

A journey into the world of inverse problems in quantum mechanics

by
ELIOT BOLDUC

Submitted for the degree of Doctor of Philosophy

Heriot-Watt University

School of Engineering and Physical Sciences

August 2017

The copyright in this thesis is owned by the author. Any quotation from the thesis or use of any of the information contained in it must acknowledge this thesis as the source of the quotation or information.

Abstract

Technology has come a long way since the birth of quantum mechanics. The science has led to computers, and now, the scientists are pushing the fundamentals further to eventually be able to construct a quantum computer from the bottom up. Quantum tomography has a vital role in this ambitious endeavour: it's the study of how one can retrieve the values describing a quantum state, like finding coordinates on a map for a given position. The challenge that the quantum tomographer faces lies in the sheer number of these values, which grows exponentially with the components of quantum system. This is the curse of dimensionality and cannot be avoided with classical means. Therefore, the tomographer is forced to come up with algorithms that scale well with the number of components, either via prior information or by reducing the problem to its simplest form. In this thesis, we devise algorithms for retrieving quantum states using a little of both approaches. We develop a direct way of retrieving quantum state-vector values by assuming that the state is pure, which is often the case in optics. In addition, we show that a simple optimisation technique, projected gradient descent, can outperform all other methods for retrieving general quantum states. Our contribution to the field is thus to provide tools that enable the tomographer to work on larger quantum states and that hopefully help her create the building blocks of a quantum computer. We touch on other somewhat related subjects such as image denoising and imaging quantum correlations.

“Not all those who wander are lost.”

J.R.R. Tolkien

I would like to dedicate this thesis to my loving Chounette.

Acknowledgements

I would like to thank: my supervisor, Prof. Jonathan Leach, for his always pertinent and useful counsel and guidance throughout this PhD; all members of the Experimental Quantum Optics Group who have contributed to the work presented here, for their amazing work and all of the lively discussions.

Thanks also to my parents and my partner for their love and encouragement, which has been key to the success of this endeavour.

ACADEMIC REGISTRY
Research Thesis Submission

Name:	Eliot Bolduc		
School:	EPS/IPAQS		
Version: <i>(i.e. First, Resubmission, Final)</i>	Final	Degree Sought:	PhD

Declaration

In accordance with the appropriate regulations I hereby submit my thesis and I declare that:

- 1) the thesis embodies the results of my own work and has been composed by myself
- 2) where appropriate, I have made acknowledgement of the work of others and have made reference to work carried out in collaboration with other persons
- 3) the thesis is the correct version of the thesis for submission and is the same version as any electronic versions submitted*.
- 4) my thesis for the award referred to, deposited in the Heriot-Watt University Library, should be made available for loan or photocopying and be available via the Institutional Repository, subject to such conditions as the Librarian may require
- 5) I understand that as a student of the University I am required to abide by the Regulations of the University and to conform to its discipline.
- 6) I confirm that the thesis has been verified against plagiarism via an approved plagiarism detection application e.g. Turnitin.

* Please note that it is the responsibility of the candidate to ensure that the correct version of the thesis is submitted.

Signature of Candidate:		Date:	14/01/2018
-------------------------	---	-------	------------

Submission

Submitted By <i>(name in capitals)</i> :	
Signature of Individual Submitting:	
Date Submitted:	

For Completion in the Student Service Centre (SSC)

Received in the SSC by <i>(name in capitals)</i> :			
<i>Method of Submission</i> <i>(Handed in to SSC; posted through internal/external mail):</i>			
<i>E-thesis Submitted (mandatory for final theses)</i>			
Signature:		Date:	

Table of contents

Abstract	i
1 Introduction	1
1.1 Inverse problems	2
1.2 Mathematical tools for linear inverse problems	3
1.3 Motion blur example	4
2 Quantum states and measurements	8
2.1 The state vector	8
2.2 The density matrix	10
2.3 Multipartite systems	13
2.4 Types of measurements	23
2.5 Brief review of existing quantum tomographic methods	29
2.6 How is the thesis organized?	33
3 Direct measurement of a state vector	34
3.1 Introduction	34
3.2 Theoretical development	35
3.3 Experimental results	38
3.4 Discussion	51
3.5 Conclusion	53
4 Tomography through gradient descent	54
4.1 Introduction	54
4.2 Theoretical development	56
4.3 Numerical results	67
4.4 Discussion	77
4.5 Conclusion	78
5 Entangled photons coincidence detection with an EMCCD camera	79
5.1 Introduction	79
5.2 Theoretical development	80
5.3 Experimental results	86
5.4 Conclusion	90
6 Low-light-level image denoising	91
6.1 Introduction	91
6.2 Theoretical development	92
6.3 Numerical Results	101
6.4 Experimental results	104
6.5 Conclusion	105
7 Contributions to other published work	106

List of Publications

Papers

1. M. Agnew, E. Bolduc, K.J. Resch, S. Franke-Arnold, J. Leach. *Discriminating single-photon states unambiguously in high dimensions*. Physical Review Letters **113**, 020501 (2014)
2. T. Roger, S. Vezzoli, E. Bolduc, J. Valente, J.J.F. Heitz, J. Jeffers, C. Soci, J. Leach, C. Couteau, N.I. Zheludev, D. Faccio. *Coherent perfect absorption in deeply sub-wavelength films in the single-photon regime*. Nature Communications **6** (2015)
3. E. Bolduc, G. Gariepy, J. Leach. *Direct measurement of large-scale quantum states via expectation value of non-Hermitian matrices*. Nature Communications, **7** (2016)
4. E. Bolduc, D. Faccio, J. Leach. *Acquisition of multiple photon pairs with an EM-CCD camera*. Journal of Optics **19** 054006 (2017)
5. E. Bolduc, G. Knee, E. Gauger, J. Leach. *Projected gradient descent algorithms for quantum state tomography*. arXiv 1612.09531 (2016) [accepted, NPJ Quantum Information]
6. E. Bolduc, S. Chan, M. Agnew, M.J. Padgett, J. Leach. *A fast algorithm for low-light-level image denoising at video rate*. Manuscript in preparation (2017)

First author and speaker conference contributions

1. E. Bolduc, M. Agnew, J. Leach. *Video-rate denoising of low-light-level images acquired with a SPAD camera*. Oral presentation, Photonics North 2016 (May 2016)
2. E. Bolduc, M. Agnew, Y. Altmann, J. Leach. *Real-time denoising of low-light-level images acquired with a SPAD camera*. Oral presentation, IONS Quebec 2016 (May 2016)
3. E. Bolduc, G. Gariepy, J. Leach. *Tomographic measurement of a million-dimensional entangled state*. Invited oral presentation, Symposium on structured light and orthogonality in imaging, Rank Prize Funds (November 2014)
4. E. Bolduc, G. Gariepy, J. Leach. *Tomographic measurement of a (105x105)-dimensional entangled state*. Oral presentation, Frontiers in optics 2014 (October 2014)
5. E. Bolduc, G. Gariepy, J. Leach. *Tomographic measurement of a 11,025-dimensional entangled state*. Oral presentation, Photon14 (September 2014)

Chapter 1 – Introduction

In 1927, the act of "measuring" took a whole new meaning. Werner Heisenberg and others derived the so-called uncertainty principle [1, 2], one consequence of which was that an observer could never know the complete state of a quantum system. This came as a shock to many of the great minds of the 20th century physics, including Einstein who was relentlessly critical of quantum mechanics, especially regarding the treatment of entangled systems [3]. But to this day (as of 2017), Heisenberg's uncertainty principle still stands strong.

As a consequence of the uncertainty principle, one can never deduce the complete state of a single quantum system because a measurement disrupts the system in such a way that the information not provided by the measurement outcome is erased, at least in part. But for the experimentalists who wish to measure all the information about meticulously prepared quantum states, all is not lost. Experimentalists have fooled Nature into revealing her secrets by performing many different measurements on identically prepared quantum states. To gather more information than is possible on a single system, experimentalists perform a series of actions:

- prepare a quantum system in a particular state,
- perform the first measurement,
- discard the first, now perturbed, quantum system,
- prepare a second quantum system in the same way the first was prepared,
- perform a complementary measurement on the second system.

By repeating the process many times using complementary measurements, that is measurements that are not redundant, one can gather all the information that is necessary to completely describe the identically prepared quantum systems.

Once the experimentalist detains all the measurement outcomes, the game is not over. In quantum theory, the state of a system is mathematically modeled by a set of complex numbers that make up either a state vector or a state matrix, more commonly known as a density matrix. To map the set of measurement outcomes to the correct state vector or density matrix turns out to be a relatively hard problem. This mapping is difficult to make compared to the converse mapping: predicting the measurement outcomes with a known state. Because of this difficulty, quantum state retrieval is said to be in a class of "inverse problems", the general concept of which we detail in the next section.

A large part of this thesis addresses the retrieval of quantum states from known measurement outcomes via classical algorithms, which we will treat in Chapters 2, 3 and 4. The

latter concept is not to be confused with quantum algorithms, which are solely applied with the help of quantum computers rather than classical computers or modern laptop computers. Chapter 5 is an outlier since the subject examined in this part of the thesis is not an inverse problem, but is related to the other chapters in that quantum correlations are crucial to it. Chapter 6 presents another type of inverse problem in the context of quantum mechanics: the recovery of an image from data containing a very low number of photons.

1.1 Inverse problems

Locating the provenance of an earthquake is a perfect example of an inverse problem in classical mechanics. When two tectonic plates strike each other, a matter wave propagates to the surface of the earth. This wave can travel at different speeds depending on the type of matter it travels in. Once it reaches the surface, seismographs record the vibration intensities at many locations. With this data and a few mathematical tools for solving inverse problems, one can find the precise location and timing of the earthquake. Of course, it would be easier to predict the effects of the earthquake on the surface of the planet if the epicenter and intensity were known in advance. In the context of inverse problems, predicting the effects from a known cause is called the “forward problem”. However, the details of the cause of the earthquake are not known prior to the matter wave reaching the surface of the planet. One has to solve the earthquake inverse problem to uncover its epicenter and intensity.

Inverse problems occur when effects are known, but their cause is unknown. When the set of observed effects is sufficient to reconstruct the cause unequivocally, this set is said to be “informationally complete”. Provided with an informationally incomplete set of observations or measurements, one might find an erroneous solution to an inverse problem. In the game of Twenty Questions for example, the goal is to find the name of an object or a person given the answers to at most twenty yes/no questions. Each answer has the potential to reduce the space of possible solutions, until there is hopefully only one possible solution left, in which case one has reached an informationally complete set of answers. Inverse problems have also turned up in optics: medical tomographic imaging [4], looking around corners [5], image reconstruction with scattered light [6].

There exists many types of inverse problems, each with different levels of complexity. The simplest type of inverse problem would be a linear one, for which there exist a linear relationship between the cause and the effect. One example of a linear inverse problem would be reconstructing an image affected by motion blur. As a detailed example of a simple inverse problem and its solution, we provide a detailed reconstruction of such an image in Section 1.3, after going through the mathematical tools for solving such a problem.

Nonlinear inverse problems have a higher degree of complexity to them. The field of machine learning for example is filled with such problems: labelling handwritten digits [7], categorising breeds of cats and dogs [4] or the very large subfield of deep learning [8], etc. In a particular machine learning problem, one is typically not necessarily concerned by finding an exact solution, but rather by minimising the error rate of the trained algorithm. Quantum state vector reconstruction lies in the category of nonlinear inverse problems. This will be the subject of Chapter 3 and part of Chapter 4. We also touch on non-linear optimisation in Chapter 6 regarding image denoising.

An additional notable category of inverse problems is one where the relationship between the cause and the effect is linear, but the solution space is constrained. This is the case of density matrix reconstruction, which is the subject of Chapter 4. The remaining chapter (Chapter 5) is not about an inverse problem, but as mentioned above, there still is a link with the rest of the thesis: quantum correlations.

1.2 Mathematical tools for linear inverse problems

The pre-requisites for this section are linear algebra and a bit of matrix calculus. Most of the tools detailed in this section will be extensively used in Chapter 4. Throughout this section, we will apply the concepts of linear inverse problems to the case of reconstructing an image taken with a moving camera, thus producing motion blur.

For a typical linear inverse problem, we model the cause and the effect by a set of numbers organised in the vectors \mathbf{x} and \mathbf{b} , respectively. The linear relationship between the cause and the effect – the forward model – is written

$$A\mathbf{x} = \mathbf{b}, \tag{1.1}$$

where A is a transformation matrix representing a known phenomenon. The pseudo-inverse of the transformation matrix provides the solution to the inverse problem:

$$\mathbf{x} = (A^T A)^{-1} A^T \mathbf{b}. \tag{1.2}$$

However, in practice, in the presence of noise in the measurement outcomes \mathbf{b} , Eq. 1.2 often fails to yield a satisfying answer for the cause \mathbf{x} . This scenario happens whenever the matrix A is very sensitive to noise, which could either come from the data or from numerical instabilities.

We can characterise the sensitivity to noise of a matrix A through a metric called the “condition number”. This metric is defined using the singular value decomposition of the matrix $A = USV^T$, where U and V are unitary matrices ($U^T = U^{-1}$ and $V^T = V^{-1}$)

and S is a diagonal matrix containing the singular values. The metric mentioned above is equal to the ratio of the maximum to the minimum singular value $\sigma_{\max}/\sigma_{\min}$. A condition number of unity indicates that the matrix is as well-posed as can be, that is to say that the matrix and its inverse are both very robust to noise in the cause \mathbf{x} or the effect \mathbf{b} . As a side note, the case of a unit condition number implies that the matrix A is unitary. A condition number that is much higher than unity indicates that the matrix is ill-conditioned, and thus makes the above linear problem ill-posed in that small variations in the data vector \mathbf{b} produces large variations in the estimate of \mathbf{x} .

We will now apply these tools to the a deconvolution problem, which is prone to the high condition number issue mentioned above, to further clarify the meaning of A , \mathbf{x} and \mathbf{b} .

1.3 Motion blur example

In the case of motion blur deconvolution, the cause \mathbf{x} corresponds to the scene being photographed, the transformation A represents the movement of the camera and the effect \mathbf{b} is the resulting blurred image. A blurred image taken with a moving smart phone camera is shown in Fig. 1.1. We seek to reconstruct a more accurate depiction of the scene being photographed.

There is an apparent issue with the problem dimensionality at first since the cause and the effect are both supposed to be two-dimensional scenes whereas \mathbf{x} and \mathbf{b} are both one-dimensional vectors. This issue is quickly solved by considering one line of the scene at a time. If the image consists of $n \times m$ pixels, then line i of the burred image corresponds to \mathbf{b}_i and line i of the reconstructed image is written \mathbf{x}_i . We assume that the camera moves at constant speed in the horizontal direction, such that the lines i range from 1 to n .

When taking a picture with a camera, the detector acquires the scene over a finite period of time. If the camera is moving at constant speed over the duration of the acquisition, the intensity of the light reaching one pixel is highly correlated with that of its neighbors. The transformation matrix A has to accurately reflect the blurring effect over the correct number pixels, which depends on the product of the camera moving speed with the acquisition time. For a blurred image where the effect occurs over three pixels, the blurring



FIGURE 1.1: A picture taken with a camera phone moving at approximately constant speed in approximately the horizontal direction. The size of the image is $n = 146$ by $m = 611$.

transformation would read

$$A_{3,611} = \begin{bmatrix} 1 & 1 & 1 & 0 & 0 & \dots & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 & \dots & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & \dots & 0 & 0 & 0 \\ & & & & \vdots & & & & \\ 0 & 0 & 0 & 0 & 0 & \dots & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 & 0 & \dots & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 & 0 & \dots & 0 & 0 & 1 \end{bmatrix}, \quad (1.3)$$

with the notation $A_{b,m}$ including b the number of pixel over which the blur occurs and m the number of pixels in a one line of the image (also corresponding to the number of columns and rows corresponds). Here the hypothetical image would have dimensions of $n \times 611$. The transformation matrix adds the intensity of a pixel and two of its horizontal neighbors together: in this case, this is a convolution. In the case of the last few rows, the convolution is treated with the assumption that the boundary conditions are circular for convenience, but this will not lead to a physical solution for the edges of the treated image.

We now need to deconvolve the blurred image using a transformation matrix, but we do not know yet the exact number of pixels that have correlated intensities. By inspecting Fig. 1.1 closely, we estimate that the blurring effect occurs over about 50 pixels. That is to say that one pixel of \mathbf{x} is distributed over 50 pixels of \mathbf{b} . Instead of having three values of 1 on each row of A , as illustrated in the matrix of Eq. 1.3, we have in this specific case 50 consecutive 1 on each row. We now try to deconvolve each line of the blurred image using a more general version of Eq. 1.2:

$$\mathbf{x}_i = \text{pinv}(A_{50,611})\mathbf{b}_i, \quad (1.4)$$

where $\text{pinv}(\cdot)$ is the pseudo inverse operation, defined as

$$\text{pinv}(A) = \lim_{\delta \rightarrow 0} (A^T A + \delta I)^{-1} A^T. \quad (1.5)$$

The pseudo inverse operation is more general than the standard matrix inverse in that the former generally yields a finite result when applied to rank-deficient matrices. Here, the matrix $A_{50,611}$ is almost full rank, but we have to use the pseudo inverse. In practice, the pseudo inverse can be computed via the singular value decomposition:

$$\text{pinv}(A) = US^+V^T, \quad (1.6)$$

where $S_{jj}^+ = 1/S_{jj}$ if $S_{jj} \geq \epsilon$ and $S_{jj}^+ = 0$ if $S_{jj} < \epsilon$, with ϵ being a number that is typically slightly above machine precision.



FIGURE 1.2: **Naively deconvolved blurry image.** The computation is a deconvolution of Fig. 1.1, according to $\mathbf{x}_i = \text{pinv}(A_{50,611})\mathbf{b}_i$. The pseudo inverse of the blurring matrix amplifies noise greatly, such that the signal completely disappears.

The result of using the pseudo inverse to deconvolve the image of Fig. 1.1 is shown in Fig. 1.2. The pseudo inverse fails as it solely produces noise. The reason for this failure is that the convolution matrix $A_{50,611}$ has a very high condition number. In fact, this matrix has an infinite condition number because one of its singular values vanishes. However, if we ignore this one vanishing singular value, the minimum singular value of $A_{50,611}$ is 0.007238, while its maximum singular value is equal to 50; this would yield a condition number equal to 6908. This number is an indication of how much the noise is amplified by. We thus need another strategy to estimate the original scene.

A more sensible answer can be found using a regularisation technique [9]. To this effect, we simultaneously minimise both the the sum of the square errors $(A\mathbf{x}_i - \mathbf{b}_i)^T(A\mathbf{x}_i - \mathbf{b}_i)$ and the sum of the squares of the solution $\sum_j x_{i,j}^2 = \mathbf{x}_i^T \mathbf{x}_i$. This sum of the squares is high in the case of the naively deconvolved picture from Fig. 1.2. In a sense, the outcome of the deconvolution “exploded”. To keep noise spikes from occurring, we minimise the sum of the squares (or the L_2 norm) of the estimate, which also need to agree with the data. In the limit where we only minimise the sum of the squares, we get a perfectly smooth image with no signal either. Thus, there exists a tradeoff between the smoothness of \mathbf{x}_i and agreement of \mathbf{x}_i with the data \mathbf{b}_i . This tradeoff is encapsulated in a variable λ_0 weighting the sum the squares. Finally, we can express the above thoughts mathematically by minimising the quantity

$$y = (A\mathbf{x} - \mathbf{b})^T(A\mathbf{x} - \mathbf{b}) + \lambda_0\mathbf{x}^T\mathbf{x}, \quad (1.7)$$

which is done through isolating \mathbf{x} in $\partial y/\partial \mathbf{x} = 0$. We have omitted the subscripts for the sake of not crowding the equation too much. Using the identities $\partial(\mathbf{x}^T\mathbf{x})/\partial \mathbf{x} = 2\mathbf{x}$ and $\partial(A\mathbf{x})/\partial \mathbf{x} = A^T$, we find

$$\mathbf{x} = (A^T A + \lambda_0 I)^{-1} A^T \mathbf{b}. \quad (1.8)$$

This type of regularisation is a special case of what is called Tikhonov regularization [10]. The more general case involves a Tikhonov matrix Γ that multiplies \mathbf{x}_i in the regularisation term $\|\Gamma\mathbf{x}_i\|^2$, where $\|\cdot\|$ is the L_2 norm.

The deconvolution results for various values of λ_0 are shown in Fig. 1.3. A low value of lambda produces a noisy image that is more consistent with the real scene, that is a piece

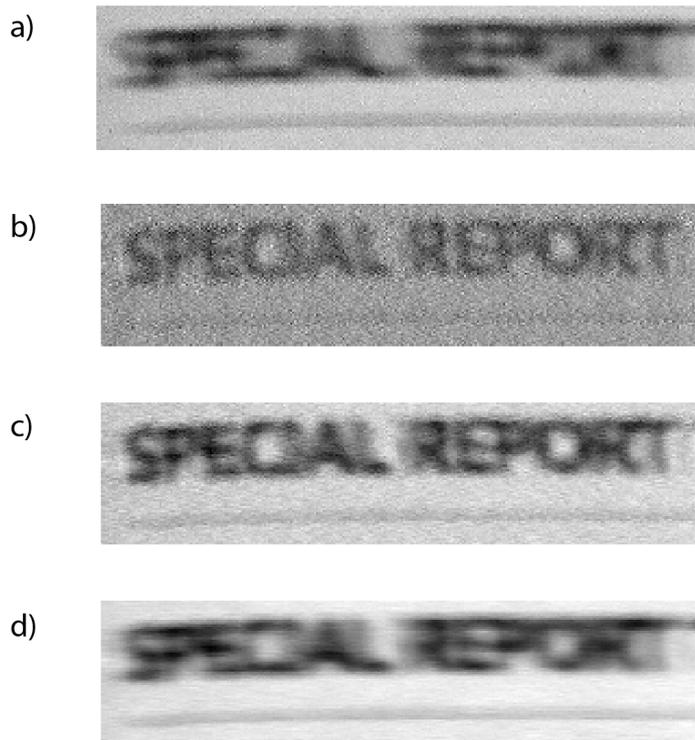


FIGURE 1.3: **Original picture (a) and deconvoluted versions using $\lambda_0 =$ b) 10 c) 80 d) 500.**

of paper on which the words “SPECIAL REPORT” are written. As we increase the value of λ_0 , the amount of noise decreases, but the deconvolution operation is made weaker.

This above deconvolution example is similar to density matrix retrieval in that they both are linear inverse problems. In the latter, the cause is the density matrix, the transformation matrix consists of a set of measurement operators and the effects are numbers of detector clicks. Notably, the derivative of Eq. 1.7 is central to Chapter 4. Before going into the heart of the thesis however, let us go through the basics of quantum mechanical states and their mathematical models.

Chapter 2 – Quantum states and measurements

At the heart of quantum mechanics lies the state vector and the density matrix. The mathematical description of the density matrix provides the most general description of a quantum system as it can include coupling with an environment, while the scope of the state vector is limited to quantum systems that are isolated from their environment, i.e. pure states. Both descriptions are important however: while the state vector has a limited scope, it is uniquely defined by a smaller set of numbers than the density matrix. The state vector can thus be recovered using fewer measurement outcomes, provided that the quantum system at hand is known to be pure prior to the measurement taking.

In the next two sections, we will outline the details of the state vector and the density matrix, and we pay particular attention to the vector notation that will be used throughout the next chapters. We also show how to generate random state vectors.

2.1 The state vector

The quantum state vector is an abstract model for describing a closed quantum system, may it be the spin of an isolated particle or the whole universe. The dimensionality d of a system corresponds to the number of orthogonal states that the system can be in. The spin of an electron can either be up or down, and therefore the electron spin is a two-dimensional system or a *qubit*. The position of a particle is a continuous infinite-dimensional degree of freedom. In this thesis, we only study the case of discrete quantum states, although the dimensionality is allowed to be infinite. One example of a discrete and infinite-dimensional degree of freedom is the number of photons in a wave-packet or light pulse.

A d -dimensional quantum state vector is uniquely defined by a set of d probability amplitudes $\{c_0, c_1, \dots, c_{d-1}\}$. State vectors live in a complex vector space, or a Hilbert space. The state vector is typically written with the *ket* notation:

$$|\Psi\rangle = \sum_{i=0}^{d-1} c_i |i\rangle, \quad (2.1)$$

which is short for

$$|\Psi\rangle = \begin{bmatrix} c_0 \\ c_1 \\ \vdots \\ c_{d-1} \end{bmatrix}. \quad (2.2)$$

We will also be using the the *bra* notation:

$$\langle\Psi| = \left[c_0^* \ c_1^* \ \dots \ c_{d-1}^* \right]. \quad (2.3)$$

It is not known exactly whether or not the state vector is a real entity or just a mathematical tool for predicting measurement outcomes; this is an active field of research [11]. For this reason, there is very little physical intuition for what the state vector elements really are.

2.1.1 The Born rule

A measurement on a d -dimensional state vector is modelled with a projections along an axis in Hilbert space. A measurement vector $|\phi\rangle$ takes the same form as a state vector:

$$|\phi\rangle = \sum_{i=0}^{d-1} a_i |i\rangle. \quad (2.4)$$

Both the state vector and the measurement vector must be normalised so that the expected outcome values can be interpreted as probabilities. Normalisation is done by ensuring that the inner product of a vector with itself is unitary: $\langle\Psi|\Psi\rangle = 1$ and $\langle\phi|\phi\rangle = 1$.

Measurement outcomes are predicted via the Born rule which states that the probability P that a system in state $|\Psi\rangle$ be found on the axis $|\phi\rangle$ is given by the absolute square of their inner product

$$P = |\langle\Psi|\phi\rangle|^2. \quad (2.5)$$

To be able to accurately measure the value of P , one has to prepare many identically prepared quantum states $|\Psi\rangle$ and repeat the same projective measurement until the desired accuracy is reached.

For a set of N measurements $\{|\phi_0\rangle, |\phi_1\rangle, \dots, |\phi_{N-1}\rangle\}$, it is trivial to calculate the corresponding outcomes $\{P_0, P_1, \dots, P_{N-1}\}$. However, it is considerable harder to solve the inverse problem, whereby we detain the measurements and their corresponding outcomes, and we seek to find the underlying state vector $|\Psi\rangle$. In fact, this situation is equivalent to the phase retrieval problem, which is not convex [12].

