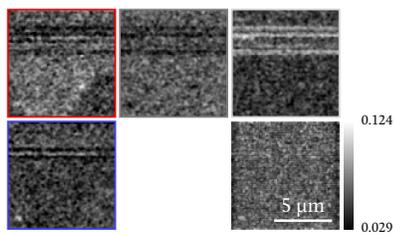
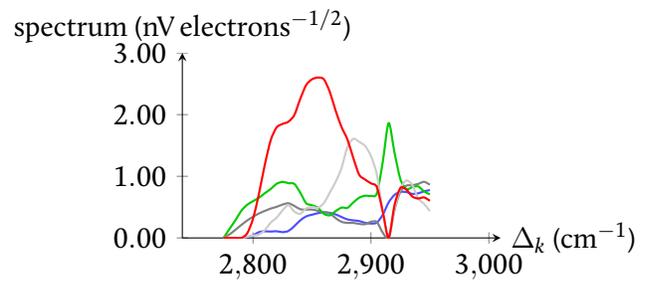
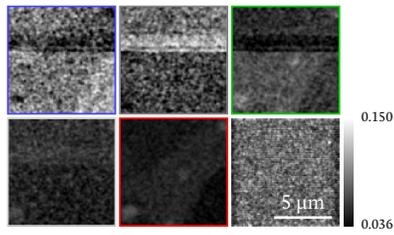
F.1. Figure 5.2



F.2. Figure 5.5



F.3. Figure 5.8



Appendix G

Parameters used for background fitting and FSC³

The following table shows the background (PBS-only region) fitting for gradient and PBS removal and the spectrum retrieval method for all the SRS images shown in figures 5.18–5.21. The images are listed in the same order of appearance as in the figures.

Lamellarity	Sample & image	Background fit	Spectrum retrieval method
Unilamellar	sm+ch 1	Linear in x & y	Spatial averaging of spectrum
	sm+ch 2		
	sm+ch 3		
	sm+ch 4		
Multilamellar	sm+ch 5	None	Spatial averaging of spectrum
	sm+ch 6		
	sm+ch 7		
Unilamellar	DOPC 1	Linear in x & y	Spatial averaging of spectrum
	DOPC 2		
Multilamellar	DOPC 3	None	Spatial averaging of spectrum
Unilamellar	Ternary 1	Linear in x & y	Unguided 2-component FSC ³
		Linear in y	Guided 2-component FSC ³
	Ternary 2	Linear in y	Unguided 2-component FSC ³
		Linear in x & y	Guided 2-component FSC ³
	Ternary 3	Linear in y	Unguided 2-component FSC ³
			Guided 3-component FSC ³
	Ternary 4	Linear in y	Unguided 2-component FSC ³
			Guided 2-component FSC ³
	Ternary 5	Linear in x & y	Unguided 2-component FSC ³
			Guided 3-component FSC ³

Appendix H

Publications and conference presentations

H.1. Publications

The following articles are in preparation. The first contains the interferometric reflectometry theory and will contain the results of future reflectometry experiments; the second contains our SRS results, as well as the results of CARS experiments performed by some of the other authors before this work began; and the third contains the detailed study of the Hodgkin-Huxley and Heimburg-Jackson models shows in chapter 2 and a proposal for experiments to help shed more light on neural communication from a biophysical perspective.

1. Nahmad-Rohen A & Langbein W (2020): *A model for the complex reflection coefficient of a collection of parallel layers*, in preparation
2. Nahmad-Rohen A, Regan D, McPhee C, Masia F, Borri P & Langbein W (2020): *Observation of lipid domain coexistence in single lipid bilayers using SRS and CARS*, in preparation
3. Nahmad-Rohen A & Langbein W (2020): *Our current understanding of action potentials and how to increase it*, in preparation

H.2. Conference presentations

Some of the work presented here has been the subject of the following conference presentations.

1. Nahmad-Rohen A, Langbein W & Borri P (2016): *Optical imaging of action potentials in live axons*, 1st Physics of Excitable Membranes Workshop, Copenhagen, Denmark (poster)

2. Nahmad-Rohen A, Regan D, Zorinants G, Langbein W & Borri P (2017): *Simultaneous measurement of thickness and refractive index by interferometric reflectometry*, 4th BioNanoPhotonics Symposium, Cardiff, United Kingdom (poster)
3. Nahmad-Rohen A, Regan D, Borri P & Langbein W (2018): *Simultaneous measurement of thickness and refractive index by interferometric reflectometry (iRef)*, 9th Photon Conference, Birmingham, United Kingdom (talk)
4. Nahmad-Rohen A, Regan D, Borri P & Langbein W (2019): *Simultaneous measurement of thickness and refractive index of lipid bilayers by interferometric reflectometry*, 32nd Focus on Microscopy Conference, London, United Kingdom (poster)
5. Nahmad-Rohen A, Regan D, Borri P & Langbein W (2019): *Hyperspectral coherent Raman imaging of coexisting domains in single lipid bilayers*, 32nd Focus on Microscopy Conference, London, United Kingdom (talk)
6. Nahmad-Rohen A (2019): *Theoretical and experimental study of lipid bilayer optics*, 1st Latin-American Research Symposium, Cardiff, United Kingdom (talk)