According to the Born rule, a measurement in the computational basis $\{|0\rangle, |1\rangle, \dots, |d-1\rangle\}$ would yield outcome $|j\rangle$ with a probability of $|c_j|^2$. A basis in dimension d is a set of d orthonormal vectors that can form a unitary matrix. In the vector notation, the Born rule can be expressed as

$$|\langle\Psi|j\rangle|^2 = \left| \begin{bmatrix} c_0^* & c_1^* & \dots & c_j^* & \dots & c_{d-1}^* \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 1 \\ \vdots \\ 0 \end{bmatrix} \right|^2 = |c_j^*|^2. \quad (2.6)$$

In general a projective measurement is given by

$$|\langle\Psi|\phi\rangle|^2 = |c_0^*a_0 + c_1^*a_1 + \dots + c_{d-1}^*a_{d-1}|^2, \quad (2.7)$$

which is a quadratic equation for the state vector elements. Hence, finding an unknown state vector $|\Psi\rangle$ from many measurement outcomes is equivalent to solving a set of quadratic equations [13]. Notably, the computational complexity of the Born rule in the context of pure states is $O(d)$ as computing an expectation value $|\langle\Psi|\phi\rangle|^2$ requires d multiplications.

So far we have introduced the notions of the quantum state vector and projective measurements. These notions are sufficient for modeling a noiseless quantum system, i.e. a pure state. Often in optics, assuming the purity of a quantum state is valid because there is very little interaction between photons and their environment. In general however, even in optics, we need a model that accounts for noisy systems and measurements, and this is the purpose of the density matrix and generalised measurements.

2.2 The density matrix

In principle, the universe could be represented with a single pure quantum state vector without loss of generality. To avoid storing or processing an unnecessary amount of information on everything contained within the universe, one models a quantum system by assuming that it is isolated from its environment. This strategy works perfectly if the quantum system is actually isolated, in which case the quantum system is accurately described as a pure quantum state. As soon as the quantum system is coupled to its environment, one must consider the joint system-environment body to observe fully coherent behavior. In the coupled system-environment case, ignoring the environment – also referred to as “tracing out the environment” – results in a loss of coherence, which is modeled by the density matrix. In a nutshell, the universe as a whole could be accurately

modeled by a pure quantum state, but in general a small subset of the universe is more accurately modeled by the density matrix because of coupling between the subset and the rest of the universe.

The density matrix ρ represents a statistical mixture of pure states:

$$\rho = \sum_{i=0}^{d-1} p_i |\psi_i\rangle\langle\psi_i|, \quad (2.8)$$

where p_i should be interpreted as positive probabilities, i.e. $\sum p_i = 1$, and the pure states $\{|\psi_0\rangle, |\psi_1\rangle, \dots, |\psi_{d-1}\rangle\}$ form an orthonormal basis. This formalism is more general than the state vector notation, and one can model the pure state of Eq. 2.1 using the density matrix with a single outer product:

$$\rho_{\text{pure}} = |\Psi\rangle\langle\Psi|. \quad (2.9)$$

If there are more than one non-zero weight, the quantum system described by the density matrix is said to be in a statistical mixture of pure states or a “mixed state” for short. The completely mixed state takes the form of the normalised identity matrix

$$\rho_{\text{fullyMixed}} = I/d, \quad (2.10)$$

which only occurs when all the weights are equal, i.e. $p_i = 1/d$.

In real experiments, the density matrix typically lies in the partially mixed state space. The metric for how close the density matrix is to being pure is the purity, given by

$$\text{Purity} = \text{Tr}[\rho^2], \quad (2.11)$$

and a pure state has unit purity. Another way to quantify the level of purity is through the Von Neumann entropy:

$$S(\rho) = -\text{Tr}(\rho \ln \rho), \quad (2.12)$$

where $\ln \rho$ is the matrix logarithm operation. The entropy of a pure state is zero, and maximal entropy is reached for a maximally mixed state. Often, the desired quantum system is pure, but as discussed above, coupling with the environment renders the system partially mixed. If the purity of the system is *close enough* to unity however, the state vector formalism can be a sufficiently accurate model. The details of what defines *close enough* depends on the application and the necessary level of accuracy.

We can deduce some of the properties of the density matrix as direct consequences of its definition. Firstly, the density matrix is always Hermitian, i.e. $\rho = \rho^\dagger$, due to the outer product symmetry. Secondly, the density matrix can always be decomposed into

its eigenvalues $\{p_0, p_1, \dots, p_{d-1}\}$ and eigenvectors $\{|\psi_0\rangle, |\psi_1\rangle, \dots, |\psi_{d-1}\rangle\}$. Because its eigenvalues are taken from a probability distribution – and thus positive – the density matrix is part of the positive matrices set. Hence, in quantum state recovery one needs to constrain the search space to positive matrices. Finally, the density matrix is normalised in that its trace is equal to one: $\text{Tr}[\rho] = 1$. This property can be derived from the facts that each eigenvector is normalised and that the weights sum to unity.

2.2.1 The Born rule reformulated for quantum state reconstruction

In the context of the density matrix, the Born rule of Eq. 2.5 takes a different form. The probability P that the quantum system be successfully projected onto the axis $|\phi\rangle$ is given by the statistical mixture

$$P = \sum_{i=0}^{d-1} p_i |\langle \psi_i | \phi \rangle|^2. \quad (2.13)$$

The projective measurements embedded in Eq. 2.5 and Eq. 2.13 are a special case of what is called “generalised measurements”. A measurement could correspond to a projection onto a mixture of axes. Similarly to quantum states, measurements are more accurately modeled by matrices. In this formalism, a projector $\hat{\Pi}_\phi$ is simply given by the outer product of its vector: $\hat{\Pi}_\phi = |\phi\rangle\langle\phi|$, where the hat signifies that the symbol it belongs to refers to a matrix or an operator. Using the matrix formalism for both the quantum state and the measurements, the Born rule takes yet another form:

$$P = \text{Tr}[\rho \hat{\Pi}_\phi], \quad (2.14)$$

which is the most common way of computing expectation values. The computational complexity of Eq. 2.14 determined by the matrix multiplication of two $d \times d$ -dimensional matrices, an operation that scales as d^3 . One can reduce the computational complexity for computing expectation values by recognising that Eq. 2.14 can be rewritten in the form

$$P = \sum_{\text{All elements}} \rho \circ \hat{\Pi}_\phi^T, \quad (2.15)$$

where the \circ operation stands for the Hadamard product, i.e. element-wise multiplication, and the T superscript for the matrix transpose. The Hadamard multiplication then becomes the bottleneck of the computation, which now scales as d^2 . For the purpose quantum state recovery, we will unwrap both the density matrix and the measurement operator into column-vectors, such that they become $\boldsymbol{\rho} = \text{unwrap}(\rho)$ and $\boldsymbol{\pi}_\phi = \text{unwrap}(\hat{\Pi}_\phi^T)$ where the bolded lowercase symbols refer to vectors. Whether one unwraps a matrix row by row or column by column does not matter as long as the wrapping operation is consistent with the unwrapping operation. The final version of the Born rule, and the one that we will be using for the density matrix formalism throughout this thesis, is given by the

simple inner product

$$P = \rho \cdot \pi_\phi, \quad (2.16)$$

Note that Eq. 2.16 is linear with respect to the density matrix elements, while the Born rule from Eq. 2.5 relating state vector elements to an expectation value is quadratic.

In the context of the density matrix formalism, the problem of quantum state reconstruction amounts to solving a constrained set of linear equations. Given N projectors $\{\pi_{\phi_0}, \pi_{\phi_1}, \dots, \pi_{\phi_{N-1}}\}$ and expectation values $\{P_0, P_1, \dots, P_{N-1}\}$, one can construct a set of N linear equations where the density matrix elements are unknown. Since the number of unknowns in the density matrix is d^2 , the minimum number of independent measurements that are informationally complete is also equal to d^2 . Recall that one must ensure that the density matrix lies in the set of positive matrices. If the expectation values are noiseless, solving this quantum state reconstruction problem simply consists of solving the unconstrained set of linear equations. However, after adding noise to the model, density matrix retrieval becomes more involved. Solving this problem in the presence of noise requires one to take the positivity constraint into account; this is the subject of Chapter 4.

2.3 Multipartite systems

A multipartite system is simply one composed of more than one particle or subsystems. Subsystems that are independent of each other are said to be *separable*. This word originates from the mathematical description of a multipartite state $|\Psi_{1,2}\rangle$. Say $|\Psi_1\rangle$ and $|\Psi_2\rangle$ are state vectors representing two independent particles, then the state of the composite system can be written

$$|\Psi_{1,2}\rangle = |\Psi_1\rangle \otimes |\Psi_2\rangle, \quad (2.17)$$

where the \otimes operation stands for the Kronecker product. The above state is separable in that the state of each subsystem can be separated on either side of the Kronecker product. For example, if each system is a qubit, i.e. $|\Psi_1\rangle = a|0\rangle + b|1\rangle$ and $|\Psi_2\rangle = c|0\rangle + d|1\rangle$, then

$$\begin{aligned} |\psi_{1,2}\rangle &= (a|0\rangle + b|1\rangle) \otimes (c|0\rangle + d|1\rangle) \\ &= ac|00\rangle + ad|01\rangle + bc|10\rangle + bd|11\rangle. \end{aligned} \quad (2.18)$$

In vector notation, the above equation takes the form

$$|\psi_{1,2}\rangle = \begin{bmatrix} a \\ b \end{bmatrix} \otimes \begin{bmatrix} c \\ d \end{bmatrix} \quad (2.19)$$

$$= \begin{bmatrix} ac \\ ad \\ bc \\ bd \end{bmatrix} \quad (2.20)$$

In the density matrix notation, a separable state is written $\rho_{1,2} = \rho_1 \otimes \rho_2$. If a multipartite system is not factorisable in the Kronecker sense, then the subsystems are *entangled* with each other. Telling whether a density matrix is separable or not is a difficult problem in general that falls outside the scope of this thesis [14, 15]. However, there exists entanglement witnesses made to ease the experimental and theoretical process for distinguishing a separable state from an entangled one [16, 17].

2.3.1 Entanglement

Entanglement is an inherent feature of quantum mechanics and at the heart of the quantum information field. An entangled system comprises of multiple parts where at least two of the parts are coupled to each other, but all parts can be coupled to each other in principle as far as quantum mechanics is concerned. Without prior knowledge, a pure quantum system made of n qubits has to be described with 2^n complex values: the state vector grows exponentially with system size (and so does the density matrix). Quantum mechanics suffers the *curse of dimensionality*, which strongly justifies the effort put into quantum state reconstruction. As the size of quantum states being controlled in the laboratory grows, the quantum state reconstruction methods have to keep up.

Bell states

The most basic entangled states are the ones representing two coupled qubits: the Bell states named after John S. Bell for being central to his seminal paper on the EPR paradox [18]. The Bell states are given by

$$|\Phi^+\rangle = \frac{|00\rangle + |11\rangle}{\sqrt{2}} \quad (2.21)$$

$$|\Phi^-\rangle = \frac{|00\rangle - |11\rangle}{\sqrt{2}} \quad (2.22)$$

$$|\Psi^+\rangle = \frac{|01\rangle + |10\rangle}{\sqrt{2}} \quad (2.23)$$

$$|\Psi^-\rangle = \frac{|01\rangle - |10\rangle}{\sqrt{2}}. \quad (2.24)$$

	$\{ 0\rangle, 1\rangle\}$	$\{ +\rangle, -\rangle\}$	$\{ \circ \rangle, \ominus \rangle\}$
$ \Phi^+\rangle$	\parallel	\parallel	\times
$ \Phi^-\rangle$	\parallel	\times	\parallel
$ \Psi^+\rangle$	\times	\parallel	\parallel
$ \Psi^-\rangle$	\times	\times	\times

TABLE 2.1: **Correlations \parallel and anti-correlations \times in the Bell states outcomes depending on which basis is used for the measurements.** From a quantum state reconstruction perspective, it would be sufficient to measure correlations in two of the three bases to distinguish the Bell states.

Any Bell state cannot be written as a Kronecker product of two individual qubits. For the first two states $|\Phi^+\rangle$ and $|\Phi^-\rangle$, one qubit is correlated with the other qubit. If an experimentalist makes measurements on both qubits in the computational basis $\{|0\rangle, |1\rangle\}$, the outcomes will be the same: either both outcomes are $|0\rangle$ or they are both $|1\rangle$. For the two last Bell states $|\Psi^+\rangle$ and $|\Psi^-\rangle$, the outcomes are anti-correlated. The phase between the kets has no impact on this particular measurement, but it does affect a measurement made in a conjugate basis, which refers to either

$$\{|+\rangle, |-\rangle\} = \left\{ \frac{|0\rangle + |1\rangle}{\sqrt{2}}, \frac{|0\rangle - |1\rangle}{\sqrt{2}} \right\} \quad (2.25)$$

or

$$\{| \circ \rangle, | \ominus \rangle\} = \left\{ \frac{|0\rangle + \imath|1\rangle}{\sqrt{2}}, \frac{|0\rangle - \imath|1\rangle}{\sqrt{2}} \right\}, \quad (2.26)$$

where the \imath symbol stands for the imaginary number $\sqrt{-1}$. The first Bell state is correlated in both the computational basis and the $\{|+\rangle, |-\rangle\}$ basis. However, the second Bell state $|\Phi^-\rangle$ yields correlated outcomes in the computational basis and yields anti-correlated outcomes in the $\{|+\rangle, |-\rangle\}$ basis. The above statements can be verified via a basis change in the definition of the Bell states:

$$|0\rangle = \frac{|+\rangle + |-\rangle}{\sqrt{2}} \quad (2.27)$$

$$|1\rangle = \frac{|+\rangle - |-\rangle}{\sqrt{2}} \quad (2.28)$$

$$|00\rangle = \frac{|+\rangle + |-\rangle}{\sqrt{2}} \otimes \frac{|+\rangle + |-\rangle}{\sqrt{2}} = \frac{|++\rangle + |+-\rangle + |-+\rangle + |--\rangle}{2} \quad (2.29)$$

$$|11\rangle = \frac{|+\rangle - |-\rangle}{\sqrt{2}} \otimes \frac{|+\rangle - |-\rangle}{\sqrt{2}} = \frac{|++\rangle - |+-\rangle - |-+\rangle + |--\rangle}{2} \quad (2.30)$$

$$|\Phi^+\rangle = \frac{|++\rangle + |--\rangle}{\sqrt{2}} \quad (2.31)$$

$$|\Phi^-\rangle = \frac{|+-\rangle + |-+\rangle}{\sqrt{2}}. \quad (2.32)$$

Whether a multipartite state is entangled or not is a binary problem with a yes or no answer, but there also exists levels of entanglement. The Bell states model pairs of maximally entangled qubits. An example of a quantum state that is not separable and not maximally entangled would be

$$|\Psi_{\text{between}}\rangle = \frac{|00\rangle + |01\rangle + |11\rangle}{\sqrt{3}} \quad (2.33)$$

for instance. The level of entanglement of a particle with another one or its environment can be assessed by measuring and calculating the purity of said particle through Eq. 2.11. If the particle is in a pure state, then it is not entangled. If the particle is in a completely mixed state, then it is maximally entangled with something. The level of entanglement of a given quantum system is inversely proportional to its purity. However, to be able to calculate the purity of a subsystem within a larger system, one requires a tool called the *partial trace*.

2.3.2 Entanglement in the orbital angular momentum degree of freedom

One of the main experiments conducted in the context of this thesis involve single photons and their orbital angular momentum (OAM). First, we will describe what the OAM degree of freedom is, and then we will apply it in the context of entanglement.

Laguerre-Gaussian modes

A photon has a spin angular momentum of \hbar at all times: it is an intrinsic property of light. However, this photon could carry orbital angular momentum or not: this is a matter how its phase structure is shaped, which is a free parameter of nature. A photon carries OAM when its phase profile rotates about its axis of propagation [19], and the amount of OAM that is carried by a photon is quantified by an integer ℓ between $-\infty$ and ∞ .

Fig. 2.3 shows the isophase patterns of three beams: the first has no OAM, the second carries OAM of $\ell = 1$, and the third carries $\ell = 2$. An isophase pattern is the set of surfaces where the phase ϕ_ℓ is equal to a constant. The isophase pattern of the fundamental transverse spatial mode of light – the Gaussian type – has a flat phase profile when taking a cross section of its beam, as shown on the top of Fig. 2.3. The flat surfaces repeat every time the light propagates for the length of a wavelength. The second beam carries OAM of $\ell = 1$, which means that its phase pattern is a helix with one branch. The helix performs a full rotation about the axis of propagation inside a wavelength. The third beam carries OAM of $\ell = 2$ and is characterised by a helix with two branches separated by 180° . In general, there are $|\ell|$ branches on a helix representing the isophase pattern of a beam carrying OAM, and the branches are separated by $(360/|\ell|)^\circ$ from their neighboring

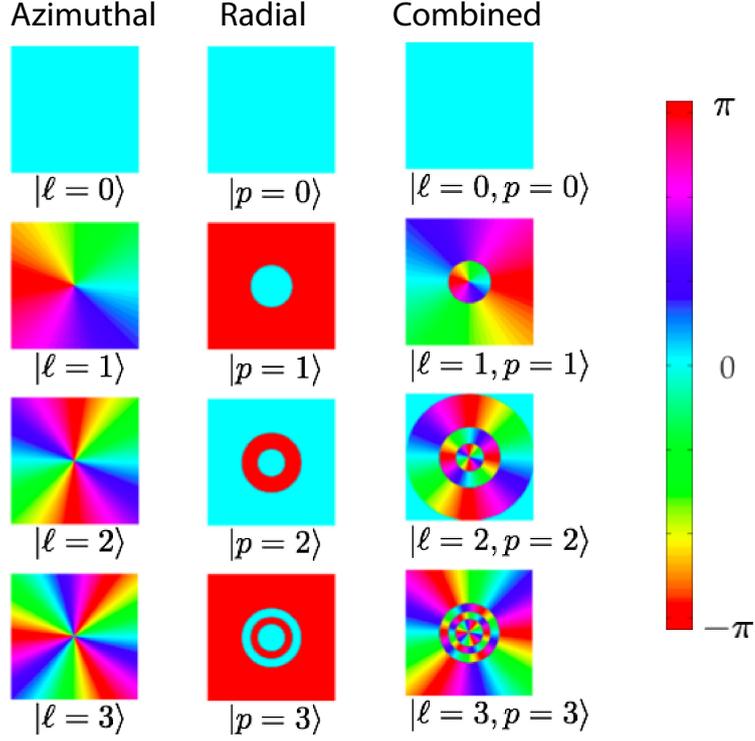


FIGURE 2.1: **Phase profiles of a few Laguerre-Gauss (LG) modes.** The azimuthal modes can be complex-valued and the solely radial modes cannot. The latter can either be real positive or real negative, and the number of zero-crossings is equal to the mode order. The intensity profiles of the azimuthal modes are shown in Figure 2.2.

branches.

At a given cross section, there is a linear relationship between the phase $\theta(\ell)$ and the angle of rotation ϕ around the propagation axis:

$$\theta(\ell) = \ell\phi + \phi_0 \quad (2.34)$$

where ϕ_0 is a constant phase that depends on the position along the propagation axis. By inspecting Eq. 2.34, we observe that all phases exist simultaneously in the center of a beam carrying a non-zero amount of OAM, thus creating a singularity in the structure. This singularity appears as a zero-intensity hole in the center of the intensity pattern, which generally give a donut-shape intensity profile to beams carrying OAM. This donut shape is typical of the Laguerre-Gaussian (LG) transverse spatial modes [19], which are eigenstates of free space that can carry OAM. Phase profiles of a few LG modes are shown in Fig. 2.1 and intensity profiles are shown in Fig. 2.2.

At the position of the waist w_0 , An LG mode is given by

$$|LG(r, \phi)_p^\ell\rangle = \frac{C_{\ell p}^{LG}}{w_0} \left(\frac{r\sqrt{2}}{w_0} \right)^{|\ell|} \exp\left(-\frac{r^2}{w_0^2}\right) L_p^{|\ell|} \left(\frac{2r^2}{w_0^2} \right) \exp(-i\ell\phi), \quad (2.35)$$

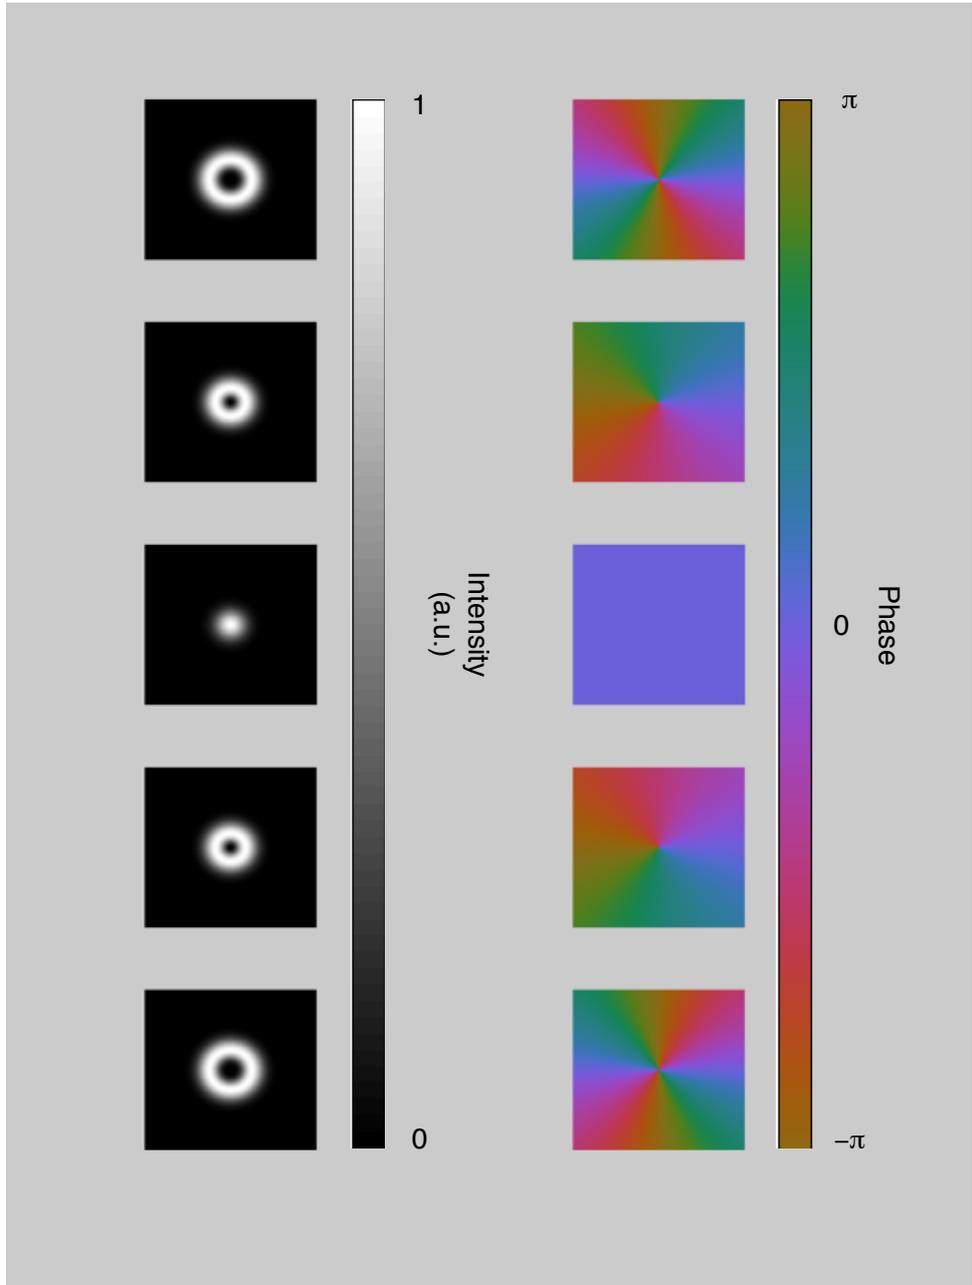


FIGURE 2.2: **Intensity and phase profiles of the azimuthal modes with ℓ ranging from -2 to 2.**

where C_{lp}^{LG} is a normalisation constant and $L_p^{|\ell|}()$ is a generalised Laguerre polynomial. These modes are crucial to the results of Chapter 3. The LG modes are eigenmodes of free space, which signifies that they are invariant under a Fourier transform. Concretely, the same equation can be used in one plane and in the far-field of this plane to describe the transverse spatial profile of LG modes. Fig. 2.3 and 2.4 illustrate these modes as a function of propagation. We let the size of the beam expand with propagation z , and we cut the structures at the radial position $r^2 = 2w(z)^2$.

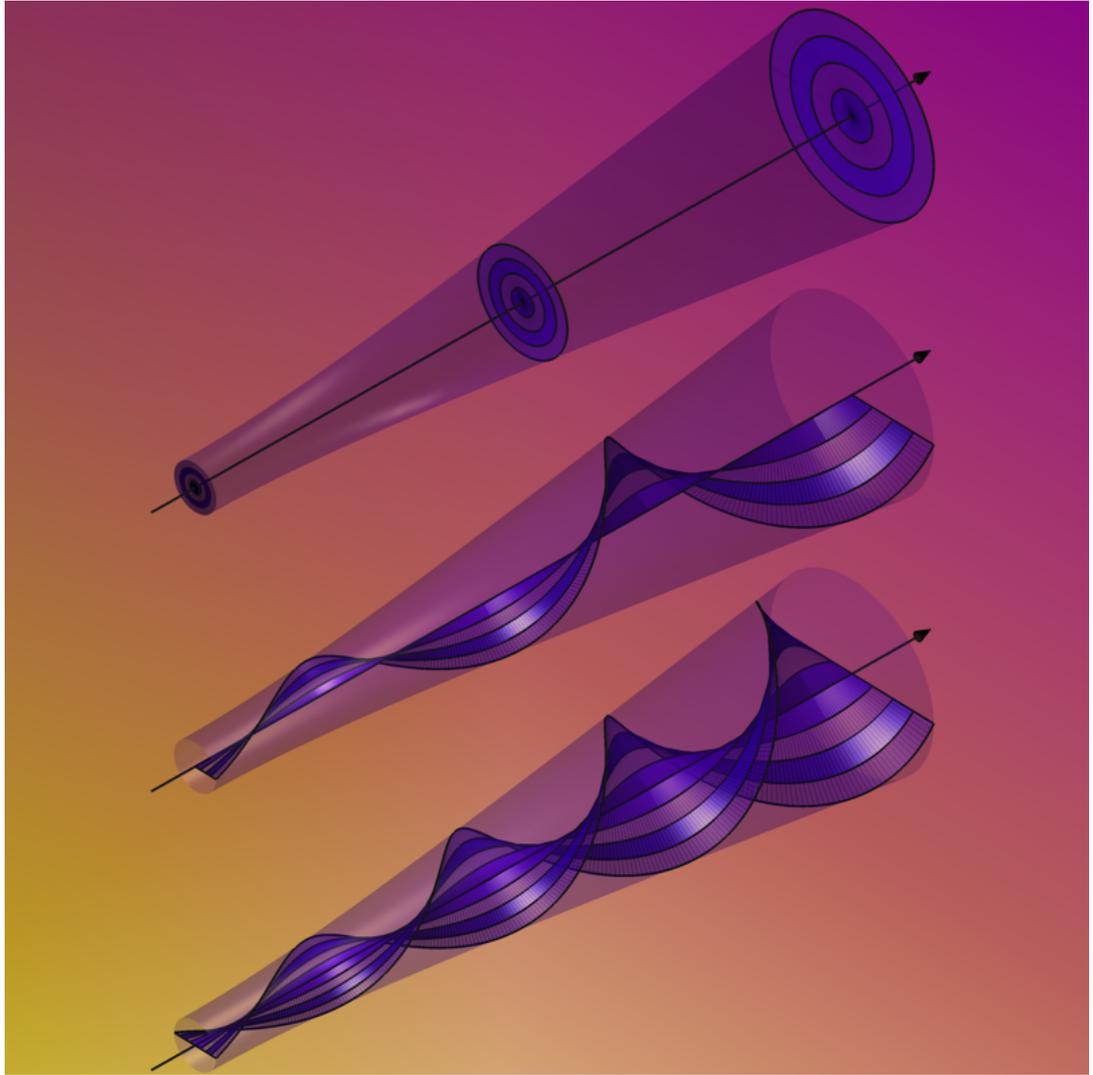


FIGURE 2.3: **Isophase patterns of expanding light beams with OAM of $\ell = 0, 1$ and 2 , from top to bottom.** The transparent cones are drawn to show the extent of helices at any given position along the axis of propagation. Alternating opaque and transparent surfaces on the isophase structures are drawn solely for the purpose of providing a better perspective. The isophase structures are cut when the intensity of the outskirts reaches $1/e^2$.

The state of a photon's transverse spatial mode can always be expressed as a superposition of Laguerre-Gaussian modes as it forms a complete basis. Mathematically, we can write the spatial state of a photon as

$$|\psi(r, \phi)\rangle = \frac{1}{N_\psi} \sum_{\ell=-\infty}^{\infty} \sum_{p=0}^{\infty} c_{\ell,p} |LG(r, \phi)_p^\ell\rangle, \quad (2.36)$$

where N_ψ is a normalisation constant and $c_{\ell,p}$ are the decomposition coefficients. In practice, we can't use an infinite number of coefficient to describe a transverse spatial state. Hence, we truncate the azimuthal and radial indices.

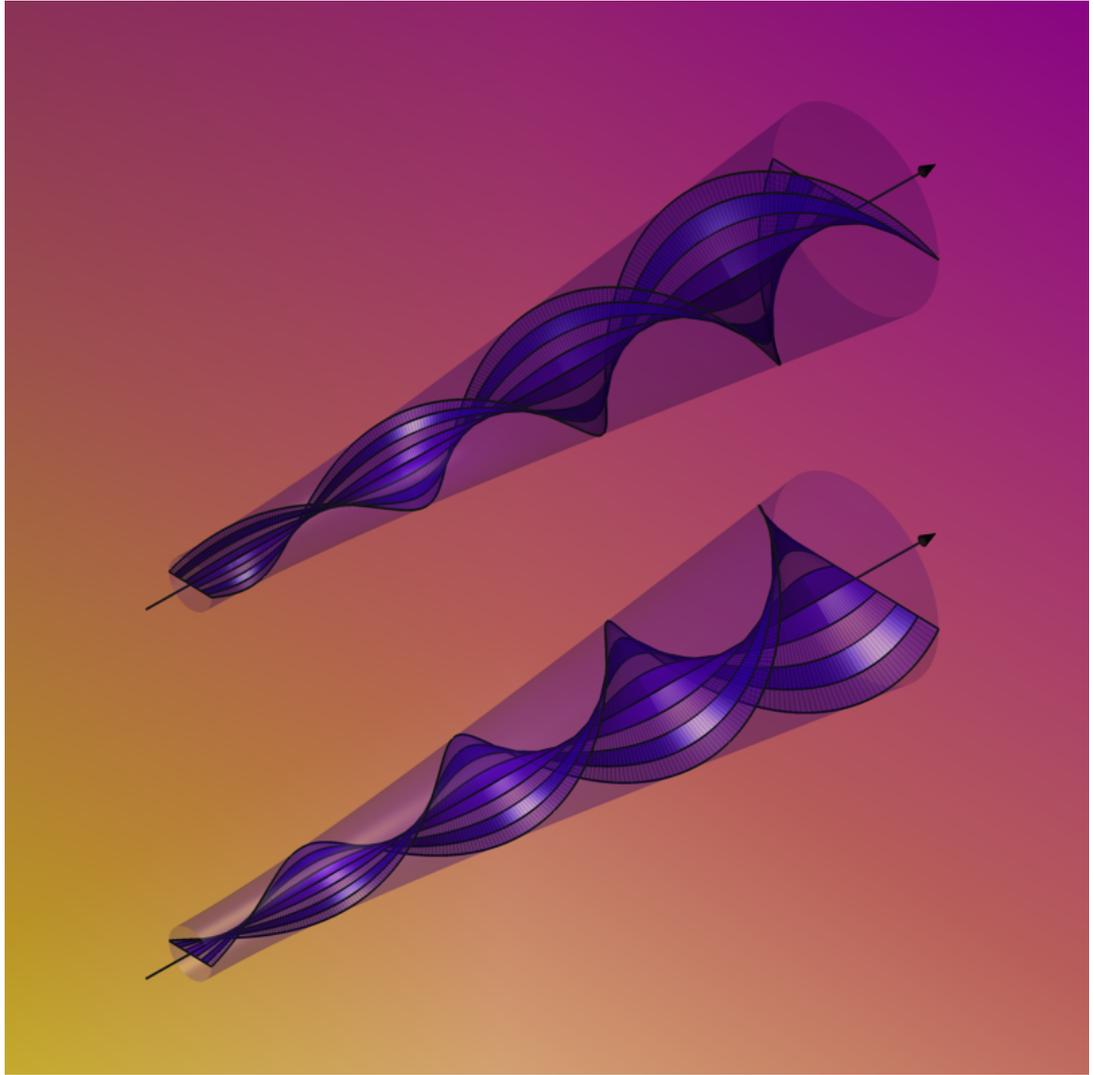


FIGURE 2.4: **Isophase patterns for light beams carrying OAM of $\ell = -2$ and 2 , shown from top to bottom.** The difference between the two patterns lies in the direction of rotation.

Entangled photon pair in the LG basis

Additionally, it is possible to express a joint spatial state of a pair of entangled photons using this basis. This time, however, one must specify which variable describes which photon. As a consequence, we need to add indices to all variables:

$$|\Psi(r_1, \phi_1, r_2, \phi_2)\rangle = \frac{1}{N_\Psi} \sum_{\ell_1=-\infty}^{\infty} \sum_{p_1=0}^{\infty} \sum_{\ell_2=-\infty}^{\infty} \sum_{p_2=0}^{\infty} c_{\ell_1, p_1, \ell_2, p_2} |LG(r_1, \phi_1)_{p_1}^{\ell_1}\rangle |LG(r_2, \phi_2)_{p_2}^{\ell_2}\rangle, \quad (2.37)$$

where N_Ψ is a normalisation constant and the $c_{\ell_1, p_1, \ell_2, p_2}$ are the complex coefficients of the joint spatial modes. In Chapter 3, the goal is to recover these complex coefficients in the most direct way possible. When two kets are side by side, it implies a Kronecker product \otimes between them.

Physical meaning can be extracted from the absolute square of a given coefficient: $|c_{\ell_1, p_1, \ell_2, p_2}/N_\Psi|^2$ is the probability of observing photon 1 in state $|LG(r_1, \phi_1)_{p_1}^{\ell_1}\rangle$ and pho-

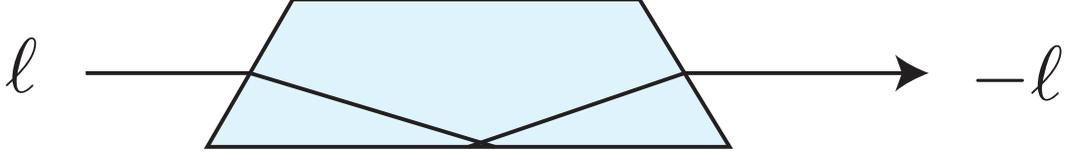


FIGURE 2.5: **Dove prism.** When an OAM-carrying beam traverses a dove prism, the helix direction is inverted.

ton 2 in state $|LG(r_2, \phi_2)_{p2}^{\ell_2}\rangle$ simultaneously. To obtain this probability experimentally, one would have to make a measurement using the joint projector:

$$|\Phi\rangle\langle\Phi| = |LG(r_1, \phi_1)_{p1}^{\ell_1}\rangle|LG(r_2, \phi_2)_{p2}^{\ell_2}\rangle\langle LG(r_1, \phi_1)_{p1}^{\ell_1}| \langle LG(r_2, \phi_2)_{p2}^{\ell_2}|, \quad (2.38)$$

the outcome of which is given by

$$|c_{\ell_1, p_1, \ell_1, p_2} / N_\Psi|^2 = \text{Tr}(|\Phi\rangle\langle\Phi||\Psi\rangle\langle\Psi|). \quad (2.39)$$

Here, the state at hand is pure, i.e. it is written as a state-vector. Hence, we can rewrite this probability as $\langle\Psi|\Phi\rangle\langle\Phi|\Psi\rangle$. The latter version only requires a matrix-vector and an inner product, while the Eq. 2.39 requires a matrix-matrix multiplication and a trace operation. Computationally, $\langle\Psi|\Phi\rangle\langle\Phi|\Psi\rangle$ is preferable to Eq. 2.39, a fact that we take advantage of in Chapter 4.

Spatial light modulators

There multiple ways to control the OAM of a light beam in the laboratory. One simple way is to send the beam through a dove prism, see Fig. 2.5. This inverts the sign of the azimuthal index, but does not change its absolute value. To be able to arbitrarily modify the phase and intensity profile of a light beam, we make use of spatial light modulators (SLMs) [20]. See Fig. 2.6.

An SLM is a computer-controlled array of pixels made of birefringent liquid crystal. Each pixel can be addressed via an applied voltage at its boundaries. The effect of the pixels on an incident light beam is to locally modify the phase of the beam. Hence, one can, not only invert the sign of the azimuthal index, but impart an arbitrary amount of OAM to a light beam. In fact, it is also possible to impart radial momentum to a light beam. The SLM is the ultimate tool for controlling transverse spatial modes. In order to add ℓ amount of OAM to a light beam, we need to display the phase profile of a beam carrying $-\ell$ amount of OAM on the SLM. In addition, we can turn a fundamental mode of light ($\ell = 0$ and $p = 0$) into a radial momentum carrying beam by displaying the LG mode with $p > 0$ on the SLM.

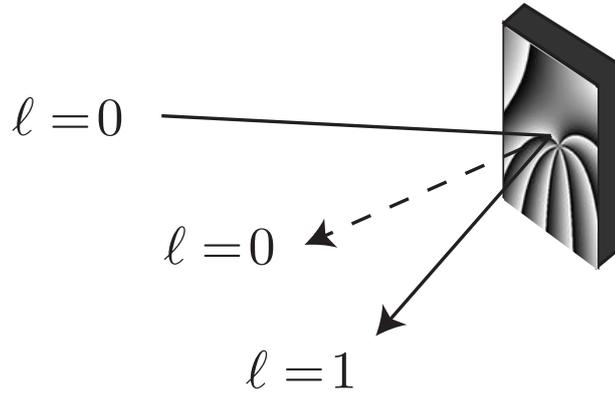


FIGURE 2.6: **Spatial light modulator.** An SLM can impart an arbitrary amount of OAM, within the limits of its capabilities. The reason the SLM hologram looks like the end of a fork is that, in practice, we display the combination of the $\ell = -1$ LG mode and a grating to separate the light that interacted with the hologram from the light that did not. SLMs are not 100% efficient, and part of the input energy reflects off the SLM like it would off a mirror, OAM unchanged.

The bottlenecks in producing high order radial LG modes with an SLM are that 1) the modes can grow bigger than the active surface of the SLM, and 2) the spatial frequency can surpass the available resolution on the SLM. Both the radial order and the azimuthal order of LG modes have corresponding larger radial profiles (see Fig. 2.2 for a visual representation of the larger azimuthal orders).

One way of making a projective measurement in the LG basis, is to use an SLM combined with a single-mode fiber. The single-mode fiber can only accept the fundamental mode of light in its core, that is $\ell = 0$ and $p = 0$. The concept of the measurement is explained in more details in the caption of Fig. 2.7. This setup is used in Chapter 3.

When it comes to quantum measurements, there are many choices to pick from, especially if we possess a tool as versatile as an SLM. With the combination of an SLM and a single mode fiber, we can project a spatial state on any combination of LG modes that we wish. The disadvantage of this apparatus is that we can only make one projective measurement at a time. Other devices, such as mode sorters, allow the experimentalist to make simultaneous projections by spatially sorting state components [21, 22]. The next section will describe all the approaches that one can take to perform a measurement in quantum mechanics.

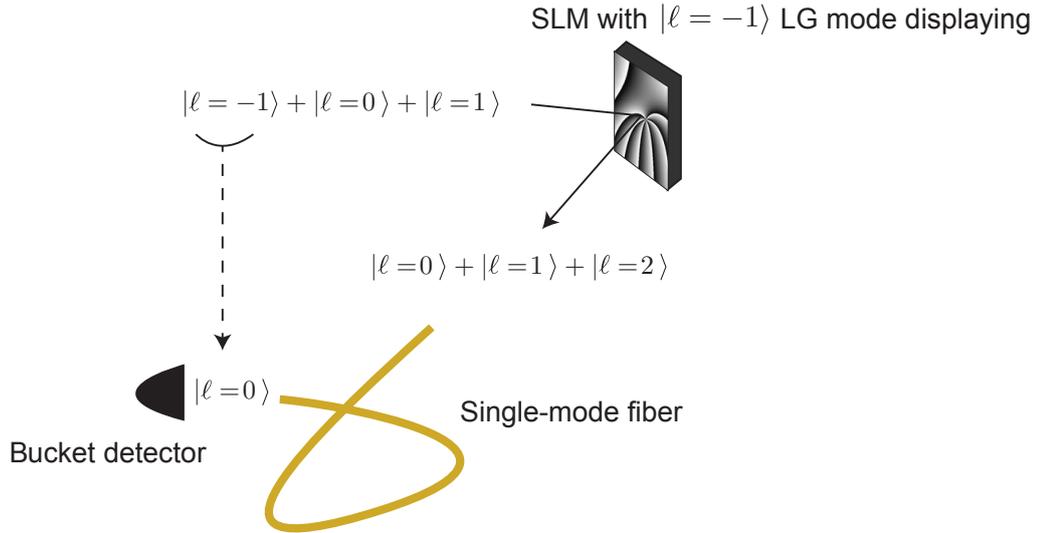


FIGURE 2.7: **Apparatus for projecting a superposition of LG modes into one of its eigenmodes.** Here, we want to measure the $\ell = -1$ and $p = 0$ component of an incoming light beam. Hence, we display an $LG_{p=0}^{\ell=-1}$ mode on the SLM screen. We project the resulting mode superposition onto the entrance facet of a single mode fiber, which filters out all but the fundamental mode. We then detect the light intensity or the number of photons with a bucket detector at that output of the single-mode fiber. The result is thus proportional to the $\ell = -1$ component of the incoming beam as we originally planned.

2.4 Types of measurements

Notably, physicist Wolfgang Pauli was one of the first to address the problem of quantum tomography when he asked whether the information of the position and the momentum of a particle was sufficient to uniquely reconstruct its quantum state [23]. As it will become evident in Chapters 3 and 4, the answer is in the negative. Later, Fano tackled the problem formally by defining the concept of “quorum” as the group of measurements that is sufficient to uniquely determine a density matrix [24].

Perhaps the most famous measurement operators that are informationally complete are the Pauli matrices:

$$\sigma_z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, \quad (2.40)$$

$$\sigma_y = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}, \quad (2.41)$$

$$\sigma_x = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad (2.42)$$

The Pauli matrices are Hermitian, as any observable should be. They are also full rank, i.e. their rank is equal to their dimension of 2. The rank of a Hermitian matrix is the number of non-zero eigenvalues that it has. For the sake of understanding what a quantum

measurement is, it is informative to decompose the Pauli matrices into rank-1 matrices

$$\sigma_z = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} - \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}, \quad (2.43)$$

$$\sigma_y = \frac{1}{2} \begin{bmatrix} 1 & -i \\ i & 1 \end{bmatrix} - \frac{1}{2} \begin{bmatrix} 1 & i \\ -i & 1 \end{bmatrix}, \quad (2.44)$$

$$\sigma_x = \frac{1}{2} \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} - \frac{1}{2} \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}. \quad (2.45)$$

Rank-1 matrices can always be constructed with an outer product. Therefore, the Pauli matrices can be expressed as a sum of outer products. In the computational basis $\{|0\rangle, |1\rangle\}$:

$$\sigma_z = |0\rangle\langle 0| - |1\rangle\langle 1| \quad (2.46)$$

$$\sigma_y = |+\ i\rangle\langle -\ i| - |+\ i\rangle\langle -\ i| \quad (2.47)$$

$$\sigma_x = |+\rangle\langle -| - |+\rangle\langle -| \quad (2.48)$$

When one makes a ‘‘Pauli measurement’’ of σ_z , what it concretely means is that we measure the $|0\rangle$ component of a state and its $|1\rangle$ component. Then, we subtract the $|1\rangle$ outcome from the $|0\rangle$ outcome. Only then do we have the result of a Pauli measurement. We can either measure the two components sequentially or simultaneously. Of course, it is more efficient to measure them simultaneously, but sometimes it’s impossible, as is the case for the SLM and single-mode fiber combination of Section 2.3.2. From the outcomes of Pauli measurements performed on a qubit, we can reconstruct this qubit density matrix ρ in the following way

$$\rho = \frac{1}{2} (I + \text{Tr}(\sigma_z \rho) \sigma_z + \text{Tr}(\sigma_y \rho) \sigma_y + \text{Tr}(\sigma_x \rho) \sigma_x). \quad (2.49)$$

The above equation does not prove very useful in practice, because the mere presence of noise can make the reconstructed density matrix ρ non positive-semidefinite. Chapter 4 provides robust ways of reconstructing the density matrix.

There exists a generalisation of Eq. 2.49 for many-qubit systems, where the many qubit are potentially entangled. We don’t use this generalisation in this thesis because of the reason mentioned in the above paragraph, but we do use the generalisation of a Pauli measurement of many-qubits systems in Chapter 4. For example, a two-qubit system can be determined using the Kronecker product of Pauli matrices with each other and with the

identity. A subset of the 16 possible combinations of Kronecker products is given by

$$I \otimes \sigma_z = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix} \quad (2.50)$$

$$\sigma_z \otimes \sigma_z = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (2.51)$$

$$\sigma_y \otimes \sigma_z = \begin{bmatrix} 0 & 0 & -i & 0 \\ 0 & 0 & 0 & i \\ i & 0 & 0 & 0 \\ 0 & -i & 0 & 0 \end{bmatrix} \quad (2.52)$$

$$\sigma_x \otimes \sigma_z = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \\ 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \end{bmatrix}. \quad (2.53)$$

The rest of the possible Kronecker products are

$$I \otimes \sigma_y, \quad \sigma_z \otimes \sigma_y, \quad \sigma_y \otimes \sigma_y, \quad \sigma_x \otimes \sigma_y, \quad (2.54)$$

$$I \otimes \sigma_x, \quad \sigma_x \otimes \sigma_x, \quad \sigma_x \otimes \sigma_x, \quad \sigma_x \otimes \sigma_x, \quad (2.55)$$

$$I \otimes I, \quad \sigma_x \otimes I, \quad \sigma_x \otimes I, \quad \sigma_x \otimes I. \quad (2.56)$$

For an arbitrary number of qubits n_q in a system, the number of possible combinations grows exponentially as 4^{n_q} . This is the curse of dimensionality. There is no way to reduce this number without making assumptions with a classical approach. There is however an efficient way of reconstructing large systems without making 4^{n_q} measurements and without relying on prior knowledge: that is through quantum principle component analysis [25]. This method will take over classical techniques when a quantum computer will be fully functional and stable, but for the moment, we have to rely on classical techniques. In Chapter 4, we make extensive use of the Pauli matrices and their Kronecker products.

2.4.1 Positive-operator valued measure or POVM

The formal definition of a POVM consists of a measurement operator set that gives the identity when we sum the operators. For example, the Pauli operator decompositions

from Eq. 2.57 to 2.59 can form POVMs.

$$I = |0\rangle\langle 0| + |1\rangle\langle 1| \quad (2.57)$$

$$I = |+\iota\rangle\langle -\iota| + |+\iota\rangle\langle -\iota| \quad (2.58)$$

$$I = |+\rangle\langle -| + |+\rangle\langle -| \quad (2.59)$$

Any basis in any dimension forms a POVM. The proof is straightforward. A basis is described by a unitary matrix U , which by definition obeys $U^\dagger U = I$, and $U^\dagger U$ happens to also be the sum of the basis operators. The rank-1 decomposition of Eq. 2.50 to 2.56 also form a POVM.

Structure in a set of measurement operators is usually beneficial. One type of POVM that has much intrinsic structure and is particularly useful is one that provides the most amount of information per operator: the SIC-POVM, which is the subject of the following subsection.

Symmetric, informationally complete, positive operator valued measure or SIC-POVM

A SIC-POVM is a set of d^2 rank-one POVM elements that are symmetric in the sense that the angular distance between any two different projectors is a constant: the overlap between the eigenvector of a POVM element and the eigenvector of another POVM element is $1/(d+1)$, See Fig. 2.8.

In Chapter 4, we generalise the concept of SIC-POVM to a set of N quasi-symmetric projectors, where $N < d^2$. These sets can be useful in the presence of prior information about a quantum state that we wish to reconstruct, which would signify that d^2 measurements are overcomplete or that a lower number of measurements are necessary to uniquely reconstruct a quantum state.

A SIC-POVM is optimal in the sense that, in the absence of prior information, a measurement outcome gives the maximum amount of information about the state being probed. However, implementing a SIC-POVM in the laboratory can prove challenging. A more accessible approach is using basis measurements to probe a quantum state. Mutually unbiased bases are optimal when one is limited to implementing basis measurements in the laboratory.

Mutually unbiased bases

Mutually unbiased bases (MUBs) are a set of bases which have maximum incompatibility with each other: the overlap between a vector from one basis and a vector from another basis is $1/d$. While SIC-POVM are not built out of bases, MUBS are. MUBS are known

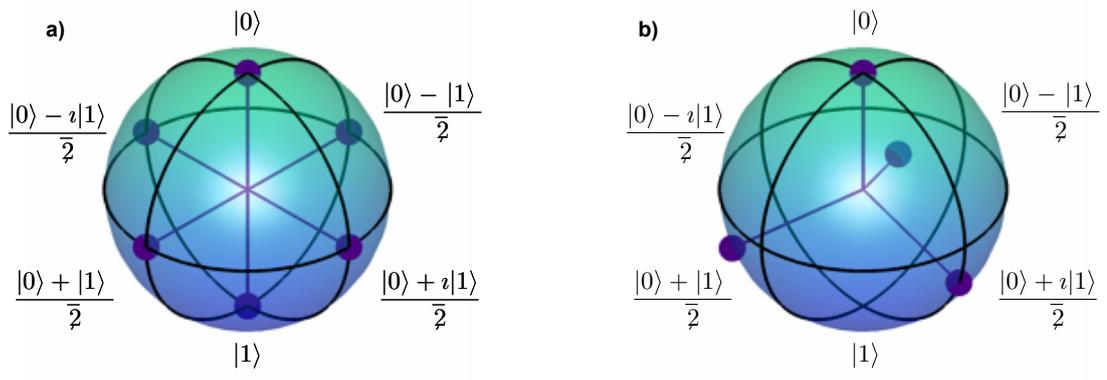


FIGURE 2.8: **Bloch spheres representation of a) Pauli measurement rank-1 decomposition and b) SIC-POVM elements for a qubit.** The computational basis $\{|0\rangle, |1\rangle\}$ corresponds to the North-South axes of the Bloch sphere, while the superposition bases correspond to the two other axes. The SIC-POVM is perfectly symmetrical and forms a tetrahedron in the Bloch sphere representation. The angle between each vector and each other vector is a constant. In the two-dimensional case, this constant is equal to 109.4712° . Because the SIC projections are maximally separated, they provide more information per measurement than the Pauli projectors.

to exist for dimensions equal to integer powers of prime numbers [26], but whether or not they exist in all dimensions is still an open question. When it is possible to implement a set of MUBs for quantum tomography, they are proved to be optimal when limited to bases measurements [26].

For a qubit, the projectors forming the rank-1 decomposition of any of the Pauli matrices (Eq. 2.57 to 2.59) also form a MUB. The two-dimensional bases are given by

$$\text{basis 1 : } \{|0\rangle, |1\rangle\} \rightarrow \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad (2.60)$$

$$\text{basis 2 : } \{|+\rangle, |-\rangle\} \rightarrow \begin{bmatrix} 1/\sqrt{2} & 1/\sqrt{2} \\ i/\sqrt{2} & -i/\sqrt{2} \end{bmatrix} \quad (2.61)$$

$$\text{basis 3 : } \{|+\rangle, |-\rangle\} \rightarrow \begin{bmatrix} 1/\sqrt{2} & 1/\sqrt{2} \\ 1/\sqrt{2} & -1/\sqrt{2} \end{bmatrix}. \quad (2.62)$$

We have already shown that each matrix forms a basis and is thus a unitary matrix. We now show that the above matrices form MUBs by calculating the overlap between each vector:

$$\text{basis 1 and 2 : } \left| \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}^\dagger \begin{bmatrix} 1/\sqrt{2} & 1/\sqrt{2} \\ i/\sqrt{2} & -i/\sqrt{2} \end{bmatrix} \right|^2 = \begin{bmatrix} 1/2 & 1/2 \\ 1/2 & 1/2 \end{bmatrix} \quad (2.63)$$

$$\text{basis 1 and 3 : } \left| \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}^\dagger \begin{bmatrix} 1/\sqrt{2} & 1/\sqrt{2} \\ 1/\sqrt{2} & -1/\sqrt{2} \end{bmatrix} \right|^2 = \begin{bmatrix} 1/2 & 1/2 \\ 1/2 & 1/2 \end{bmatrix} \quad (2.64)$$

$$\text{basis 2 and 3 : } \left| \begin{bmatrix} 1/\sqrt{2} & 1/\sqrt{2} \\ i/\sqrt{2} & -i/\sqrt{2} \end{bmatrix}^\dagger \begin{bmatrix} 1/\sqrt{2} & 1/\sqrt{2} \\ 1/\sqrt{2} & -1/\sqrt{2} \end{bmatrix} \right|^2 = \begin{bmatrix} 1/2 & 1/2 \\ 1/2 & 1/2 \end{bmatrix}, \quad (2.65)$$

where the $|\cdot|^2$ operation refers to the absolute square of each matrix element. We thus confirm that the overlap between any vector in a basis with any vector in an other MUB is equal to $1/2 = 1/d$.

In the laboratory, one could implement a MUB measurement by adding a unitary transformation, which – as the name suggests – transforms a state into any other state of the same space without loss, behind the measurement apparatus. This is easier said than done in the case of the spatial transverse mode of light since such a unitary transform is out of technological reach, but suggestions as to how it could be done have been given [27].

In this thesis, we do not investigate tomographic procedures with mutually unbiased bases, but the latter can be applied with the novel methods that we detail in Chapter 4. However, the procedure layed out in Chapter 3 is restricted to a fixed set of measurements that does not include MUBs.

2.4.2 Local measurements VS non-local measurements

In the presence of multiple bodies in a quantum system, the experimentalist can either perform local or non-local measurements. Local measurements only involves operations on individual bodies, while non-local measurements involves the interaction of many bodies together. Any measurement operator that is separable into Kronecker products is a local measurement, and if the operator is not separable, then it refers to a non-local measurement. For example, the Pauli Kronecker operators of Eq. 2.50 to 2.56 refer to local measurements. The Bell states shown in Eq. 2.21 to 2.24 can be converted into non-local measurement operators as they are non-separable. For an in-depth discussion of non-local measurements, see [28].

The experiments detailed in Chapter 3 only involve local measurements as it was easier for us to implement them in our laboratory. The algorithms of Chapter 4, however, can be readily applied to sets of non-local measurement operators, provided that these can be experimentally implemented.

2.4.3 Condition number of a measurement set

Sometimes, experimentalists have access to a limited range of quantum transformations and are thus restricted in the types of measurements that they can perform. Other times, it is simply easier to implement random measurements than to implement structured measurements. In any case, there are a few pitfalls that must be avoided when choosing a set of measurements.

One way of assessing the quality of a measurement set is using the condition number, as in the deconvolution example of Chapter 1. To do so, we stack the group of measurement operators into a single matrix. We have to unwrap each operator into a vector first, and then lay the vectors on top of each other to form a large matrix, which we call the “measurement matrix”. Notably, we use measurement matrices in Chapters 4 and 6. The condition number of a measurement matrix is given by

$$C_\sigma = \frac{\sigma_{\max}}{\sigma_{\min}} \quad (2.66)$$

where σ_{\max} and σ_{\min} are the maximum and minimum singular values, respectively. When the measurement matrix has a unit condition number, the corresponding inverse problem – reconstructing an unknown state – is generally easier than case of a large condition number.

It might be tempting for an experimentalist to ignore the structured measurements that we have seen so far, and implement a set of random measurements to achieve their goal. However, this would be a mistake as random measurements operators yield high condition numbers and therefore render the reconstruction problem either more difficult or, in some cases, impossible. The structured measurement sets – Pauli operators, SIC-POVMs and MUBs – all have low condition numbers, which provides a good reason to make use of them.

2.5 Brief review of existing quantum tomographic methods

The problem of quantum tomography can be stated as mapping a set of outcomes n_i to the best possible estimate of the density matrix ρ given that we know what measurements $\text{Tr}(\rho\hat{\Pi})$ took place and that the set of measurements $\{\hat{\Pi}_1, \hat{\Pi}_2, \dots, \hat{\Pi}_N\}$ is complete, i.e. there is a unique solution. The search space must be limited to positive-definite unit-trace matrices.

2.5.1 Maximum likelihood estimation

The field of quantum tomography has improved very much since Wolfgang Pauli asked the question about the completeness of the position and momentum of a particle [23], especially in the last two decades. The framework of maximum likelihood in the context of quantum tomography was adopted around the year 2000 with [29]. The likelihood \mathcal{L} of an estimate \mathbf{x} given some data \mathbf{n} is given by the product of all conditional probabilities:

$$\mathcal{L} = \prod_i P(x_i|n_i), \quad (2.67)$$

where we have assumed that all data points n_i are drawn from identical distributions. The main idea in this framework is to find the most likely estimator \mathbf{x}^* given the observations \mathbf{n} . This particular estimator must then maximise the likelihood (Eq. 2.67). The most widely adopted philosophy however is to minimise the negative of the log likelihood: the estimator \mathbf{x}^* that maximises the likelihood also minimises the negative log likelihood

$$-\log \mathcal{L} = -\sum_i \log P(x_i|n_i), \quad (2.68)$$

since the log of a product is also a sum of logarithms. The conditional probabilities depend on the precise setting of the experiment. The maximum likelihood framework is the starting point for the theory sections of Chapters 4 and 6. In the former chapter for example, we model shot noise, which has a Poisson distribution.

2.5.2 Gradient descent

There exists myriads of optimisation algorithms, but perhaps the most widely used is gradient descent. It owes its popularity to its low time complexity: calculate the gradient of a scalar function with respect to a density matrix with d^2 elements only takes time $O(d^2)$. Gradient descent is an iterative algorithm for converging to the minimum of a scalar function f that can depend in general on many parameters.

Say for example that the scalar function in question is the above negative log likelihood

$$f(\rho) = -\nabla_{\rho} \log \mathcal{L}(\rho) \quad (2.69)$$

and that it depends on a density matrix ρ . Then, to iteratively minimise the negative log likelihood with gradient descent, one subtracts the gradient from the current estimate ρ_k at iteration number k :

$$\rho_{k+1} = \rho_k - \nabla_{\rho} f(\rho) = \rho_k + \gamma \nabla_{\rho} \log \mathcal{L}(\rho), \quad (2.70)$$

where γ is a descent rate that is typically lower than one.

There are many flavours of the gradient descent algorithm, and we will explore some of them in Chapters 4 and 6. Many of the methods below also make use of some variation of gradient descent.

2.5.3 Full tomography algorithm

Here, we give the details of a few algorithms where no assumptions is made about the density matrix.

Density matrix tomography via Cholesky decomposition

Multiple techniques have also been built upon the maximum likelihood framework. Notably, James *et al.* have developed a quantum tomography using this framework with a Cholesky decomposition of the density matrix [30]. The Cholesky decomposition of a density matrix is given by

$$\rho = TT^\dagger. \quad (2.71)$$

The main idea underlying the algorithm of James *et al.* is to keep the density matrix estimate in the positive-definite space during a gradient descent procedure by making use of the fact that any positive-definite matrix has a unique Cholesky decomposition. However, this reparametrisation of the density matrix maps the constrained linear problem of density matrix estimation to an unconstrained non-linear problem: the outcome estimates

$$\text{Tr}(TT^\dagger \hat{\Pi}) \quad (2.72)$$

depend quadratically on the elements of the Cholesky decomposition matrix T . The problem is turned into a set of d^2 quadratic equations, where d is the density matrix dimensionality (the integer $d^2 - 1$ is the number of unknowns in the density matrix). The problem of solving a set of quadratic equations is NP-hard in general [12], which means that there are no known polynomial-time solutions to it for worst-case scenarios. As a sidenote, this problem is known to be easy when the number of equations is at least four times greater than the number of unknowns [12], but one would rather solve the density matrix reconstruction problem with the minimum number of measurements. Shang *et al.* reported that the reparametrisation of the density matrix into a Cholesky decomposition rendered the problem ill-conditioned and made the convergence slow [31].

The diluted iterative algorithm

Rehacek *et al.* also developed a gradient-based iterative algorithm for quantum tomography [32]. Some of the authors have refined their algorithm over the years [33], the last version of which we use as a benchmark in Chapter 4. See Chapter 4 for the details of the diluted iterative algorithm.

Quantum tomography via convex optimisation

Convex optimisation is a wide research field that treats problems that are convex, i.e. that have no local minima. As quantum tomography can be shown to be a convex problem [34], convex optimisers such as CVX can be readily applied to quantum state reconstruction [35]. The advantage of convex optimisers is that they can readily include convex constraints such as the positivity of matrices. We use the MATLAB implementation of CVX to benchmark our methods in Chapter 4. We give more details about convex optimisation in Chapter 4.

2.5.4 Pure-state tomography algorithm

Here, we give the details of a few algorithms where no assumption is made about the density matrix. This is the subject of Chapter 3 and a small part of Chapter 4.

In quasi-pure state tomography, one typically can use a subset of the measurements that would be sufficient to retrieve the density matrix without prior information. However, one question arises: which subset of the possible measurement should one use? The answer is to use a subset of measurements which are as far from each other as possible. We will explore this avenue in Section 4.3.1.

Compressed sensing

When the density matrix is known to have very few non-zero eigenvalues, the size of a complete set of measurements is greatly reduced [34]. By adding one line to the code using CVX, we can turn a full tomography algorithm with no assumptions to a quasi-pure state tomography algorithm. The line of code in question tells the CVX package to minimise the nuclear norm of the density matrix. The nuclear norm corresponds to the sum of the singular values. The algorithm can converge to the correct density matrix estimate even when the number of measurements is much lower than d^2 , which is the value for a complete set under no assumptions. The output of the compressed sensing algorithm is a density matrix whose eigenvalues are mostly close to zero.

The measurements that are typically used in the compressed sensing approach are chosen from a subset of informationally complete MUB measurements. The elements of the subset are chosen randomly, but the whole subset is still well-conditioned because the maximum overlap between a pair of MUB measurements is $1/d$. In general, the goal is to gain as much information per measurement as possible. This is achieved by choosing measurements that are close to being orthogonal.

Phase retrieval via Wirtinger flow

The phase retrieval problem is closely related to the pure-state tomography problem. In fact, the mathematics are exactly the same. In particular, the phase retrieval occurs in X-ray crystallography, where one probes an unknown crystal with a coherent X-ray beam. The experimentalist measures the intensity of the X-ray beam in the Fourier plane of the crystal. From many such intensity patterns, the problem consists of retrieving the underlying structure of the three-dimensional crystal.

In 2015, Candes *et al.* developed an algorithm for solving the phase retrieval problem via an initialisation procedure followed by gradient descent until convergence [36]. The

initialisation procedure finds an initial density matrix ρ_0 , which is the sum of the measurement operators weighted by the outcomes – this is very similar to Eq. 2.49. The initial density matrix is allowed to go outside the space of positive-definite matrices. They then find the eigenvector with the highest eigenvalue through the power method, i.e. repeated matrix-vector multiplications. Then they let gradient descent do the rest. Although the authors know that the problem is non-convex, they show that their algorithm converges to the correct answer provided that the number of measurements is greater than five times the dimensionality of the system. We independently developed a similar algorithm for pure-state quantum tomography, see Section 4.3.1.

Low-rank tomography through randomised SVD

Although it is simulated, one of the most impressive achievements in the field of quantum tomography is the reconstruction of a 16-qubit density matrix [37]. The number of unknowns in a density matrix this size amounts to a staggering $2^{32} = 4\,294\,967\,296$. One of the key parts of the algorithm needed to achieve this is to force the rank of the density matrix to be lower than some low integer r . As a consequence, one can write the SVD of the density matrix (or its eigenvalue decomposition) with compressed matrices. Another key part of the algorithm is the randomised SVD to be able to bring a density matrix estimate back to the positive-definite space. This algorithm resembles the ones we describe in Chapter 4. However, ours were developed independently and for a different purpose: full tomography.

2.6 How is the thesis organized?

At the start of every chapter, I state my specific contributions to the work being presented. These parts and the conclusions (Chapter 8) are the only ones where I use the first person, because the rest of the text represents work that was done in a team effort. After my contribution statements, I will lay the details of how the work presented in the chapter fits in the big scheme of the thesis. Then, a literature review is presented followed by the body of the work.

Chapter 3 – Direct measurement of a state vector

This work was initiated by myself. The goal is to provide a robust procedure for directly measuring the quantum wavevector. Dr Genevieve Gariepy and I derived the theory and I performed all the experiments with the help of my PhD advisor, Dr Jonathan Leach. This work is published in Nature Communications [38].

Quantum tomography of pure states consists of a constrained inverse problem, where the constrain lies in the space of normalised state vectors. This work does not apply to general quantum state tomography, but our method greatly simplifies the retrieval of quantum state vectors in that there is no post-processing of the recorded data required other than a linear weighted sum. This allows us to record and reconstruct the largest state-vector to date, to the best of our knowledge.

3.1 Introduction

One of the forthcoming challenges in the field of computing is harnessing the potential processing power of quantum technologies. Being able to measure the properties of a quantum system is at the very core of this challenge, and the quantum states to be characterized are often very large. When a quantum state is simply too large to be fully determined in a reasonable time, making assumptions about the state can render the task possible, especially when the state is close to being pure. In this chapter, we describe a novel, simple and fast way to characterize high-dimensional quantum states with certified purity. We demonstrate our method experimentally by characterizing the state vector of photon pairs spatially entangled in more than 100,000 dimensions. To fully measure a state of that magnitude with our method, the acquisition time is a few days rather than the centuries expected for assumption-free tomography. What saves time is the fact that it is direct in that it requires trivial post-processing of the large amount of acquired data. We apply the results to the computation of the Schmidt basis of the experimental system that we used and to the computation of whether the state at hand violates Bell inequalities or not [39–41]. The concepts of Schmidt decomposition and Bell inequalities will be explained in due time.

Propelled by the promise of great computational speedup, much progress has been made in the past few years towards producing and characterizing very large entangled quantum states [42–45]. While the high dimensionality of the state space is a necessary feature of practical quantum computers, this property comes at a price: large quantum

states quickly become intractable. The number of parameters that define a state scales exponentially with the components of a many-body system. In order to push the limit of the characterizable state spaces, one can make assumptions about the measured state.

Often, the outcome of standard quantum state tomography, i.e., the density matrix, gives redundant information. In particular, this is the case when the state is permutationally invariant [46], a matrix product state [47], or when it is almost pure [48–51]. The assumption that a state is quasi pure is already employed in a variety of methods: matrix-product tomography [47], weak-value assisted tomography [50, 52], and tomography via compressive sensing [34, 53]. The improvement provided by the pure-state assumption significantly enlarges the size of the state space that can be accessed through tomography. This is the assumption we use here: we assume that the measured quantum state is pure, which means that its density matrix should have at most one large eigenvalue. We then show that our method is valid even if the purity is not equal to unity. In fact, the purity can be as low as 0.8 and the method will still yield an average fidelity with the true state of 0.99. In other words, if there are a few non-zero eigenvalues in the density matrix, the latter can be accurately represented by a state-vector without great loss of fidelity.

In a general problem of pure state determination, the challenge is to measure those complex amplitudes representing the state out of real-valued measurements output. The core of our technique is to build an ensemble of what we call “column operators” that, when applied to the state, each give directly one of the complex amplitudes to be measured. However, these column operators are not observable, i.e. we cannot retrieve their output with a single measurement. We therefore decompose each column operator by a weighted addition of strong projective measurements. We then apply those measurements on the state and add the results according to the decomposition. Normalising the measurements then allows us to retrieve the exact complex amplitudes – in the absence of noise – that describe the state.

3.2 Theoretical development

To mathematically describe our direct measurement scheme, let us start with the simpler case of a single quantum system in a d -dimensional Hilbert space. We express a discrete state-vector as

$$|\Psi\rangle = \sum_{j=0}^{d-1} c_j |j\rangle, \quad (3.1)$$

where $c_j = \langle j|\Psi\rangle$ are the unknown complex coefficients we seek to measure in the basis of eigenstates $|j\rangle$. In order to probe these coefficients, we introduce the column operators $\hat{C}_j = |\nu\rangle\langle j|$, where $|\nu\rangle$ is an arbitrary reference vector. A column operator is associated

with a random variable whose expectation value

$$\langle \widehat{C}_j \rangle = \langle \Psi | \nu \rangle c_j \quad (3.2)$$

is proportional to a complex state-vector expansion coefficient. The key step is to measure the expectation values of the column operators \widehat{C}_j , although they are not observables in general. The reference vector is completely arbitrary, and for simplicity, we set $|\nu\rangle = |0\rangle$.

We now need to decompose the column operator into a set of observables. The column operators are not Hermitian in general. The idea of constructing non-Hermitian matrices from observables is not a new one [54]. There are many solutions to the construction of a given column operator. This decomposition is a crucial part of the method. Finding the minimal decomposition of a column operator is a non-trivial task, and to solve it, we use the differential evolution algorithm.

Differential evolution

In order to find the efficient decomposition of the many-body column operator from Eq. 3.5 in terms of projective measurements, we ran a differential evolution algorithm (see Algorithm 1) [55]. We set each vector element of the projectors and each complex weight as free parameters. We find the analytical solution by adjusting the parameters to the nearest rational numbers. The algorithm does not converge when we set the number of projectors to lower than five. This leads us to believe that five is the minimum number of local projectors for the exact construction of the two-body column-operator $\widehat{C}_{\ell_1, k_1}^{\ell_2, k_2} = |0, 0\rangle\langle \ell_1, k_1| \otimes |0, 0\rangle\langle \ell_2, k_2|$, where ℓ_1, k_1, ℓ_2 , and k_2 are not equal to zero.

The construction we find is:

$$\widehat{C}_j = |0\rangle\langle j| = \frac{2}{3} \sum_{k=0}^2 e^{i2\pi k/3} |s_{jk}\rangle\langle s_{jk}|, \quad (3.3)$$

where $|s_{jk}\rangle = (|0\rangle + e^{i4\pi k/3}|j\rangle)/\sqrt{2}$ and i represents the imaginary unit. For the case of $j = 0$, the column operator is simply the projector $|0\rangle\langle 0|$. By measuring the expectation values of all the column operators, we retrieve the first row of the density matrix.

As mentioned above, the expectation value of the j^{th} column operator is proportional to the corresponding complex coefficient. We can thus express the measured state vector as

$$|\Psi_M\rangle = \sum_{j=0}^{d-1} \langle \widehat{C}_j \rangle |j\rangle = \langle \Psi | \nu \rangle \sum_{j=0}^{d-1} c_j |j\rangle, \quad (3.4)$$

Algorithm 1 Differential evolution algorithm for finding the operator decomposition of \widehat{C}_{11} in a two-qubit system

- 1: Initialise hyperparameters: $F = 1$, $CR = 0.5$, population number $\beta = 50$
 - 2: Dimensionality of the problem: $D = 2 \cdot 5 \cdot 2 \cdot 5 = 100$
 - 3: Initialise $10 \times 50 = 500$ random state vectors $\Phi_{\alpha 1}^{(\beta)}$ and $\Phi_{\alpha 2}^{(\beta)}$
 - 4: Initialise target column operator: $\widehat{C}_{11} = |1\rangle\langle 1| \otimes |1\rangle\langle 1|$
 - 5: Compute estimated column operators: $\widehat{C}_{11}^{(\beta)} = \frac{4}{5} \sum_{\alpha=0}^4 |\Phi_{\alpha 1}^{(\beta)}\rangle\langle \Phi_{\alpha 1}^{(\beta)}| \otimes |\Phi_{\alpha 2}^{(\beta)}\rangle\langle \Phi_{\alpha 2}^{(\beta)}|$
 - 6: Arrange variables in a vector: $\mathbf{x}^{(\beta)} = [\text{Re}(|\Phi_{\alpha 1}^{(\beta)}\rangle), \text{Im}(|\Phi_{\alpha 1}^{(\beta)}\rangle), \text{Re}(|\Phi_{\alpha 2}^{(\beta)}\rangle), \text{Im}(|\Phi_{\alpha 2}^{(\beta)}\rangle)]$
 - 7: Note: state vector normalisation is not enforced from this point on: the normalisation constant will directly give the corresponding column operator weight
 - 8: **while** $\min_{\beta} \left\| \widehat{C}_{11}^{(\beta)} - \widehat{C}_{11} \right\|_2 > 10^{-5}$ **do**
 - 9: **for** every member of the population $x^{(\beta)}$ **do**
 - 10: Pick three other distinct members at random $\mathbf{x}^{(a)}$, $\mathbf{x}^{(b)}$ and $\mathbf{x}^{(c)}$
 - 11: Pick a random index R out of the $D = 100$ variables
 - 12: **for** every variable $x_i^{(\beta)}$ **do**
 - 13: Draw a number r_i from a uniform distribution $U(0, 1)$
 - 14: **if** $r_i < CR$ or $i = R$ **then**
 - 15: Compute a potential new position $y_i = x_i^{(a)} + F(x_i^{(b)} - x_i^{(c)})$
 - 16: **else** do not move current variable: $y_i = x_i^{(\beta)}$
 - 17: **if** $\left\| \widehat{C}_{11}^{(\beta)}(\mathbf{y}) - \widehat{C}_{11} \right\|_2 < \left\| \widehat{C}_{11}^{(\beta)}(\mathbf{x}^{(\beta)}) - \widehat{C}_{11} \right\|_2$ **then**
 - 18: Replace population member with candidate $\mathbf{x}^{(\beta)} = \mathbf{y}$
 - 19: **Return** $\text{argmin}_{\mathbf{x}^{(\beta)}} \left\| \widehat{C}_{11}^{(\beta)} - \widehat{C}_{11} \right\|_2$
-

with $|\nu\rangle = |0\rangle$. By normalizing the measured state, we retrieve the initial state $|\Psi\rangle$. As long as the overlap between the reference vector and the initial state is non-zero, normalization is possible. Thus, before performing the above procedure, we ensure that the measurement of the expectation value of the reference projector $|\nu\rangle\langle\nu|$ gives rise to a significant number of clicks on the detection apparatus, which ensures a non-zero overlap with the state. If this is not the case, we chose another reference vector; the higher the expectation value of the projector $|\nu\rangle\langle\nu|$ is, the more signal we get throughout the whole procedure.

All but one of the column matrices must be composed with at least three projectors each, as in Eq. 3.3; the remaining matrix, \widehat{C}_0 , can be composed by a single projector. The minimum number of projections required to construct all column operators is therefore $3d - 2$, a number that is consistent with existing work concerning minimal informationally-complete sets of projectors for uniquely determining single-system pure states [48, 49].

In cases where a quantum state is known to include two potentially coupled systems, each of which exists in a d -dimensional subspace, the construction of a column operator (found using the differential evolution algorithm detailed below) can be made with 5 joint

projectors:

$$\hat{C}_{pq} = |0\rangle\langle p| \otimes |0\rangle\langle q| = \frac{4}{5} \sum_{k=0}^4 e^{i2\pi k/5} |s_{pk}\rangle\langle s_{pk}| \otimes |s_{qk}\rangle\langle s_{qk}| \quad (3.5)$$

where $\sqrt{2}|s_{tk}\rangle = |0\rangle + e^{i4\pi k/5}|t\rangle$ and $p, q, t \in [1, 2, \dots, d-1]$. The column operators of this type account for more than 99% of the measurements shown in Figs 3.2 and 3.3. The column operators corresponding to the special cases $|0\rangle\langle p| \otimes |0\rangle\langle 0|$ and $|0\rangle\langle 0| \otimes |0\rangle\langle q|$ can be constructed with 3 joint projectors using the projector $|0\rangle\langle 0|$ on one system and the column operator from Eq. 3.3 on the other system. In this case, our method requires about $5d^2$ joint projective measurements which, in addition to the directness of the method, represents an improvement of a factor of $d^2/5$ in the number of projections required for full tomography. The greater the dimensionality of the system, the greater the improvement. Hence, we have the ability to measure the complex coefficients of the largest entangled state ever characterized.

3.3 Experimental results

As an experimental demonstration, we apply our method to the measurement of a two-photon state produced through spontaneous parametric downconversion (SPDC) in a nonlinear crystal. This process has been the subject of extensive study and is routinely used in laboratories [40, 45, 56–62]. Our experimental setup is illustrated in Fig. 3.1. The spatially entangled photon pairs are generated through SPDC with a 405-nm laser diode pumping a 1-mm-long periodically-poled KTP crystal with 50 mW of power. The generated photon pairs are entangled in their spatial degree of freedom, and their transverse profile is commonly expressed in the discrete Laguerre-Gaussian (LG) basis, described by a radial index p and an azimuthal index ℓ [19, 56]. Although the LG basis is infinite, tomographic measurements have previously been limited to a subset of the order of ten LG modes, leading to a total dimensionality of a few hundreds at most [57, 62]. As a preliminary result, we consider 1023 LG modes on each photon, and consequently obtain a more accurate description of the generated state. We choose a decomposition where the azimuthal index ℓ ranges from -15 to 15 and the radial index p ranges from 0 to 33:

$$|\Psi\rangle_{LG} = \sum_{\ell_1=-15}^{15} \sum_{p_1=0}^{33} \sum_{\ell_2=-15}^{15} \sum_{p_2=0}^{33} c_{p_1, p_2}^{\ell_1, \ell_2} |LG_{p_1, p_2}^{\ell_1, \ell_2}\rangle, \quad (3.6)$$

yielding a total number of 1,046,529 complex amplitudes.

To perform projection measurements on each photon of the pair, we separate the two photons with a 90° prism and image the plane of the crystal to a Holoeye spatial light modulator (SLM) with a magnification of 10. On the SLM, we display two holograms

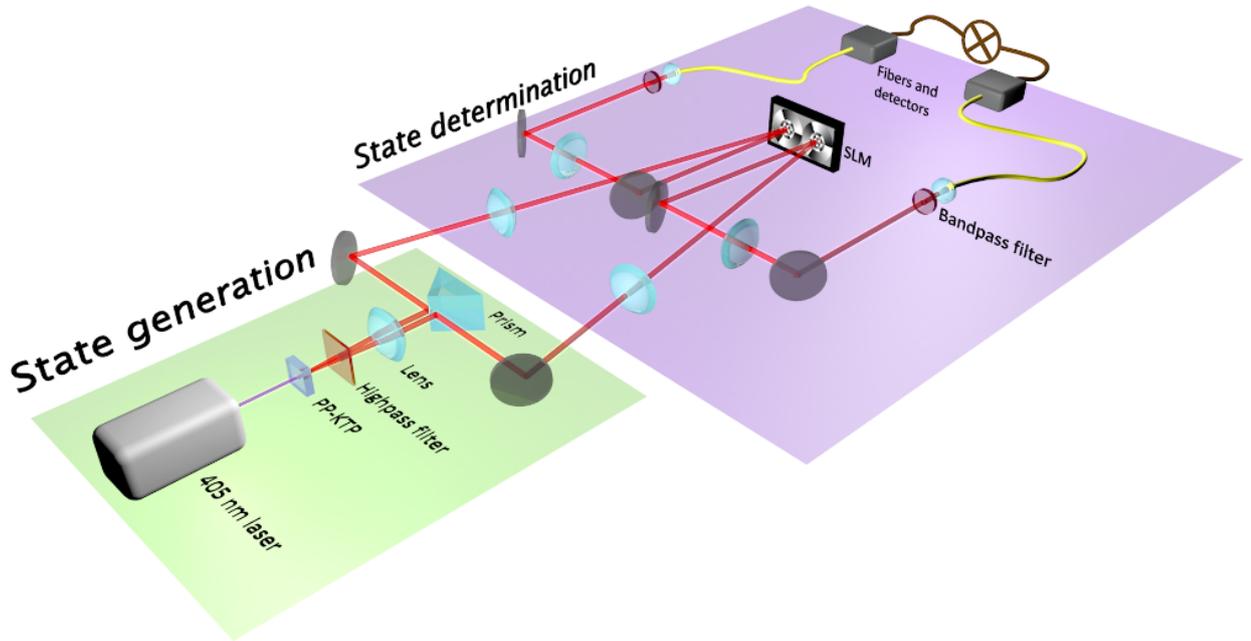


FIGURE 3.1: **Experimental setup used for the direct measurement of the high-dimensional spatial mode of the entangled SPDC state.**

corresponding to the desired projectors simultaneously, one on each side of the SLM for each photon of a pair, to control their amplitude and phase profiles independently. We make use of intensity masking when we display superposition modes [20]. We then image the plane of the SLM with a magnification of 0.0004 to two single mode fibres in order to project the joint mode in the fundamental one. All measurements are performed in coincidence with two single photon avalanche detectors, with a timing window of 25 ns and an integration time between 1 and 3 seconds. The count rate of the fundamental mode is 850 coincidences per second, and the threshold for measuring a complex coefficient is set to around 20 coincidences per second.

We show in Fig. 3.2 the results of the preliminary measurement. Coincidence counts were measured for one month and we retrieve both the amplitude and phase of the complex coefficients. We computed only the complex coefficients close to the diagonal, as the others are expected to be near zero. This measurement allowed us to identify that the dimensionality used for the measurement did not encompass the complete state: as evidenced in Fig. 3.2c, the measured amplitudes do not reach a value close to zero as we get to the end of the radial index p range. To correct for this, we changed the magnification between the crystal and the fibre to ensure that we capture the complete state. For the next experiment, we thus choose an azimuthal index ℓ ranging from -15 to 15 and the radial index p ranging from 0 to 10:

$$|\Psi\rangle_{LG} = \sum_{\ell_1=-15}^{15} \sum_{p_1=0}^{10} \sum_{\ell_2=-15}^{15} \sum_{p_2=0}^{10} c_{p_1, p_2}^{\ell_1, \ell_2} |LG_{p_1, p_2}^{\ell_1, \ell_2}\rangle. \quad (3.7)$$

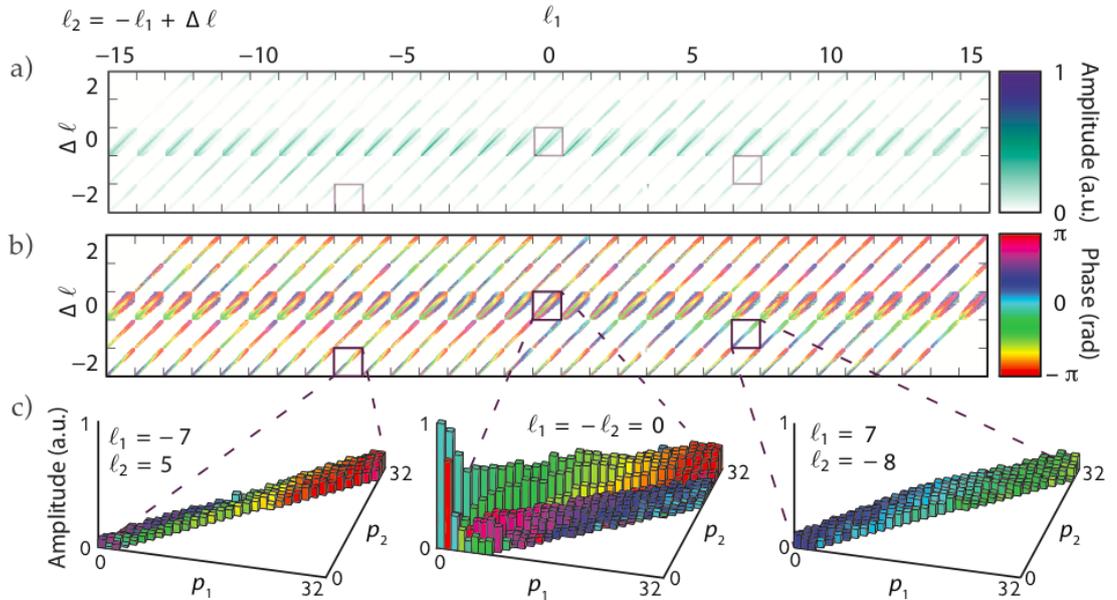


FIGURE 3.2: **Measured complex coefficients of the SPDC spatial field in the LG basis.** This measurement took one month of continuous coincidence detection, but the complex coefficients were constructed online, that is as the experiment was running.

Our final result is shown in Fig. 3.3. From this measurement, we retrieve all the complex amplitudes of the measured state. We find that the probabilities of finding the state in a given eigenstate of the LG basis given by our reconstruction are consistent with the directly measured probabilities using simple eigenprojectors (see Fig. 3.4). In the following subsections, we discuss the validity of the obtained result and detail some interesting computations that are made available from the knowledge of all complex amplitudes.

3.3.1 Comparison with assumption-free tomography

We now need to ascertain the validity of the quantum state measured through column operators. To do so, we compare a small subset of the resulting state-vector to the corresponding density matrix measured with another method.

To perform the density matrix recovery, we tried two existing methods: 1) the Cholesky decomposition method [30] and 2) the diluted maximum-likelihood algorithm [33]. We aimed to recover a 5×5 -dimensional density matrix, but we were unsuccessful with both of the above techniques. The first optimisation method kept falling in local minima, and we never managed to make the second converge. We were thus forced to construct a new algorithm for the sole purpose of benchmarking the fidelity of the state-vector with the density matrix. We briefly present this algorithm in the following few paragraphs, and we fully detail its implementation and explore its capabilities in Chapter 4.

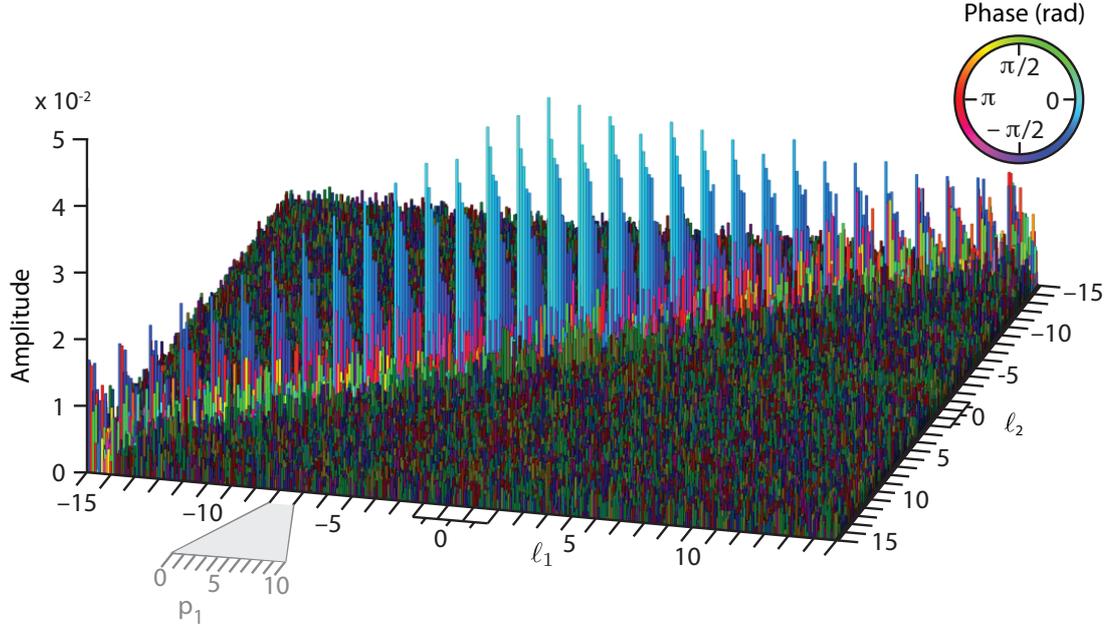


FIGURE 3.3: **Measurement of the probability amplitude of the SPDC transverse spatial mode in the LG basis.** Between two integer values of ℓ_1 (or two integer values of ℓ_2), the radial index p_1 (p_2) ranges from 0 to 10. The amplitudes are constructed using Eq. 3.5, the elements of which are coincidence frequencies. The total time for measuring the SPDC state in the LG basis is approximately two weeks.

The full tomography problem that we wish to solve consists of minimising the objective function

$$f(\rho) = \left\| \frac{\mathcal{A}(\rho) - b}{\sqrt{\mathcal{A}(\rho)}} \right\|_2^2 \quad (3.8)$$

with the constraint that the density matrix ρ must be positive semi-definite, that is its eigenvalues must either be positive or null. Concretely, b is the vector of experimentally obtained counts, and $\mathcal{A}(\rho)$ is the vector of estimated probabilities: element k of $\mathcal{A}(\rho)$ corresponds to an estimated outcome $\text{Tr}(\hat{\Pi}_k \rho)$ for a given estimate of ρ and measurement operator $\hat{\Pi}_k$. The division is performed element-wise, and maximum likelihood analysis is used to derive the term inside the 2-norm [30]. To approximate the Poisson noise associated with our experiment, we use in the maximum likelihood analysis a Gaussian noise profile with a standard deviation equal to the square-root of the counts. This approximation proves accurate in the case of high counts.

To minimise Eq. 3.8, we use a similar full tomography algorithm to the ones presented in Chapter 4. This algorithm borrows concepts from both the optimisation field and the machine learning field. We express the density matrix as a vector x of length d^2 . We express the N measurement operators in a similar way, writing them as vectors and stacking them in a $(N \times d^2)$ -dimensional matrix A . These measurement operators are random projectors that lie in a two-dimensional subspace of the Hilbert space, because these operators are easy to implement with an SLM. To find a solution, the algorithm relies on the

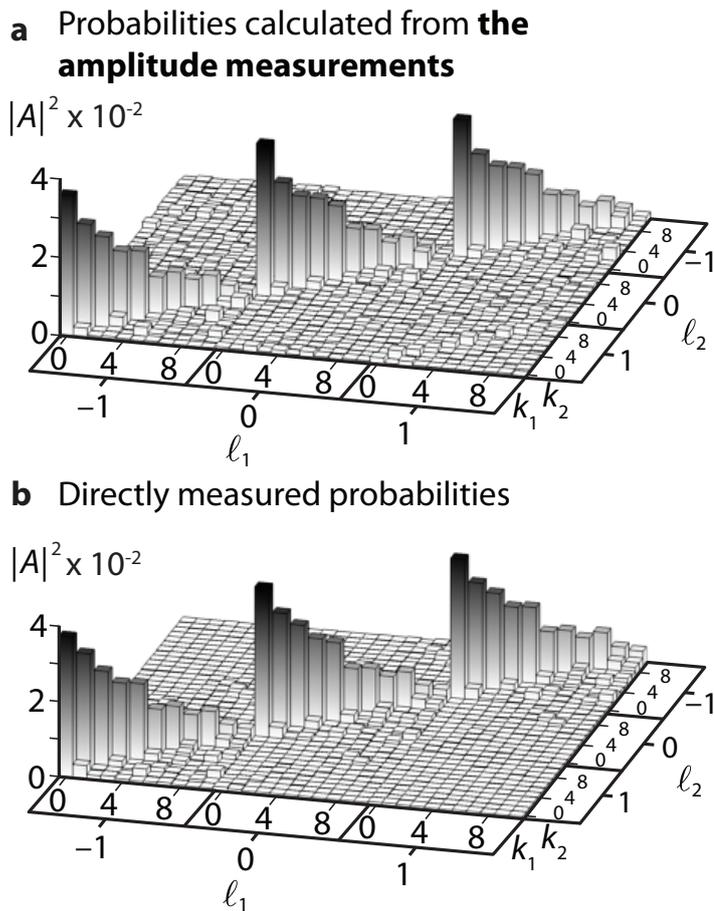


FIGURE 3.4: **Normalised probability of occurrence of some of the joint modes**. These 3d plots should be identical in theory: a) the first is calculated from the probability amplitudes retrieved using our method and b) the second is measured by setting the SLM holograms directly as the joint modes and recording the counts. The probabilities as measured in the two different ways are close enough to conclude that it is a positive sanity check.

gradient of the objective function, given by:

$$\nabla f(x) = 2A^\dagger \left[\left(\frac{Ax - b}{Ax} \right) \left(2 - \frac{Ax - b}{Ax} \right) \right]. \quad (3.9)$$

In this formulation, vector-vector operations are performed element-wise and matrix-vector operations are performed normally. We then find the density matrix that minimises Eq. 3.8 iteratively using the following algorithm:

$$x_k \leftarrow \mathcal{P}(x_{k-1} - \gamma \nabla f(x_{k-1})), \quad (3.10)$$

where γ is a learning rate and has a value between 0 and 1. The operator $\mathcal{P}(\cdot)$ represents a projection of the density matrix estimate onto the subset of positive semi-definite matrices. This projection is performed by replacing all negative eigenvalues by zero. To reduce the number of iterations required to find the solution to the constrained problem, we use as a

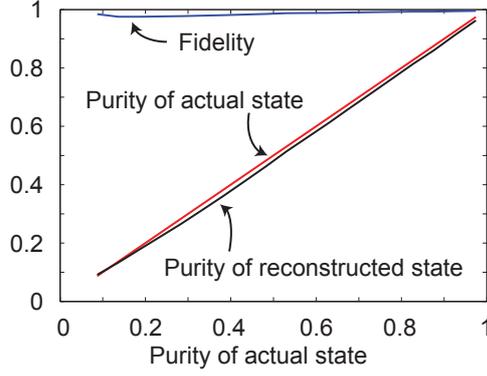


FIGURE 3.5: **Simulation of the full tomography procedure.** This simulation is a sanity check to prove that the benchmark algorithm is correct. We perform 5000 Monte Carlo simulations on and (5×5) -dimensional density matrices of varying purity. We use Poisson noise in the simulation and find an average error on the count rate of 2%. We use a projected gradient descent algorithm similar to the ones described in Chapter 4 to recover the density matrix. The fidelity between the true and recovered state, shown in blue, is (0.986 ± 0.007) on average. The uncertainty is computed as the standard deviation over all 5000 simulation runs. We use the following parameter in the simulations: $N = 1000$, $\gamma = 0.3$ and 10000 iterations.

start point for the algorithm the solution x we find for the unconstrained problem, where x does not represent a physically realisable density matrix,

$$x_0 = (A^\dagger A + I/10)^{-1} A^\dagger b. \quad (3.11)$$

Note that x is not a physically realisable density matrix as it is the solution to the unconstrained problem. The identity term in Eq. 3.11 is used as a regularisation term to minimise the influence of the noise on the inverse problem solution x .

Unlike our direct reconstruction method, this algorithm makes no assumption on the purity of the state at hand. Without going into an exhaustive theoretical analysis, we can provide evidence of the algorithm's validity by performing Monte Carlo simulations. We here perform simulations on 5000 quantum states of random purity, and show the results in Fig. 3.5. The eigenvalues of the random density matrices are fixed and set to be exponentially decreasing. We set an average number of counts of 2000, with a noise level $(\sum_i \sqrt{b_i}/b_i)/N$ of 2%. Our results show that the fidelity $\text{Tr}[\sqrt{\sqrt{\rho}\rho_{\text{act}}\sqrt{\rho}}]$ between the true density matrix ρ_{true} and the recovered density matrix ρ is 0.98 on average.

Having verified the validity of the algorithm on simulations, we now apply this projected gradient algorithm to the experimental data. We find that the state recovered using this full tomography algorithm has a fidelity with the directly measured state of (0.985 ± 0.004) , as shown in Fig. 3.6. In this case, the average purity of the state is found to be (0.96 ± 0.02) . There is thus a high consistency between the results obtained using our direct measurement approach and using the full tomography method. We also note that the directly measured state has a fidelity of (0.992 ± 0.002) with the primary

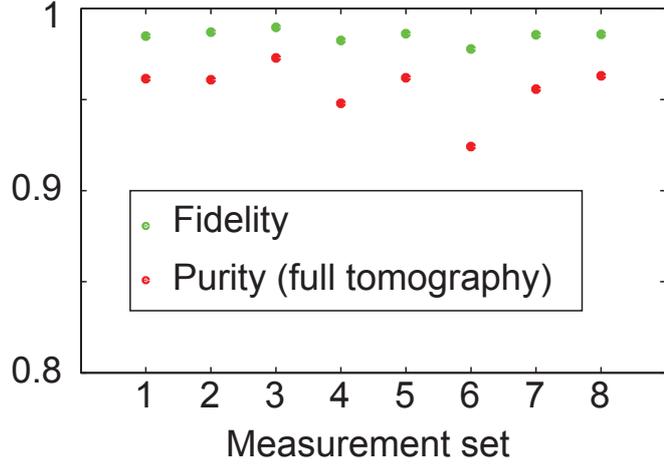


FIGURE 3.6: **Comparison of direct approach and full tomography.** We perform the data processing using the direct approach once on a (5×5) -dimension subset in the LG basis. We perform the full tomography procedure eight times on the same subset but with different random measurement operators. The purity of the recovered density matrices using the full tomography is shown in red. The average fidelity between the states obtained with the direct and full tomography methods (data points in green) is (0.986 ± 0.007) .

eigenvector of the density matrix recovered with full tomography.

Generating random state vectors

In this and the next chapter, we need to test state vector recovery algorithms on virtual quantum states. As such, we need to uniformly sample state vectors from the Hilbert space, in the Haar sense. This is done by producing randomly generated unitary matrices whose eigenvalue phase angles are uniformly distributed [63]. More explicitly, such a unitary matrix would have the following property: if $\lambda_{U,i}$, where i ranges from 0 to $d - 1$, are the eigenvalues of the unitary matrix U_{uniform} , the phase angle of an eigenvalue is defined by $\theta_i = \arctan[\text{Im}(\lambda_{U,i})/\text{Re}(\lambda_{U,i})]$. These phase angles θ_i must be uniformly distributed between $-\pi$ and π .

We can generate a unitary matrix with this property using a simple algorithm [63]. Firstly, we generate a random matrix M_{source} whose elements have amplitudes and phase angles that are sampled from a uniform distribution between 0 and 1 and between 0 and 2π , respectively. Secondly, we perform an orthogonal-triangular decomposition on this matrix: $M_{\text{source}} = QR$, where Q is a unitary matrix and R is an upper-triangular matrix. Thirdly, we extract the sign of the diagonal elements of the upper-triangular matrix R and form a new diagonal matrix D with these signs. Finally, the uniformly sampled unitary matrix is given by

$$U = QD. \quad (3.12)$$

By definition, the columns of U form a basis, i.e. a set of d orthogonal states $UU^\dagger = I$. As a consequence of the above algorithm, this basis is randomly sampled in the Hilbert

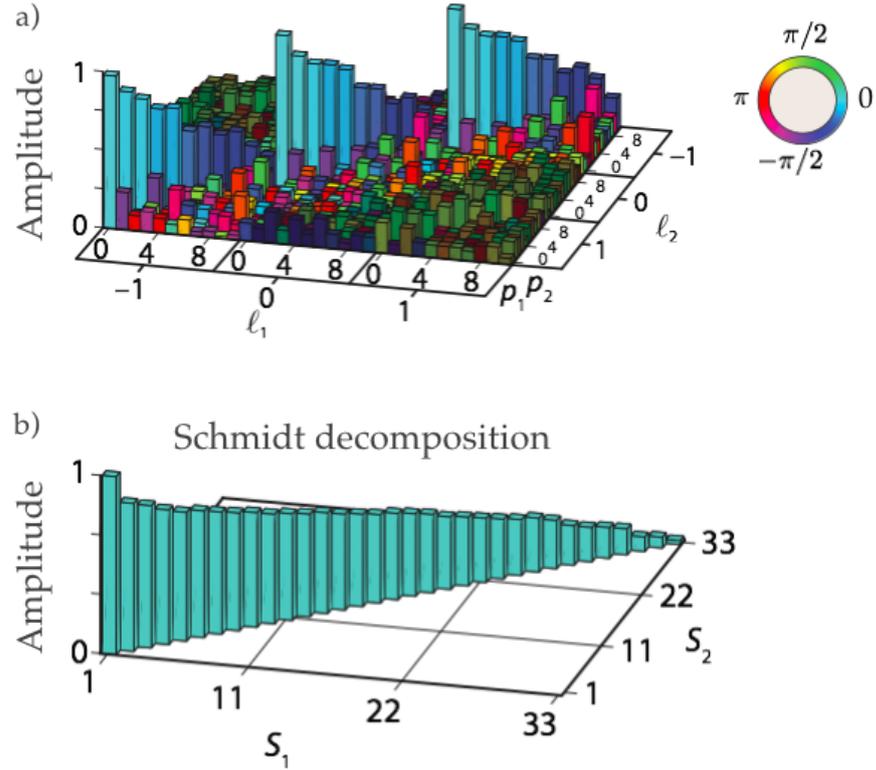


FIGURE 3.7: **Schmidt decomposition of the measured state.** a) Subset of Fig. 3.3 where the azimuthal indices range from -1 to 1, and the radial indices range from 0 to 10. The amplitudes are not normalised. From a), we compute b) its Schmidt decomposition with Eq. 3.13. The diagonal elements of b) are what would be called “singular values” of a) in linear algebra.

space such that it can be used in a simulation of state vector recovery without inducing any kind of bias in the results.

3.3.2 Application to the computation of the Schmidt decomposition

Our experiments are performed using the LG basis to express the entangled two-photon state. If we were to use this state basis for communication, errors would occur anytime photon 1 and photon 2 are detected in modes $|\ell_1, k_1\rangle$ and $|\ell_2, k_2\rangle$, respectively, with $\ell_1 \neq -\ell_2$ or $k_1 \neq k_2$ [58]. In other words, all off-diagonal elements of the amplitude matrix shown in Fig. 3.3 would register as errors, with the probability of getting an error in the communication being equal to the sum of the off-diagonal probabilities. It is theoretically possible to remove any error by finding a basis $\{|S_1, S_2\rangle\}$ that yields perfect correlations [64], using the Schmidt decomposition. In the Schmidt basis, the state becomes:

$$|\Phi_S\rangle = \sum_i \sqrt{\lambda_i} |S_{1,i}, S_{2,i}\rangle, \quad (3.13)$$

where λ_i are the probabilities associated with the Schmidt modes $|S_{1,i}, S_{2,i}\rangle$. The Schmidt basis is completely equivalent to the unitary matrices found through a singular value decomposition (SVD). In other words, if we call M the amplitude matrix from Fig. 3.7 a),

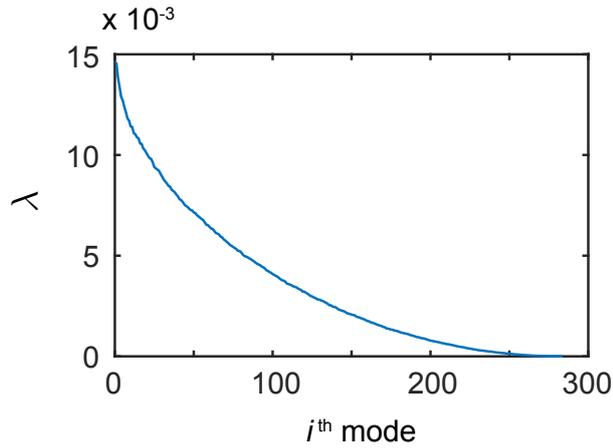


FIGURE 3.8: **Probability of occurrence of the Schmidt modes** The Schmidt modes and their associated probabilities are calculated from the entire directly measured SPDC state vector. The Schmidt number is given by $1/\sum_i \lambda_i^2$ and, in this case, equals to 142 [65].

then this matrix can be decomposed as

$$M = USV^\dagger, \quad (3.14)$$

where S is a diagonal matrix of real singular values (the Schmidt matrix), U is the basis in which we detect photon 1 and V is the basis in which we detect photon 2. Since S is diagonal, photon 1 and 2 are perfectly correlated in this joint basis. The extent of the diagonal is a good metric of how much entanglement is present.

We perform a singular value decomposition on the directly measured state to find numerically the Schmidt modes. Fig. 3.8 shows the probability of occurrence λ_i of the Schmidt modes, and the first ten Schmidt modes of the (33×33) subset from Fig. 3.7b are shown in Fig. 3.9.

3.3.3 Application to Bell inequalities in high dimensions

As an example of a quantum mechanical calculation that can be performed by knowing the state-vector, we verify that the state measured in the LG basis violates the Bell inequalities. A state that violates a Bell inequality is fundamentally interesting because it indicates strong entanglement and non-locality of quantum mechanics [39, 66]. From a more practical point of view, violation of the Bell inequalities with entangled particles certifies that a third party is not eavesdropping, hence guaranteeing the security of a quantum channel in a quantum cryptography protocol [67]. Bell inequality tests have been applied to up to (14×14) -dimensional systems [40].

Here, we apply the d -dimensional Bell operators to d -dimensional subsets of the measured state in the LG basis; see blue and red points in Fig. 3.10. We then make the Schmidt decomposition of the measured state to ensure that the state is expressed in a

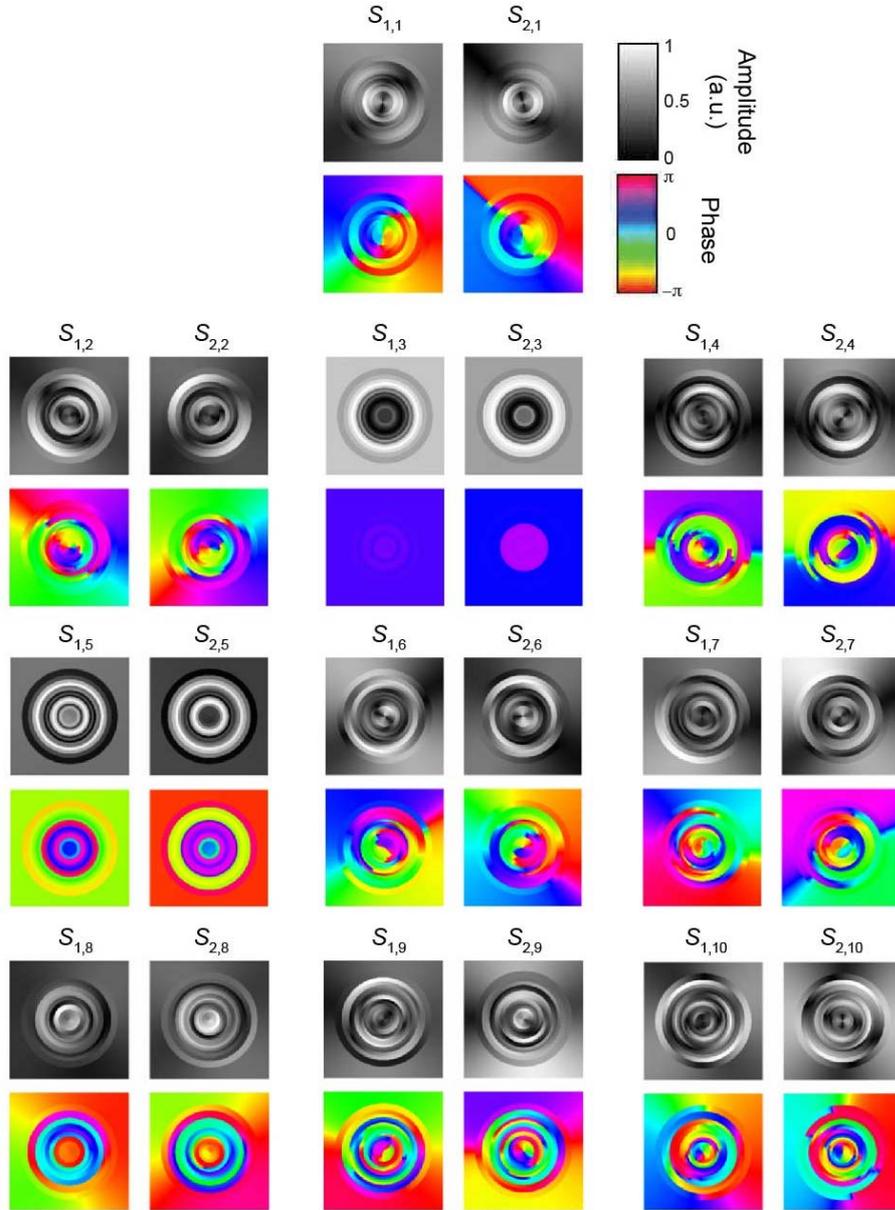


FIGURE 3.9: **Amplitude and phase profiles of the first ten Schmidt modes.** We first process the experimental data to reduce the statistical noise contribution, inherent to the measurements: we set the value of coefficients with amplitude below 0.04 to zero. The Schmidt modes forms a basis in which we would theoretically find the photons in the joint photon states as being highly correlated. Parameters such as the size of the pump beam on the crystal, the alignment of the setup, optical aberrations and the detection geometry all affect the specific form of the Schmidt modes.

basis that has the property of being maximally correlated (see subsection 3.3.2) [64]. We apply the Bell operators on the measured state expressed in this new basis and find that the resulting Bell values approaches the maximum possible violation called the Tsirelson bound [41]. We simulate the Bell test with the measured SPDC state in a few hours, the result of which is shown in Fig. 3.10. Because we have certified that the state that we measured is quasi-pure, the results of the Bell inequality test strongly suggest that the generated state is violating Bell inequalities for dimensions up to (100×100) .

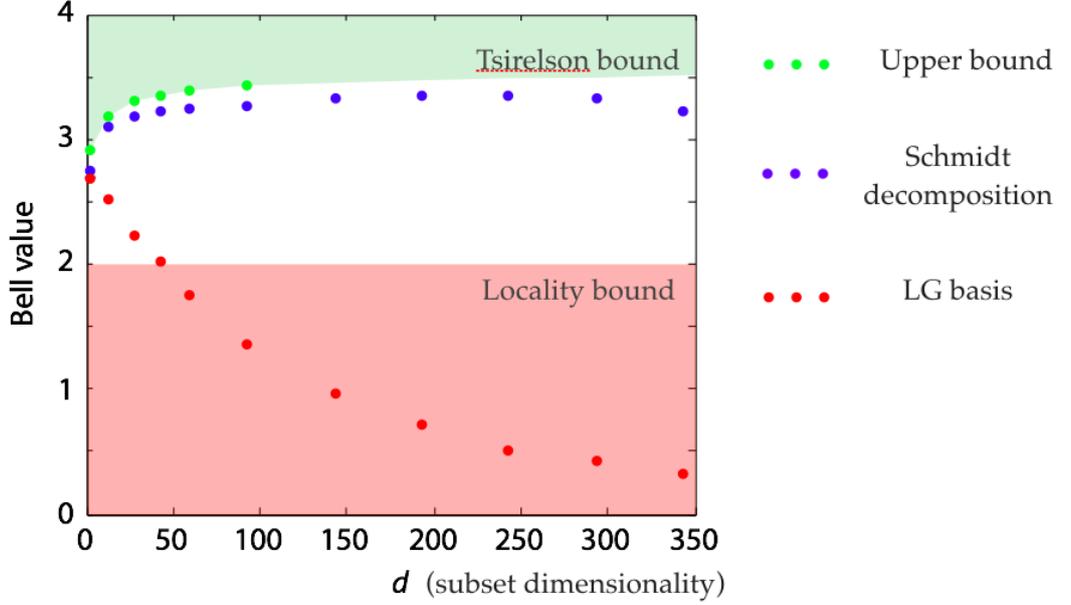


FIGURE 3.10: **Theoretical Bell analysis applied on the measured state.** We simulate a Bell test on d -dimensional subsets of the measured state. When the numerical result S of a Bell test is above 2, the state proves secure for high-dimensional quantum key distribution. Tsirelson’s bound is the theoretical maximum outcome of a Bell test, computed from a generalisation of the Bell inequalities in any dimension [41]. The bound was calculated for dimensions up to $d = 100$ (green points), and the boundary of the green and white regions corresponds to an estimate of the bound in higher dimensions.

3.3.4 Purity of the spatial part of the SPDC state

Our method was developed under the assumption of working with a pure state, so that the photon pair we measure needs to be pure for the method to yield accurate results. However, the photon frequencies Ω_1 and Ω_2 are also found to be correlated in a SPDC two-photon field. For the transverse spatial part of the state to be pure, it needs to be uncorrelated to the frequency or energy part of the state. Thus, the spatial part of the state will be pure only if it can be written in separable form:

$$\begin{aligned}
|\Phi'\rangle &= \int d\mathbf{k}_1 d\Omega_1 d\mathbf{k}_2 d\Omega_2 \Phi'(\mathbf{k}_1, \Omega_1, \mathbf{k}_2, \Omega_2) |\mathbf{k}_1, \Omega_1, \mathbf{k}_2, \Omega_2\rangle \\
&= \int d\mathbf{k}_1 d\Omega_1 d\mathbf{k}_2 d\Omega_2 H(\mathbf{k}_1, \mathbf{k}_2) G(\Omega_1, \Omega_2) |\mathbf{k}_1, \Omega_1, \mathbf{k}_2, \Omega_2\rangle, \quad (3.15)
\end{aligned}$$

where k_i is the transverse wave vector of photon i .

To obtain a state that is separable in the spatial and frequency parts, one can project the frequency of each photon into very narrow bandwidth states. However, this has the undesired effect of lowering the count rates. A more efficient way to obtain a separable state is to work in the collinear regime [68]. To do so, a 50/50 beam-splitter is used to separate the photons and the collection modes are centered in the SPDC field. In this case, there is still a loss of 50% arising from the use of a beam-splitter. This loss can be circumvented by using a prism rather than a beam-splitter to separate the photons, and by making sure the collection modes are as close to the center as possible. We use the formalism of References [68,69] to calculate the purity of the entangled two-photon state vector. One parameter that affects the purity of the state is the average angle ϕ at which the photons leave the crystal, with respect to the pump's optical axis. Fig. 3.11 shows the dependency between the purity and the deviation angle ϕ . Based on this result and our experimental parameters, we estimated that we experimentally generate a state with a purity greater than 0.99. Although we did not confirm this result in a dedicated experiment, it would be possible to do so using a method that relies on non-local measurements to directly measure the purity of a two-photon state [70].

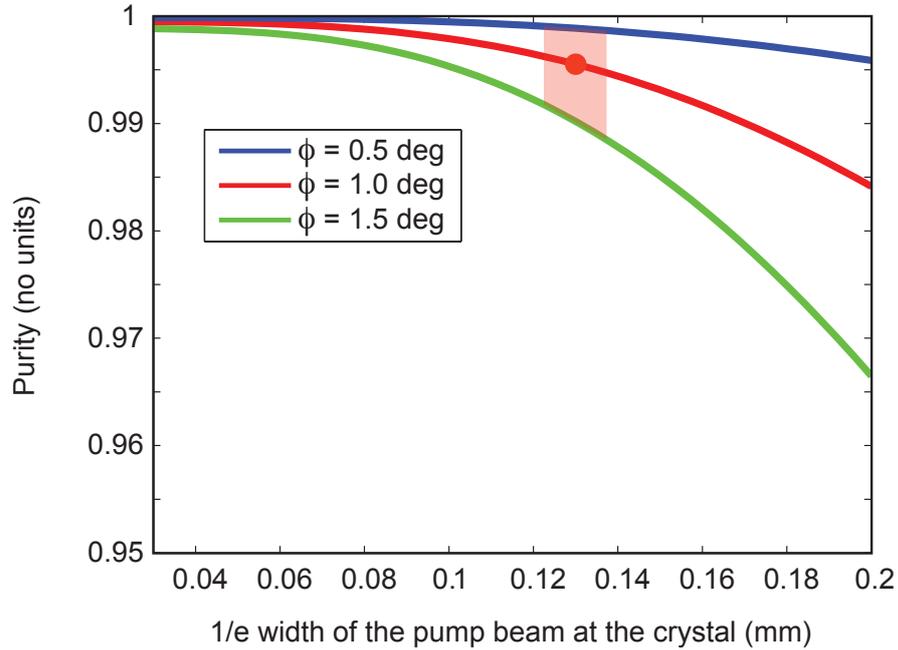


FIGURE 3.11: **Purity of the spatial part of the two-photon state.** We use the following experimental parameters: pump beam width at the plane of the crystal = 0.13 ± 0.01 mm; collection angles for both collection modes = $\phi = 1.0 \pm 0.5$ degree; pump spectrum bandwidth at $1/e = 0.25$ nm; spectral filters bandwidths at $1/e = 4$ nm; collection modes widths at $1/e = 200$ μm . For simplicity, we assume the spectrum of the pump and filters and the collection modes are Gaussian-distributed. Our experimental conditions are highlighted by the red-coloured region, where the purity is most likely between 0.99 and 1. This calculation is in line with the formalism of Reference [68].

3.3.5 Non-unit purity

In the case where our experimental setup does not provide a pure state, our direct measurement procedure can still yield valuable information for a large range of parameters. To study this regime, we perform Monte Carlo simulations on various types of density matrices. For a specific simulation run, we find the vector $|j\rangle$ out of the computational basis $\{|1\rangle, |2\rangle, \dots, |d-1\rangle\}$ that yields the highest number of counts, when simulating the measurement in this basis. We use this vector as the reference vector $|\nu\rangle$. By discarding the shot noise contribution, we can directly take the normalised j^{th} column of the density matrix to obtain the direct measurement outcomes. We then compute the primary eigenvector of the density matrix, in order to calculate its fidelity with the simulated outcomes. To do so, we write the density matrix as

$$\rho = \lambda_0 |q_0\rangle\langle q_0| + \sum_{k=1}^{r-1} \frac{1 - \lambda_0}{r - 1} |q_k\rangle\langle q_k| \quad (3.16)$$

where λ_0 is the primary eigenvalue, $|q_0\rangle$ the primary eigenvector, and $r > 1$ the rank of the density matrix. The purity of the density matrix is

$$\text{Tr}[\rho^2] = \lambda_0^2 + (1 - \lambda_0)^2 / (r - 1). \quad (3.17)$$

We use random density matrices of varying purities and ranks in our simulations, with the obtained results shown in Fig. 3.12. We find the lowest fidelity between the primary eigenvector and the recovered state when the density matrix is of rank $r = 2$, representing a system composed of a mixture of two pure states. If there are a higher number of eigenvectors, with the eigenvalues evenly distributed amongst them, the results of the direct measurement approach more reliably correspond to the primary eigenvector. In particular, the difference between the primary and secondary (second highest) eigenvalue, or spectral gap, is the parameter that most influences the quality of the recovered state. When this gap is high enough, we recover the primary eigenvector with high fidelity. In particular, we obtain a fidelity higher than 0.99, in the absence of noise, for states of purity higher than 0.81. Given the form in which we express the density matrix in Eq. 3.16, the spectral gap increases as the rank increases, which explains our findings that the fidelity increases with rank.

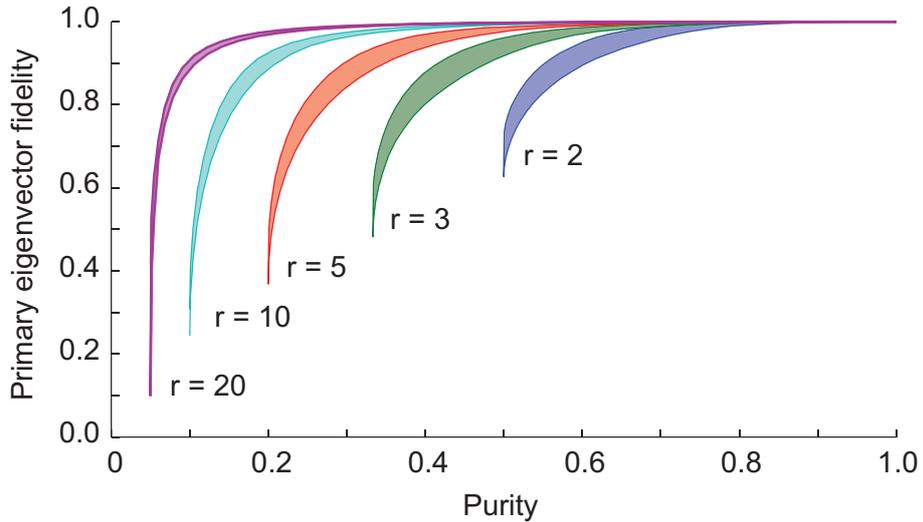


FIGURE 3.12: **Primary eigenvector fidelity with the direct measurement approach recovered state.** We simulate random density matrices of varying purity and apply our direct measurement procedure to recover the state. For each value of purity, we perform the simulation 100 times to compute the average fidelity between the directly recovered state and the primary eigenvector of the true density matrix. The hard colour lines are drawn at one standard deviation from the mean, so that the fidelity is most likely contained within the shaded areas. We show the results for ranks $r = \{2$ (blue), 3 (green), 5 (red), 10 (cyan), 20 (purple) $\}$ of the density matrix. We perform these simulations on a system of dimension $d = 40$, but find that the result shows negligible variations for other valid dimensions ($d \geq r$).

3.4 Discussion

The column operator formalism resembles that of weak-value assisted tomography described in reference [50]. Lundeen *et al.* show that, if properly chosen, weak values can be proportional to the expansion coefficients of the state-vector, such that weak values can be used to operationally define the state-vector. Weak-value tomography aims at retrieving a state that can be expressed in the same way as used in our method (Eq. 3.1):

$$\Psi = \sum_{j=1}^d c_j |j\rangle. \quad (3.18)$$

Rather than expressing the state using a combination of column-vector expectation values that can be constructed from strong measurements (Eq. 3.4), the weak-value method is based on the measurements of weak values of projectors, defined as

$$\langle \hat{A}_j \rangle_{weak} = \frac{\langle f|j\rangle \langle j|\Psi\rangle}{\langle f|\Psi\rangle} \quad (3.19)$$

To perform such a measurement, the state Ψ is first slightly perturbed in the basis of states $|j\rangle$, then projected onto a state $|f\rangle$ of a conjugate basis. This procedure leads to a coupling of the basis $|j\rangle$ onto an independent degree of freedom, over which the measurements must be performed. Once the weak values associated with each projector $|j\rangle$ have been constructed from a set of measurements (similarly to the way column-operator expectation values are constructed from strong projections), the state is then

retrieved as

$$\Psi = \nu^{-1} \sum_{j=1}^d \langle \hat{A}_j \rangle_{weak} |j\rangle, \quad (3.20)$$

where ν is a normalisation constant. Compared to our direct measurement method, the experimental procedure required to retrieve weak values is somewhat involved [52, 71] and always yields an approximation to the weak values, even in the limit of infinite signal [72].

Our work is the result of an effort to simplify this experimental procedure and find an exact solution, such that it could easily be generalized and applied to characterize quantum states composed of more than one system and to reliably monitor quantum dynamics. We lay the foundations of a method that i) exactly yields the coefficients of the state-vector and ii) only requires strong projective measurements, the most basic and widely used type of measurement in quantum mechanics. We can thus operationally define the quantum state coefficient as the expectation value of the corresponding column operator, a definition that holds for many-body systems.

One critique of the column operator approach as described above is that is not performed with generalised measurements, i.e. positive-operator valued measure (POVM). However, it would be possible to combine the direct measurement approach with the efficient measurement procedures associated with the use of POVMs. A simple procedure, described in Reference [48], allows one to transform any informationally complete set of N projection operators $\{P_1, P_2, \dots, P_N\}$ into a POVM. The projectors transformed as follows form a POVM:

$$G^{-1/2} P_k G^{-1/2} \quad (3.21)$$

$$\text{with } G = \sum_k P_k \quad (3.22)$$

The informationally complete set of projectors used in our experiment can be turned into a POVM, from which we can reconstruct the probability amplitudes associated with a POVM-based measurement, thus confirming that we can apply the direct measurement procedure using a POVM. As the technology to perform such a measurement in the spatial degree of freedom is not yet available, we did not implement this approach in our experiments.

3.5 Conclusion

As the dimensionality of quantum systems increases with research and technology, it is important that the measurement techniques are up to speed with the demanding quantum protocols. Not only does pure-state tomography allow access to large scale quantum systems, but it requires trivial data processing, especially compared to maximum likelihood estimation [29]. For these reasons, pure-state tomography will prove to be a useful asset in the quest for understanding and manipulating very-large-scale quantum systems.

Chapter 4 – Tomography through gradient descent

This work is a continuation of Chapter 3 in that, after completing the direct measurement work, I could not find a proper algorithm for benchmarking the direct measurement method. I had to develop my own algorithm for reconstructing general quantum states of large dimensions. The existing algorithms were too slow, so I combined ideas from the machine learning field (gradient descent with momentum) and the optimization field (projected gradient descent) to be able to benchmark the outcome of the direct measurement procedure with a general quantum state. Later, George Knee, Erik Gauger and Jonathan Leach all helped me in refining and characterising the projected gradient descent algorithms. During the characterisation process, we became aware of similar work published as we were working on this problem [73]. Our manuscript is accepted by Quantum Information (Nature Publishing Group).

The inverse problem addressed in this chapter is similar to that of Chapter 3, except that we now allow the quantum state to be a mixture of many pure states. Instead of having a search space that is the space of all state-vectors, we now have search space encompassing all the density matrices, which have unit trace and positive (or null) eigenvalues.

4.1 Introduction

As we discussed in the above chapters, the problem of quantum tomography is the inverse problem of reconstructing an unknown quantum state from the outcomes of a well-defined set of measurement on a large number of identically prepared states. Situations where quantum tomography may be required in experimental procedures include the need to demonstrate or verify the existence of complex quantum states, such as highly entangled states, states that violate a Bell inequality or even the initial state required in a gate-based quantum computer. These type of states are of great use in the fields of quantum-based communication, metrology and computation. Since complex states are of increasing interest for the experimentalists, the experimental methods to generate and measure larger and more complex states have greatly progressed in recent years. In particular, quantum control of a high number of qubits has been demonstrated by various groups [74–76].

As the dimensionality of the quantum states exploited in experiments increases, it becomes harder to retrieve a quantum state from a set of measurement outcomes in a timely manner. The challenge of reconstruction arises from the unfavourable scaling of the prob-

lem with large states: in the case of qubits, the number of real parameters required to describe the joint state of n qubits scales as 2^{2n} . Even for single states with a single degree of freedom, one can face the challenging task of reconstructing a state of large dimensionality: single photons characterised by their orbital angular momentum, for example, can be represented with a very large dimensionality, and we have shown in Chapter 3 that it can be characterised as a 100,000 dimensional state vector [38, 43, 71]. For the reconstruction methods to keep up with the increase in dimensionality in experimental techniques, one has to devise sophisticated data processing techniques.

In this chapter, we focus our efforts on the theoretical development of new reconstruction algorithms and numerical simulations to put them to the test. Our numerical simulations are performed assuming an informationally complete set of measurement. As discussed in Chapter 1, given an informationally complete set, the true quantum state ρ_{true} can be retrieved exactly if blessed with ideal conditions where no noise, including statistical noise, is present. However, our numerical simulations are carefully designed to represent the experimental conditions of quantum tomography measurements. In particular, as the dimensionality of the state grows, the number of detector counts for each state parameter decreases for a given number of total counts. In optical systems, low detection efficiencies of experimental setups also increases the noise of the measured data set. Our algorithms therefore address the problem of recovering the most likely estimate of the unknown state given a noisy data set.

Most methods used for quantum state reconstruction are also based on numerical techniques. Here, methods we use as benchmarks are the diluted iterative algorithm (DIA) [32, 33] and convex programming [35]. These two algorithms have been demonstrated to converge to the maximum likelihood state ρ_{ML} . They however present some limitations: in the case of the DIA, the algorithm has been observed to converge slowly in high dimensions [73, 77]. In the case of convex programming, existing solvers such as SDPT3 and SeDuMi are computationally expensive as they scale poorly with non-sparse matrix dimensionality [78, 79].

As an alternative to these iterative methods, a non-iterative quantum tomography method that can be used under a specific set of conditions was devised by Smolin *et al.* [80]. If using traceless measurement operators in a context where the noise is Gaussian, linear inversion can be used to find the unconstrained maximum likelihood state, from which the maximum likelihood state ρ_{ML} is retrieved in a single projection step. This algorithm was also implemented using a GPU to recover a simulated density matrix representing a 14-qubit system [81]. However, this technique is limited to the conditions stated above, which motivates our search for efficient techniques that are applicable to more general tomography problems.

In this work, we develop and study methods based on projected gradient descent (PGD) as efficient techniques to solve quantum tomography problems. PGD is a technique widely known and used in optimisation theory, but was only recently applied to the context of recovering density matrices [37, 73]. The recent work of Goncalves *et al.* showed that using projected gradient descent with backtracking (PGDB) for quantum reconstruction is a promising avenue, but a deeper analysis is required to fully assess the performance of PGD techniques for quantum tomography, in particular on data sets representing practical measurements. In addition to the approach of PGDB of Reference [73], we present in this chapter two new PGD algorithms for quantum state reconstruction: projected gradient descent with momentum (PGDM) and Fast Iterative Shrinkage-Thresholding Algorithm (FISTA) [82]. To the best of our knowledge, PGDM is a novel algorithm, and FISTA has never been applied to quantum tomography.

This chapter is structured as follows: we begin the theoretical development by detailing the measurement and problem formalism. We then define the three PGD algorithms and discuss their performance by studying their convergence profiles and running time and benchmarking them against DIA and SDPT3. Finally, we show the results of state reconstruction using simulations of realistic quantum tomography experiments that include noise. We run the algorithms multiple times on density matrices of fixed purity generated randomly over a broad range of Hilbert space dimensions [83]. We record for each numerical experiment the running times and fidelities between the recovered state and the actual one ρ_{true} , and use as a figure of merit the time taken by the algorithm to converge sufficiently close to the maximum likelihood state.

4.2 Theoretical development

We detail in this section the quantum state reconstruction problem that we tackle, and the proposed projected gradient descent algorithms to efficiently solve this problem.

4.2.1 Problem formulation

We consider here a general quantum state, described by the density matrix ρ . In the case of a d -dimensional state, ρ is a $d \times d$ positive semidefinite matrix with unity trace. The probability p_i of obtaining a measurement outcome, given a quantum state ρ and a general Hermitian measurement operator Π_i , is given by the Born rule in its general form:

$$p_i = \text{Tr}(\Pi_i \rho). \quad (4.1)$$

The experimental outcomes are not a direct observation of the probabilities p_i , but rather a number of detection, or number of clicks, n_i as recorded by the detector after a

finite number of trials. When no noise is present, the average number of clicks is related to the probabilities via a simple multiplicative factor r :

$$n_i = rp_i. \quad (4.2)$$

In realistic conditions, there is a discrepancy between n_i and rp_i due to both inherent statistical uncertainty and noise in the measurement device. As we have seen in Chapter 1, a naive reconstruction method such as a simple linear inversion will lead to an unphysical quantum state in the presence of significant noise, thus violating fundamental physical properties such as positivity and unit trace. One therefore needs to implement steps in the algorithm that ensures physicality of the recovered state. Multiple techniques have taken the avenue of restricting the search for a state in a space of physical solutions (see section 2.5). However, since searching in a restricted space may lead to ill-conditioned problems and slow convergence, it can be advantageous to instead allow the search to wander temporarily into unphysical territory, before projecting back into a physical state. To paraphrase Tolkien, not all density matrix estimates who wander are lost.

Eq. 4.1 can be rewritten as:

$$\mathbf{p} = A \text{vec}(\rho) \quad (4.3)$$

by expressing the probabilities p_i and measurement operators into matrix and vector format, as follow:

$$A = \begin{pmatrix} \text{vec}(\hat{\Pi}_1)^T \\ \vdots \\ \text{vec}(\hat{\Pi}_N)^T \end{pmatrix}, \quad \mathbf{p} = \begin{pmatrix} p_1 \\ \vdots \\ p_N \end{pmatrix}, \quad (4.4)$$

where N is the total number of operators. The detector clicks n_i can similarly be expressed into vector format:

$$\mathbf{n} = \begin{pmatrix} n_1 \\ \vdots \\ n_N \end{pmatrix} \quad (4.5)$$

It is possible to estimate the multiplicative factor r that links \mathbf{n} and \mathbf{p} only using the measurements \mathbf{n} , in the case where at least a subset \mathcal{Z} of the measurement matrix A forms a POVM (as defined in section 2.4.1). In that case, the sum of the probabilities p_i belonging to the subset \mathcal{Z} , is independent of the state ρ :

$$\sum_{j \in \mathcal{Z}} p_j = \frac{\mathcal{Z}_N}{d} \quad (4.6)$$

where \mathcal{Z}_N is the number of projectors in the POVM. Note that in the case where the subset \mathcal{Z} forms a basis, the sum equals to unity. In any case, the multiplication factor is found to

be:

$$r = \frac{d}{N_{\mathcal{Z}}} \sum_{j \in \mathcal{Z}} n_j, \quad (4.7)$$

Moreover, the average number of clicks per outcome is r/d , and the total number of clicks for this POVM is $rN_{\mathcal{Z}}/d$.

The condition number of the measurement matrix A , as previously defined in Eq. 2.66:

$$C_{\sigma} = \frac{\sigma_{\max}}{\sigma_{\min}} \quad (4.8)$$

determines the accuracy of the maximum likelihood state [84]. Ill-conditioned measurement matrices, which have a high condition number, have been found to be used in the fields of detector tomography [85, 86] and superconducting artificial atoms [87].

The computation of the vector of probabilities \mathbf{p} (Eq. 4.3) requires in general $O(Nd^2)$ floating-point operations. The computational complexity can be lower in the case where the operators originate from tensor products [31] or outer product, as we detail in the following subsection.

Rank-1 projectors

In our simulations, we use measurement operators that are rank-1 projectors, which has the benefit of relaxing the memory requirements. Instead of having to work with matrix operators, the measurements can now be expressed as d -dimensional vectors. In this specific case where the operators are rank-1 projectors, expressed as $|\phi_i\rangle\langle\phi_i|$, we find the Born rule takes the form:

$$p_i = \langle\phi_i|\rho|\phi_i\rangle \quad (4.9)$$

and the measurement matrix is expressed as:

$$\mathcal{A} = \begin{pmatrix} \langle\phi_1| \\ \vdots \\ \langle\phi_N| \end{pmatrix}. \quad (4.10)$$

Using this notation, we express the vector of estimated probabilities

$$\mathbf{p}_i = \sum_j \mathcal{A}_{ij} \sum_v \mathcal{A}_{iv}^* \rho_{vv}. \quad (4.11)$$

The measurement matrix expressed in Eq. 4.10 is $[N \times d]$ -dimensional, while the general form of Eq. 4.4 is $[N \times (d \times d)]$ -dimensional: we thus gain a factor d in RAM memory requirement using rank-1 projectors rather than full-rank measurement operators.

4.2.2 Maximum likelihood

To evaluate the closeness between the estimated probabilities p_i and measurement outcomes n_i , we rely on maximum likelihood analysis. We use here a likelihood function defined for a multinomial experiment, which represents well the above-formulated problem:

$$\mathcal{L}(\rho) \propto \prod_i p_i^{n_i}. \quad (4.12)$$

To find the quantum state ρ_{ML} that is the most consistent with the observed data n_i , one has to maximise this function such that

$$\rho_{\text{ML}} = \max_{\rho} \mathcal{L}(\rho). \quad (4.13)$$

Alternatively to maximising the likelihood function $\mathcal{L}(\rho)$, since any scale or shift by a constant can be ignored for the maximisation process, one can also minimise the following:

$$\rho_{\text{ML}} = \min_{\rho} -\log \mathcal{L}(\rho). \quad (4.14)$$

We thus define the cost function to be minimised as

$$\mathcal{C}(\rho) = -\log \mathcal{L}(\rho) \quad (4.15)$$

This approach is the most widely used one in tomography problems [29, 30, 32, 88]. The gradient of this negative log-likelihood function is given by

$$\nabla \mathcal{C}(\rho) = -\mathcal{G}^\dagger \mathcal{A}, \quad (4.16)$$

with $\mathcal{G}_{ij} = \mathcal{A}_{ij} n_i / p_i$.

When working with a large number of trials, this cost function is well approximated by the Poisson-approximated Gaussian likelihood function:

$$\mathcal{C}_P(\rho) \approx -\log \mathcal{L}_P(\rho) = \boldsymbol{\nu}^T \boldsymbol{\nu} \quad (4.17)$$

$$\text{with } \nu_i = \frac{r p_i - n_i}{\sqrt{n_i}}. \quad (4.18)$$

This likelihood function is also a better model than the multinomial likelihood in one particular situation, namely when the number of events has shot noise within it. The multinomial model is only accurate when the number of events or the total number of detector clicks is deterministic.

In the assumption of Poisson-distributed data, the variance for outcome i is equal to the number of clicks n_i . The elements ν_i therefore correspond to the ratio of the error $(rp_i - n_i)$ to the expected error $\sqrt{n_i}$. Given the expected noise on the outcomes, the true density matrix would yield a cost function per outcome \mathcal{C}/N equal to one. A value greater than one is sign of a poor density matrix estimate or of an incomplete noise model. On the other hand, a value smaller than one indicates an estimate that overfits to the noise in the data. In general, one finds a maximum likelihood density matrix that overfits the data slightly [89]. However, a better estimate cannot be achieved in the absence of prior knowledge.

Furthermore, to compute the gradient of the cost function, we use a Gaussian likelihood function

$$\mathcal{L}_G(\rho) = \prod_i \exp \left[-\frac{(rp_i - n_i)^2}{\sigma^2} \right] \quad (4.19)$$

$$-\log \mathcal{L}_G(\rho) = \sum_i \frac{(rp_i - n_i)^2}{\sigma^2} \quad (4.20)$$

We omit the normalisation factor for the sake of simplicity. The normalisation factor has no effect on the direction of the gradient, which is what we seek. Here, the standard deviation σ on a particular measurement outcome is assumed to be the same over all measurements. This assumption is not true in general, but it is true for the particular case of Pauli measurements [80]. Even when this assumption is broken however, the resulting likelihood is not completely wrong; it is merely not optimal.

$$\frac{d \log \mathcal{L}_G(\rho)}{d\rho_{xy}} = \frac{1}{\sigma^2} \sum_i 2(rp_i - n_i)r \frac{dp_i}{d\rho_{xy}}. \quad (4.21)$$

From Eq. 4.11, we have

$$\frac{dp_i}{d\rho_{xy}} = \sum_{xy} \mathcal{A}_{iy} \mathcal{A}_{ix}^*. \quad (4.22)$$

Returning to Eq. 4.21 gives

$$\frac{d \log \mathcal{L}_G(\rho)}{d\rho_{xy}} = \frac{2r}{\sigma^2} \sum_{ixy} (rp_i - n_i) \mathcal{A}_{iy} \mathcal{A}_{ix}^*. \quad (4.23)$$

Let us define a matrix $G_{iy} = (rp_i - n_i) \mathcal{A}_{iy}$. We then have

$$\frac{d \log \mathcal{L}_G(\rho)}{d\rho_{xy}} = \frac{2r}{\sigma^2} \sum_{ixy} G_{iy} \mathcal{A}_{ix}^*, \quad (4.24)$$

which is a matrix multiplication between G and \mathcal{A} . Finally, the gradient is compactly written as

$$\nabla \log \mathcal{L}_G(\rho) = \frac{2r}{\sigma^2} G^\dagger \mathcal{A}^*, \quad (4.25)$$

This gradient is the one used in the algorithms described in the next subsection, and its computation requires $O(Nd^2)$ floating-point operations.

4.2.3 Projected Gradient Descent Algorithms

We can now detail the PGD algorithms we use for density matrix reconstruction. In all algorithms, we use the cost function of Eq. 4.17 and seek to find the best estimate of the quantum state given the observations n_i . In particular, we are looking for an estimate that is a $d \times d$ positive semi-definite matrix of unit trace. Projected gradient descent is well adapted to that sort of problem as it consists of two main steps that address both of these goals.

We start the descent from a well-chosen physical state. In this work, we always use as a starting point the completely mixed state $\rho_0 = I/d$. The first step is the gradient descent and consists of going in the downhill direction of the cost function. This step is likely to yield a nonphysical matrix. To bring the matrix estimate back into the physical space defined by the above constraints, we perform a projection step. In this second step, the estimate is projected onto its closest point in the solution space. These two steps are then repeated until the cost function converges to a given low value which we set as threshold for the algorithm. This threshold is typically close to unity, and is set so that the estimate produces estimated probabilities that are within a standard deviation of the data on average. There is very little benefit in refining the estimate past this point. Fig. 4.1 illustrates this evolution of the density matrix through gradient descent and projection steps, in the case of a qubit going through six iterations of PGD.

In this process, a recursion relation defines the density matrix at a given iteration with respect to the density matrix at the iteration before. The simplest PGD algorithm can be written as [38]

$$\rho_{k+1} = \mathcal{P}[\rho_k - \delta \nabla \mathcal{C}(\rho_k)], \quad (4.26)$$

where δ is the step size of the descent and $\mathcal{P}[\cdot]$ is the projection operator that finds the closest matrix in the solution space to the matrix in its argument. Various mathematical definitions of ‘closest’ can be implemented at this step. One of the definitions that can be used is the *Frobenius norm*, for which the projection $\mathcal{F}[\cdot]$ is derived in Reference [80]. However, we found that implementing this norm into a PGD algorithm yields suboptimal results as the estimate tends to approach the maximum likelihood state, but to then stall and never reach it. Another option would be the *nuclear norm* with projection $\mathcal{M}(\cdot)$. This projection sets all negative eigenvalues to zero; we have also observe that this type

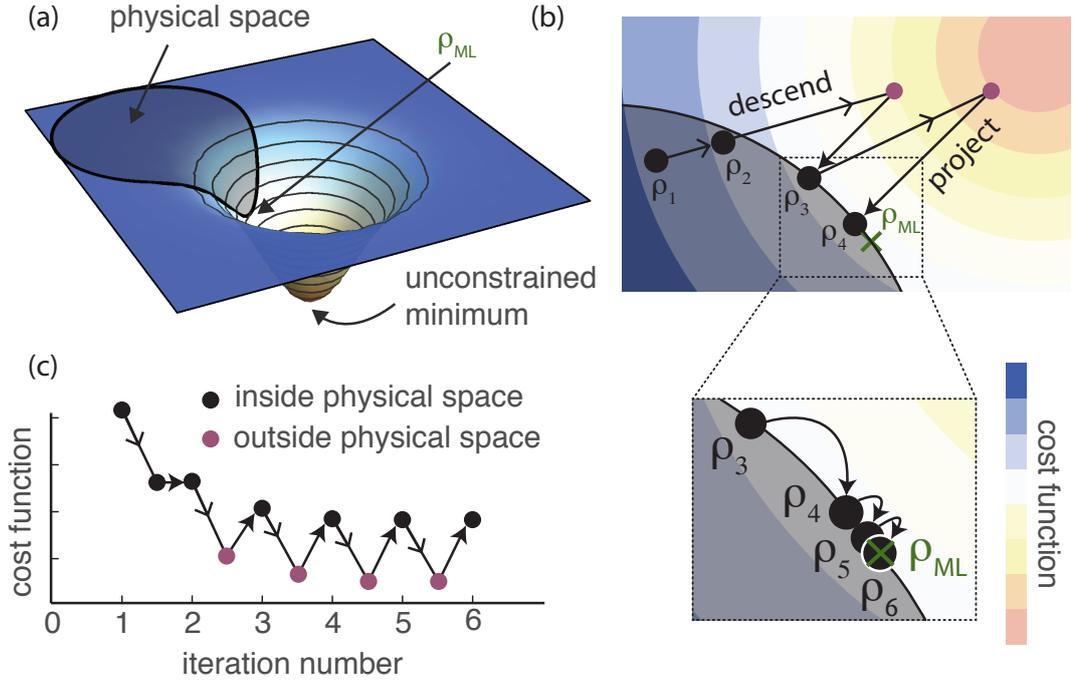


FIGURE 4.1: **Projected gradient descent.** a) The physical space is a convex subset of the larger space of unconstrained matrices. The minimum of the cost function often lies outside of the physical space. b) The PGD process applied on a qubit shows a first step in the gradient direction, that can yield a non physically-allowed density matrix. The projection then brings the estimate back into the constrained search space, (i.e. the Bloch sphere in the case of a qubit). c) This iterative process lowers the cost function until it reaches a sufficiently small value, point at which the final density matrix estimate is as close as desired to the maximum likelihood state ρ_{ML} . Thanks to George Knee, who co-created this figure with the author.

of projection fails to work in practice in a PGD algorithm.

We have therefore established that the projection that yields the best performance in a PGD algorithm is the *simplex projection* $\mathcal{P}[\cdot] \rightarrow \mathcal{S}[\cdot]$. This projector transforms the eigenvalues of the density matrix in such a way that they are all nonnegative and that they sum to unity [90]. The simplex projection is the most sophisticated one, which could explain its supremacy over $\mathcal{F}[\cdot]$ and $\mathcal{M}(\cdot)$. It is based on the knowledge that the diagonal elements of the density matrix represent probabilities, so that they should be nonnegative and sum to one. By contrast, the only characteristic of the density matrix used in $\mathcal{F}[\cdot]$ and $\mathcal{M}[\cdot]$ is its positivity ($\mathcal{M}[\cdot]$ also requires explicit normalisation). Each projection algorithm is laid out in Algorithms 2 to 4.

In the case where the multiplicative factor r is known or computed using Eq. 4.7, the version of the simplex projector $\mathcal{S}[\cdot]$ described in detail in Reference [73] applies. Otherwise, the projection must preserve the trace of the argument, so it is performed over the space of positive matrices [80]. In particular, if the unit-trace property is not enforced at every iteration of the algorithm, the multiplication factor r can be set as a free parameter. During the reconstruction process, r would then be incorporated in the density matrix trace ($\text{Tr}[\rho_{\text{free-trace}}] = r$), and the free-trace density matrix would be normalised

only after its recovery. The latter method reveals itself particularly useful in cases where r cannot be easily estimated using Eq. 4.7 because no subset of the measurement matrix forms a POVM, which is often the case for an ill-conditioned measurement matrix.

Algorithm 2 Projection onto the simplex in MATLAB notation: $\mathcal{S}[\cdot]$

```

1: Compute eigenvalue decomposition of non-physical density matrix  $\rho_{np} = U\Lambda U^\dagger$ .
2: Vectorise and sort eigenvalues in descending order:  $\lambda = \text{sort}(\text{diag}(\Lambda), 'descend')$ 
3: for  $j = 1 : d$  do
4:    $\lambda(1:(d+1-j)) = \lambda(1:(d+1-j)) + (1 - \text{sum}(\lambda(1:(d+1-j)))) / (d+1-j)$ 
5:   if  $\text{abs}(\text{sum}(\lambda(\lambda < 0))) == 0$  then
6:     break
7:   else  $\lambda(d+1-j) = 0$ 
8:  $\Lambda^+ = \text{diag}(\lambda)$ 
9: return  $\rho_p = U\Lambda^+U^\dagger$ 

```

Algorithm 3 Projection onto the closest positive semi-definite matrix, according to the 2-norm: $\mathcal{M}[\cdot]$

```

1: Compute eigenvalue decomposition of non-physical density matrix  $\rho_{np} = U\Lambda U^\dagger$ .
2: Vectorise and sort eigenvalues in descending order:  $\lambda = \text{sort}(\text{diag}(\Lambda), 'descend')$ 
3:  $a = 0$ 
4: for  $i = d : -1 : 1$  do
5:   if  $\lambda(i) + a/i \geq 0$  then
6:     for  $j = 1 : i$  do
7:        $\lambda_{\text{new}}(j) = \lambda(j) + a/i$ 
8:     break
9:   else
10:     $\lambda_{\text{new}}(i) = 0$ 
11:     $a = a + \lambda(i)$ 
12:  $\Lambda^+ = \text{diag}(\lambda_{\text{new}})$ 
13: return  $\rho_p = U\Lambda^+U^\dagger$ 

```

Algorithm 4 Projection onto the closest positive semi-definite matrix, according to the nuclear norm: $\mathcal{F}[\cdot]$

```

1: Compute eigenvalue decomposition of non-physical density matrix  $\rho_{np} = U\Lambda U^\dagger$ .
2: Replace negative eigenvalues with zeros to construct the positive semi-definite matrix  $\Lambda^+$ .
3: return  $\rho_p = U\Lambda^+U^\dagger$ 

```

The following subsections present in details the three PGD algorithms used for quantum tomography, which are all extensions of Eq. 4.26.

Projected Gradient Descent with Momentum

Here, we borrow a technique from the field of machine learning, the momentum-aided gradient descent, to enhance the basic PGD algorithm presented above [91]. In this tech-

nique, we store a running weighted-average of the log-likelihood gradient that we denote M_k . This running average M_k keeps in memory the previous directions of descent of the algorithm, which we use to compute a better estimate of the next descent direction. The algorithm's recursion relation is

$$\begin{aligned} M_{k+1} &= \zeta_k M_k - \gamma_k \nabla \mathcal{C}(\rho_k), \\ \rho_{k+1} &= \mathcal{S}(\rho_k + M_k), \end{aligned} \tag{4.27}$$

where γ_k is the step size, ζ_k is a term that can be interpreted as a level of ‘‘inertia’’ for the descent direction. These metaparameters can either be dependant on the iteration number k or be set as fixed parameters, constant through all iterations of the algorithm.

Momentum in a gradient descent algorithm allows one to take a step in a direction without undoing the few previous steps. The conjugate gradient algorithm is similar in that one step is completely orthogonal to the previous one. The conjugate gradient algorithm has been applied to quantum tomography in [31].

Algorithm 5 gives the full pseudo-code for the projected gradient descent with momentum (PGDM) algorithm. The symbol \circ used in this algorithm and the following ones represent the Hadamard product, which is an element-wise multiplication.

Algorithm 5 PGDM

- 1: $k = 0$
 - 2: **Initial estimate and momentum matrix:** $\rho_0 = I, M_0 = 0$
 - 3: *currentMagnitude* = $\lceil \log_{10} \mathcal{C}_P(\rho_0) \rceil$
 - 4: **Set step size and inertia:** $\gamma = (2rd)^{-1}, \zeta = 0.95$
 - 5: **while** $\sum_{j=1}^{20} |\mathcal{C}_P(\rho_j) - \mathcal{C}_P(\rho_{j-1})| > 10^{-5}$ **do**
 - 6: **Projection:** $\rho_k = \mathcal{S}(\rho_k)$
 - 7: **Estimate probabilities:** $\mathbf{p}_k = \sum_j [\mathcal{A} \circ (\mathcal{A}^* \rho_k)]_{i,j}$
 - 8: **Calculate log-likelihood:** $\mathcal{C}_P(\rho_k) = \boldsymbol{\nu}^T \boldsymbol{\nu} / N$
 - 9: **Compute gradient:** $\nabla \mathcal{C}_G(\rho_k) = 2G^\dagger \mathcal{A}^*$
 - 10: *currentMagnitude* = $\lceil \log_{10} \mathcal{C}_P(\rho_k) \rceil$
 - 11: **if** *currentMagnitude* < *previousMagnitude* **then**
 - 12: **Update inertia:** $\zeta_k = (1 - (1 - \zeta_k) * 0.95)$
 - 13: *previousMagnitude* = *currentMagnitude*
 - 14: **Update momentum:** $M_{k+1} = \zeta_k M_k - \gamma \nabla \mathcal{C}(\rho_k)$
 - 15: **Update density matrix:** $\rho_{k+1} = \rho_k + M_{k+1}$
 - 16: $k = k + 1$
 - 17: **Final projection:** $\rho_{\text{final}} = \mathcal{S}(\rho_{k+1})$
 - 18: **Return** ρ_{final}
-

Fast Iterative Shrinkage-Thresholding Algorithm

PFISTA is an algorithm that was initially developed for classical image denoising [82]. Here, we introduce and refine this method for application to quantum state tomography.

In contrast with other PGD implementation, the gradient step in PFISTA does not always follow the descent direction: the estimate ρ_k is allowed to change in a way that increases the log-likelihood function. As we will demonstrate, this algorithm can nevertheless descend much faster than the algorithm from [73], presented in the next subsection.

The recursion relation used in PFISTA is as follow:

$$\rho_{k+1} = \mathcal{S} \left[\rho_k + \frac{k-2}{k+1}(\rho_k - \rho_{k-1}) - \delta \nabla \mathcal{C}(\rho_k) \right], \quad (4.28)$$

where δ is the step size. This algorithm uses a linear combination of two previous estimates to cheaply compute the next one. It has been shown in the context of image processing that this method allows for an improvement in time complexity [82].

Algorithm 6 gives the full pseudo-code for the fast iterative shrinkage-thresholding algorithm (PFISTA).

Algorithm 6 PFISTA

- 1: $k = 0$
 - 2: Initial estimate and momentum matrix: $\rho_0 = I, M_0 = 0$
 - 3: Set step size: $\delta = (10d)^{-1}$
 - 4: **while** $\sum_{j=1}^{20} |\mathcal{C}_P(\rho_j) - \mathcal{C}_P(\rho_{j-1})| > 10^{-6}$ **do**
 - 5: Projection: $\rho_k = \mathcal{S}(\rho_k)$
 - 6: Estimate probabilities: $\mathbf{p}_k = \sum_j [\mathcal{A} \circ (\mathcal{A}^* \rho_k)]_{i,j}$
 - 7: Calculate log-likelihood: $\mathcal{C}_P(\rho_k) = \boldsymbol{\nu}^T \boldsymbol{\nu} / N$
 - 8: Compute gradient: $\nabla \mathcal{C}_G(\rho_k) = 2G^\dagger \mathcal{A}^*$
 - 9: $\rho_{k+1} = \rho_k + (k-2)(\rho_k - \rho_{k-1})(k+1)^{-1} - \delta \nabla \mathcal{C}_G(\rho_k)$
 - 10: $k = k + 1$
 - 11: Final projection: $\rho_{\text{final}} = \mathcal{S}(\rho_{k+1})$
 - 12: **Return** ρ_{final}
-

Projected gradient descent with backtracking

The PGDB algorithm is described in detail in [73]. This algorithm's main characteristic is that it focuses on finding the maximum step size to reduce the negative log-likelihood. It has the following recursion relation:

$$\rho_{k+1} = (1 - \alpha)\rho_k + \alpha \mathcal{S}[\rho_k - \nabla \mathcal{C}(\rho_k)], \quad (4.29)$$

where α is a parameter that is optimised through backtracking at each step. Unless a stationary point is reached, thus satisfying the stopping criterion, a decrease of the negative log-likelihood is guaranteed at each iteration of the algorithm.

Algorithm 7 gives the full pseudo-code for the projected gradient descent with backtracking (PGDB) algorithm.

Algorithm 7 PGDB

```
1:  $k = 0$ 
2: Initial estimate and momentum matrix:  $\rho_0 = I, M_0 = 0$ 
3: Set metaparameters:  $\mu = 1, \ell = 10^{-4}$ 
4: while  $\sum_{j=1}^{20} |\mathcal{C}(\rho_j) - \mathcal{C}(\rho_{j-1})| > 10^{-5}$  do
5:   Projection:  $\rho_k = \mathcal{S}(\rho_k)$ 
6:   Estimate probabilities:  $\mathbf{p}_{\mathbf{k}_i} = \sum_j [\mathcal{A} \circ (\mathcal{A}^* \rho_k)]_{i,j}$ 
7:   Calculate log-likelihood:  $\mathcal{C}(\rho_k) = \boldsymbol{\nu}^T \boldsymbol{\nu} / N$ 
8:   Compute gradient:  $\nabla \mathcal{C}(\rho_k) = 2P^\dagger \mathcal{A}^*$ 
9:    $\rho'_k = \mathcal{S}(\rho_k - \mu^{-1} \nabla \mathcal{C}_G(\rho_k))$ 
10:   $D = \rho'_k - \rho_k$ 
11:  Line search initialisation:  $\alpha_k = 1$ 
12:   $\mathcal{C}'_G(\rho_k) = \mathcal{C}_G(\rho_k) + \ell \alpha_k \text{Tr}[D \nabla \mathcal{C}_G(\rho_k)]$ 
13:  while  $\mathcal{C}_G(\rho_k + \alpha_k D) > \mathcal{C}'_G(\rho_k)$  do
14:    Line search:  $\alpha_k = \alpha_k / 2$ 
15:     $\mathcal{C}'_G(\rho_k) = \mathcal{C}_G(\rho_k) + \ell \alpha_k \text{Tr}[D \nabla \mathcal{C}_G(\rho_k)]$ 
16:  Update density matrix:  $\rho_{k+1} = \rho_k + \alpha_k D_k$ 
17:   $k = k + 1$ 
18: Final projection:  $\rho_{\text{final}} = \mathcal{S}(\rho_{k+1})$ 
19: Return  $\rho_{\text{final}}$ 
```

Code for PGD algorithms

In addition to the pseudo-codes provided in Algorithms 5, 6 and 7, the complete MATLAB scripts for our PGD algorithms are available at <https://github.com/eliotbo/PGDfullPackage>. They can be used to reproduce the results presented in section 4.3.

4.2.4 Benchmark algorithms

Diluted Iterative Algorithm

The diluted iterative algorithm (DIA) is a method that also exploits the gradient of the log-likelihood function. The algorithm's iterations are computed as follow [32, 33]:

$$\begin{aligned} R_k &= -H^{-1/2} [\nabla \mathcal{C}] H^{-1/2}, \\ \rho_k &= \frac{(I + \epsilon R_k) \rho_{k-1} (I + \epsilon R_k)}{\text{Tr}[(I + \epsilon R_k) \rho_{k-1} (I + \epsilon R_k)]}. \end{aligned} \quad (4.30)$$

See Eq. 4.16 for the form of $\nabla \mathcal{C}$.

The variable ϵ is used to minimise the log-likelihood function; it is optimised at every iteration and can be implemented in various ways [33, 92]. The matrix H is defined as

$$H = \sum_i \Pi_i / \sum_i p_i, \quad (4.31)$$

where the Π_i are the measurement operators. When all operators Π_i form a POVM, the matrix reduces to identity, up to a constant. The DIA leaves the density matrix estimate ρ_k positive at every iteration.

Studies on the DIA have shown that it converges quickly in the first few iterations, and very slowly later on [31, 73, 77, 85]. In section 4.3, we show that our results corroborate these observations.

Semidefinite programming

As detailed above, quantum tomography is a problem that can be expressed as a convex function that we seek to minimise over a convex set. In the field of numerical optimisation, once a problem is expressed in this form, it is considered as good as solved. Indeed, powerful and efficient algorithms and software packages to solve this type of problem are readily available; these are also guaranteed to converge at the global optimal solution.

In this work, we use such a software package as another benchmark for our quantum tomography algorithms. In particular, we use the CVX software environment and the SDPT3 solver which is an implementation of the infeasible path-following algorithm. It is important to note that quantum tomography, although often a convex problem, sometimes is not. The Gaussian approximation of a Poisson log-likelihood is an example of a non-convex problem [30], and the cost function in this case would be rejected by CVX and other similar software packages [35, 93]. In such a case, the programs, designed to operate in a “disciplined convex programming” paradigm, do not offer guarantee of global optimality. However, PGD techniques will still work in this case to find a local optimum.

In the most common case where the quantum tomography problem at hand is convex, it is also important to understand that these softwares are designed for a general purpose and are not optimised specifically for tomography problems: it is therefore not surprising that they do not perform as quickly as the other methods.

4.3 Numerical results

In this section, we perform quantum tomography simulations on multi-qubit systems and use all of the techniques detailed in section 4.2.3 to retrieve the best estimate of the quantum state. We find that using measurement operators that are canonical Pauli measurements lead to a good recovery of ρ_{ML} for all the techniques. We therefore concentrate on the total computation time as a figure of merit, as in practice we consistently reach high likelihoods.

For all simulations, we use an average number of events per outcome r/d of 10^4 . The true density matrices are randomly generated in the Haar sense (see Section 3.3.1, Subsubsection called Generating random state vectors for an explanation of the Haar measure), but we fix their purity to 0.5. For all algorithms except SDPT3, for which the code was not modified, we define an exit criterion based on an average gradient: if the average gradient of the last 20 iterations is low enough, the algorithm is terminated. All simulations are performed on a single thread of an Intel Xeon Haswell processor.

Examples of convergence curves obtained using PGD algorithms and the benchmark algorithms on a 6-qubit system characterised using Pauli measurements are shown in Fig. 4.2a. We observe the typical behavior of DIA, described in section 4.2.4, by which the algorithm displays a fast convergence in the early iterations, but quickly decelerates and necessitates a high number of iterations to finally meet the exit criterion. This results corroborate other recent studies on DIA [31,73,77,85]. We also find that SDPT3 requires only 10 and 15 iterations to converge, but the convergence is slow in high dimensions as each iteration has a computational complexity of $O(N^2 d^2)$.

It is also of vital importance to use PGD in a context of more practical scenarios, and not only using ideal assumptions of the measurement matrix A . With this goal in mind, we also apply the reconstruction algorithms in simulations using ill-conditioned measurement matrices [84, 85, 87], in which case the error of the final density matrix estimates necessarily increase compared to a well-conditioned problem [84]. We can construct matrices with condition number higher than unity by limiting the measurement operators to a restricted region of the Hilbert space. Here, we specifically build the measurement matrix using tensor products of the following qubit bases:

$$\begin{aligned} & \{[0, 1], [1, 0]\}, \\ & \{[\cos(\beta/2), \sin(\beta/2)], [\sin(\beta/2), \cos(\beta/2)]\}, \\ & \{[\cos(\beta/2), \mathbf{i} \sin(\beta/2)], [\sin(\beta/2), \mathbf{i} \cos(\beta/2)]\}. \end{aligned} \quad (4.32)$$

To construct the measurement matrix A , we unwrap in the usual way the projector operator associated with each of these vectors, which are illustrated in the Bloch spheres on Fig. 4.2b. Using a value of $\beta = \pi/2$ leads back to the canonical Pauli operators. For our simulations, we choose $\beta = \pi/3$ to model ill-conditioned measurements. The results of the PGD and benchmark algorithms on this ill-conditioned problem are shown in Fig. 4.2 b.

As mentioned above, both PGDM and PFISTA are algorithms in which the cost function can increase before reaching the maximum likelihood value. In contrast, Gonçalves *et al.* provide a proof that the cost function \mathcal{C} for PGDB is monotone: it never increases from one step to the next. that is to say that the cost function never increases in this

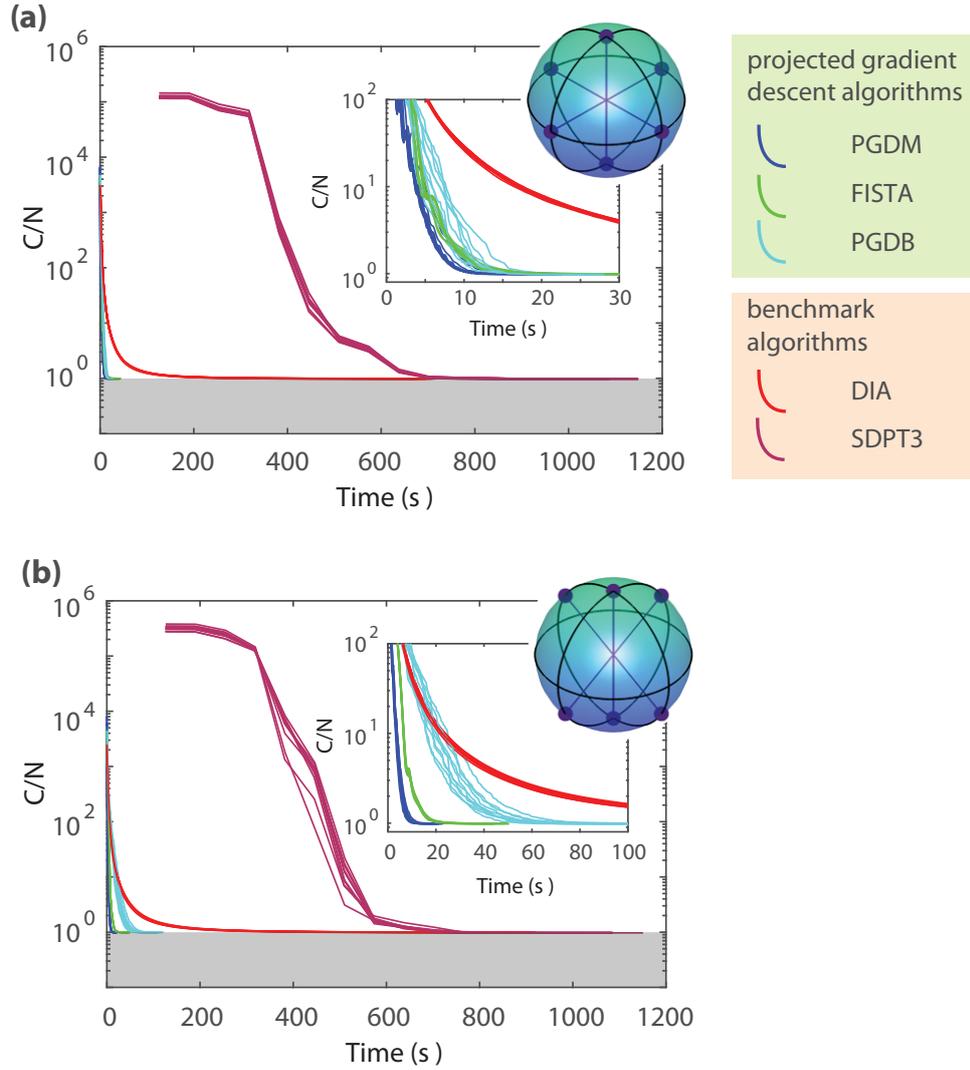


FIGURE 4.2: **Convergence curves of the cost function as a function of running time.** Simulations performed on six-qubit systems using (a) Pauli measurements, shown as vectors on the Bloch sphere, and (b) an ill-conditioned measurement matrix. The global minimum of the negative log-likelihood is expected to be found near the top of the grey regions, around $C/N \approx 1$. The algorithms converge steadily, apart from PGDM and PFISTA where the cost function goes up as it advances in iteration number, before reaching the ML state. As we can observe from the inserts, the total running time for PGDB is highly correlated with the measurement matrix condition number.

algorithm. Interestingly, we observe in Fig. 4.2b that, in the context of ill-conditioned measurement, this tends to speed up the performance of PGDM and PFISTA with respect to PGDB performance.

We also run the PGD and benchmark algorithms on a measurement and system of varying Hilbert space dimension. Fig. 4.3 shows the running time of each algorithm as a function of number of qubits in the system. We observe a speed-up of PGDM with respect to PGDB at high dimensions for both measurements with Pauli matrices and with an ill-conditioned matrix, although it is more pronounced in the case of the ill-conditioned measurement. We find that in this case, PGDM is on average about ten times faster than

PGDB for the seven and eight qubit cases.

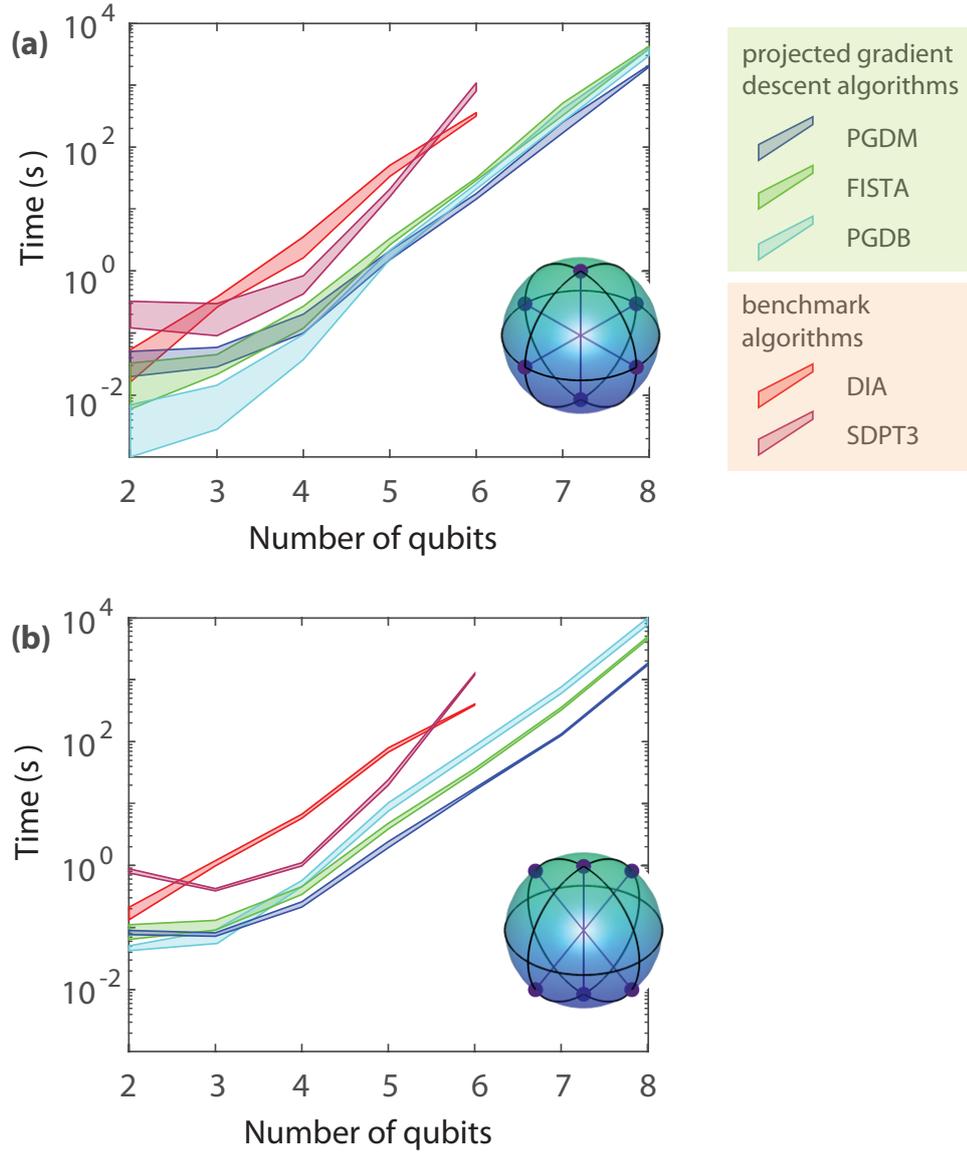


FIGURE 4.3: **Running time as a function of dimensionality.** The measurement matrix is composed of (a) Pauli measurements and (b) an ill-conditioned matrix constructed from bases relatively close to each other, as indicated on the inset. The coloured areas show the standard deviation on the running time ($\pm 1\sigma$). The gradient-based techniques have a computational complexity of $O(Nd^2)$, SDPT3 converges in time $O(N^2d^2)$.

We also observe that the number of iterations required to reach ρ_{ML} in both PGDM and PFISTA change very little with the measurement matrix condition number. As a result, these algorithms are more resilient to a varying condition number than PGDB. We illustrate this behaviour in Fig. 4.4b. On the other hand, we find that the semidefinite programming software SDPT3 always takes about 15 iterations to reach ρ_{ML} , independently of the condition number. This points to the fact that, although the preexisting PGDB method does not always outperform SPDT3, our novel variants of PGD instead provide a speed-up of up to 5x over SDPT3 and up to 50x over PGDB, as shown in Fig. 4.4c and

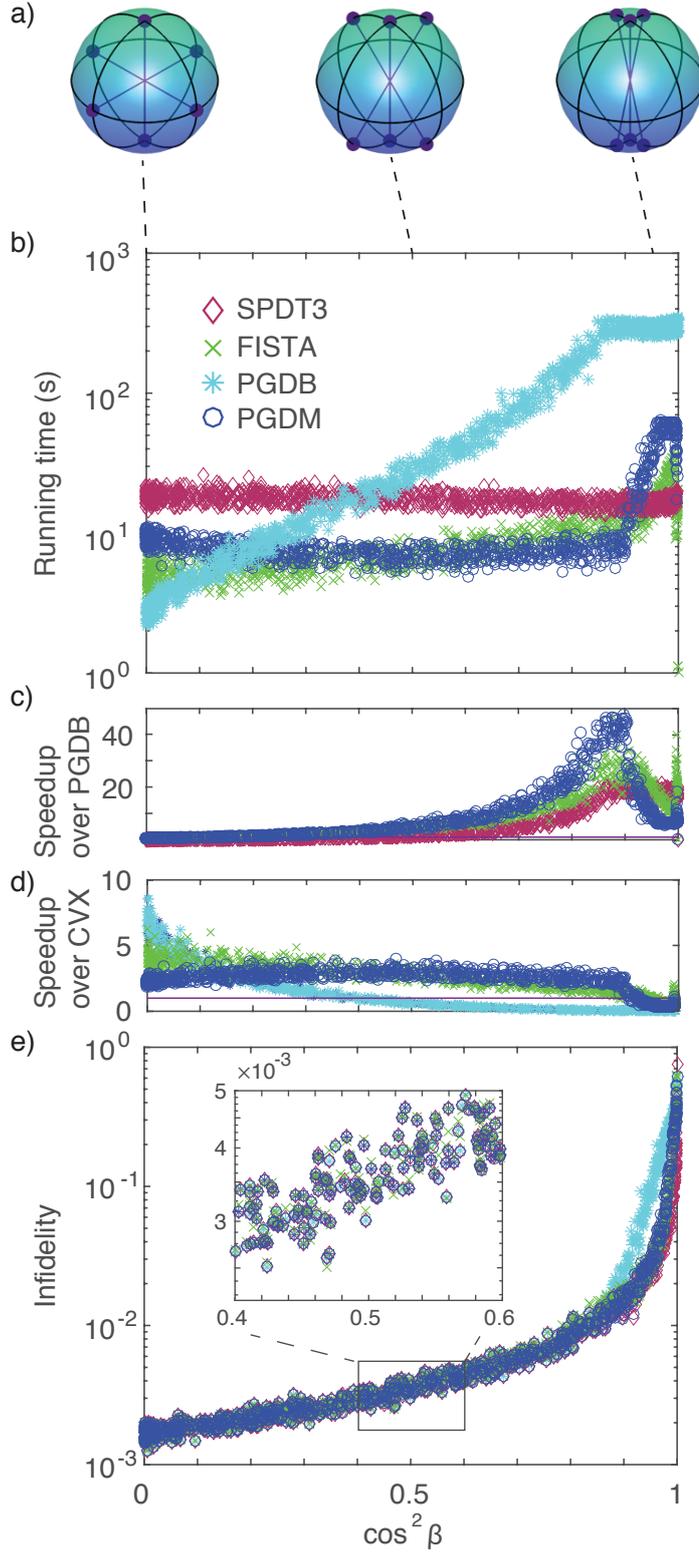


FIGURE 4.4: **Performance as a function of ill-posedness** The PGD algorithms are applied to the reconstruction of five-qubit states with a number of events per outcome of 10^4 . a) The measurement matrices, with examples illustrated on the Bloch sphere, are of varying ill-posedness. b) Running time of the different approaches. The PGDB algorithm saturates when it reaches the maximum number of iterations set in the script. c) Resulting speedup of the different approaches with respect to PGDB. d) Resulting speedup of the different approaches with respect to SPDT3. e) Infidelity between the true and recovered states. The inset shows the methods converge towards the same fidelity as they reach ρ_{ML} . The PFISTA has a slightly different behaviour as it oscillates around ρ_{ML} for a large number of iteration.

4.4d.

We also find that the ill-posedness of the measurement and the accuracy of the recovered state are monotonically related: for a fixed number of events per measurement, the more ill-posed the measurement matrix, the lower the fidelity between the true state and the recovered one. We illustrate this behaviour in Fig. 4.4e, by displaying the infidelity between the true and recovered state matrix, defined as one minus the fidelity $f(\rho_1, \rho_2) = \text{Tr}(\sqrt{\sqrt{\rho_1}\rho_2\sqrt{\rho_1}})$, as a function of the spread in Hilbert space of the basis used for constructing the ill-conditioned measurement matrix. The extreme case where $\cos^2 \beta = 0$ corresponds to mutually unbiased bases, occupying the complete Hilbert space. On the other hand, the case where $\cos^2 \beta = 1$ corresponds to a single basis measurement. Cases near both extremes are illustrated on the Bloch spheres in Fig. 4.4a. In Fig. 4.4e, we zoom on a subset of the infidelity curve as a function of ill-posedness around the intermediate angle, where $\cos^2 \pi/4 = 0.5$, where the low values of infidelity give statistical evidence that the PDG algorithms consistently converge to ρ_{ML} . The fact that the benchmark algorithms and the PGD algorithms all converge towards more or less the same infidelity is also a sign that the PGD algorithm converge towards ρ_{ML} .

4.3.1 Pure-state tomography via gradient-descent

In the case where the state being measured is known to be quasi-pure, the number of necessary measurements for retrieving the state is significantly reduced. Only a linear number of measurements are needed to uniquely identify a pure quantum state [34,38,94]. Gradient descent can also be applied in this situation. The algorithm that we describe below is similar to one developed in the context of the phase retrieval problem [36].

In this section, we assume that the state that we seek is pure, like in Chapter 3 but the following algorithm can be applied to mixed states that are known to have a high spectral gap, i.e. a large gap between the principle eigenvalue and the second highest eigenvalue of the density matrix [38]. With a large spectral gap, it is likely that an algorithm that assumes purity retrieves the principle eigenvector of the density matrix, which often corresponds to the signal that is needed.

In the context of pure states, one needs to reformulate the tomography problem with vectors instead of operators. We will now adopt the notation used for pure states: the state vector that we seek to recover is $|\psi\rangle$, the measurement vectors are $|\phi_i\rangle$ and the outcomes are still written \mathbf{n} . We define the pure-state measurement matrix \mathcal{A}_p by stacking N many measurement vectors $\langle\phi_i|$ on top of each other.

Conveniently, since the noise on measurement outcomes is still of Poissonian type, we can use the maximum likelihood function from Eq. 4.15 and replace the form of the

measurement matrix and the quantum state. The Poisson likelihood is given by

$$\mathcal{L}_{\text{Poisson}} = \prod_i \frac{p_i^{n_i} \exp(-p_i)}{n_i!}, \quad (4.33)$$

with $p_i = |\langle \phi_i | \psi \rangle|^2$. The maximum likelihood state minimizes the cost function $f_{\text{Pure}} = -\log \mathcal{L}_{\text{Poisson}}$:

$$f_{\text{Pure}} \propto \sum_i p_i - n_i \log p_i, \quad (4.34)$$

where we omit the factorial term because it does not affect the gradient. This log likelihood resembles that of a multinomial likelihood, the only difference being the single p_i term coming from the exponential in Eq. 4.33. The only terms that matter here are those with p_i in them:

$$\frac{\delta f_{\text{pure}}}{\delta |\psi\rangle_j} = \sum_i \frac{\delta p_i}{\delta |\psi\rangle_j} \left(1 - \frac{n_i}{p_i}\right) \quad (4.35)$$

$$\text{with } \frac{\delta p_i}{\delta |\psi\rangle_j} = 2|\phi_i\rangle_j \langle \phi_i | \psi \rangle \quad (4.36)$$

One difference with the density matrix cost function is that $f_{\text{Pure}}(|\psi\rangle)$ contains quadratic terms. In general, this cost function is non-convex, but we show that minimizing it still gives good approximations to state $|\psi\rangle$ in a wide range of cases. While a density matrix ρ lives in the restricted space of positive semi-definite matrices, a state vector $|\psi\rangle$ is a very general vector, i.e. with no restrictions. Hence, there is no need for the projection step in an iterative algorithm to minimize Eq. 4.34. It suffices to descend along the cost function gradient. The matrix-vector version of Eqs. 4.35 and 4.36 is given by

$$\nabla_{\psi} f_{\text{Pure}} = 2\mathcal{A}_p^\dagger [(\mathcal{A}_p |\psi\rangle) \circ (1 - \mathbf{n} \circ |\mathcal{A}_p |\psi\rangle|^{-2})], \quad (4.37)$$

where $|\mathcal{A}_p |\psi\rangle|^{-2}$ is the vector of inverse expected outcomes. The advantage of expressing the gradient without a sum is that computer architectures are optimized to perform matrix-vector multiplications.

At the beginning of our algorithm, we use a non-iterative method that gives a crude approximation to the actual state $|\psi\rangle$. The procedure consists of building the matrix $M = \sum_i |\phi_i\rangle \langle \phi_i| n_i$ and taking the principal eigenvector $|M_p\rangle$, i.e. the eigenvector corresponding to the highest eigenvalue, of M as the initial guess state vector. the primary eigenvector of M typically has a fidelity $|\langle M_p | \psi^* \rangle|$ of about 0.9 in the absence of noise and with quasi-symmetric measurements (see next subsection about Quasi-Symmetric POVMs). For the initial guess to produce the correctly scaled outcomes \mathbf{n} , it has to be scaled to the square root of the number of repeated experiments in a given experimental configurations, which can either be measured – the sum m of all events in a POVM measurement – or estimated with all the gathered measurements $m \approx dN^{-1} \sum n_i$, where d

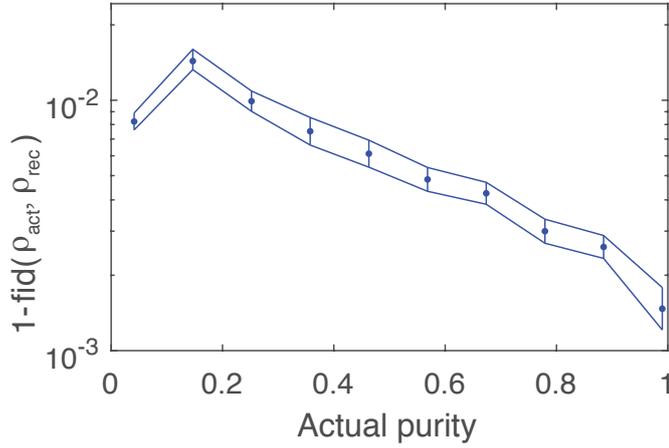


FIGURE 4.5: **Accuracy of pure state tomography via gradient descent versus the purity of the state at hand as measured through the fidelity.** This simulation is performed on a five-dimensional system with Gaussian additive noise on the generated outcomes. The density matrices ρ_{act} and ρ_{rec} are the actual and recovered density matrices.

is the system dimensionality and N is the number of measurement projectors. The latter equation is based on the fact that the average outcome of a measurement is N/d if the measurements are equally separated in the Hilbert space. The initial guess is $\sqrt{m}|M_p\rangle$.

The computational complexity of the pure-state gradient descent algorithm is $O(d^3 \log(d))$. This is to be compared with a complexity of $O(d)$ for the direct approach of Chapter 3. The advantage of the gradient descent algorithm is that it is compatible with any complete set of measurement projectors, including the very efficient and informative quasi-symmetric POVMs (next subsection).

The pure-state gradient descent algorithm is non-convex, but with a sufficient number of measurements the probability of hitting a local minimum can be made extremely small. It has been shown that $4d - 4$ linearly independent measurements are enough to uniquely identify any quantum pure state [95]. In practice however, we find that $5d$ measurement projectors is a better number for reliably finding the minimum of Eq. 4.34.

Quasi-Symmetric POVM

Symmetrical informationally-complete POVMs (SIC-POVM) are sets of d^2 measurement operators that maximise the expected variance of the data. As such, one SIC-POVM element outcome yields the maximum amount of information about the state at hand in dimension d , provided that there is no *a priori* information about said state.

If the state at hand is known to be pure or quasi-pure, as mentioned above, the tomographer does not need a set of d^2 outcomes since $4d - 4$ outcomes are sufficient to uniquely identify the unknown pure state. Therefore, it would be useful to generalise the concept of SIC-POVMs for pure states: a symmetrical or quasi-symmetrical set of projectors of

arbitrary size would be useful. This is what we do in the following few paragraphs, except that we find a close approximation to perfectly symmetrical projectors since a perfect set of arbitrary size does not exist in all dimensions [96].

Let us define $\mu = |\langle \phi_j | \phi_i \rangle|^2$. We seek to find a value for μ as a function of the dimensionality d and the number of POVM elements N . If a symmetric POVM of N elements were to exist, then we would have

$$\sum_i^N |\phi_i\rangle\langle\phi_i| = \frac{N}{d}I. \quad (4.38)$$

It follows that

$$\frac{N}{d} = \langle \phi_j | \left(\sum_i^N |\phi_i\rangle\langle\phi_i| \right) | \phi_j \rangle \quad (4.39)$$

$$= [1 + (N - 1)\mu]. \quad (4.40)$$

Isolating μ yields

$$\mu = \frac{N - d}{d(N - 1)}. \quad (4.41)$$

The Gram matrix is defined as $G = \mathcal{A}_p \mathcal{A}_p^\dagger$. The Gram matrix has unit diagonal elements. For perfectly symmetric states, the offdiagonal elements should satisfy the equation $|G_{v,w}|^2 = \mu$. Hence, we define the reference matrix R whose elements are all equal to μ except for the diagonal elements, which are equal to one. We want to find the measurement matrix \mathcal{A}_p that yields a Gram matrix G closest to the reference matrix R . To solve this problem, we use the following cost function

$$f = \sum_{v,w} (|G_{v,w}|^2 - R_{v,w})^2, \quad (4.42)$$

the gradient of which is given by

$$\frac{df}{d\mathcal{A}_p} = 4\mathcal{A}_p^\dagger [G \circ (G \circ G^\dagger - R)], \quad (4.43)$$

where the operation \circ corresponds to element-wise multiplication (or the Hadamard product).

Minimizing Eq. 4.42 can be done through gradient descent, but this function is not convex. Numerical experiments show, however, that it is possible to find POVMs whose elements have very little redundancy. In other words, there is not much informational difference between the solution found through a local minimum of Eq. 4.42 and a perfectly symmetric POVM. To precisely quantify the informational difference, one would have to

calculate the Fisher information associated with the measurement matrices. Note that the rows of the solution S_p to minimizing Eq. 4.42 do not exactly form a POVM, although they might be close to forming a POVM. To remedy this predicament, we find the closest unitary matrix Q to S_p in the Frobenius norm [Keller:1975]

$$Q = S_p(S_p^\dagger S_p)^{-1/2}. \quad (4.44)$$

The rows of Q now form a POVM; this can be proved simply using the rows \mathbf{q}^v of Q :

$$\sum_{v=1}^N \mathbf{q}^v \mathbf{q}^{v\dagger} = Q^\dagger Q = (S_p^\dagger S_p)^{-1/2} S_p^\dagger S_p (S_p^\dagger S_p)^{-1/2} = I_d \quad (4.45)$$

Finally, we normalize the rows of Q to find a close set of measurement vectors to a symmetric system

$$\mathcal{A}_p = [I_N \circ (QQ^\dagger)]^{-1/2} Q. \quad (4.46)$$

Using the above procedure, a typical maximum value for the Gram matrix offdiagonal elements is 2μ . Fig. 4.6 shows the Gram matrix for a) an initial random guess to the set of projectors and b) the set of quasi-symmetrical projectors found through gradient descent. A Gram matrix is the same as an overlap matrix, for which an element (i, j) corresponds to $|\langle \phi_i | \phi_j \rangle|^2$. A perfectly symmetrical set of projectors would yield the R matrix mentioned above: the diagonal element are all equal to each other and the off-diagonal are all equal to each other.

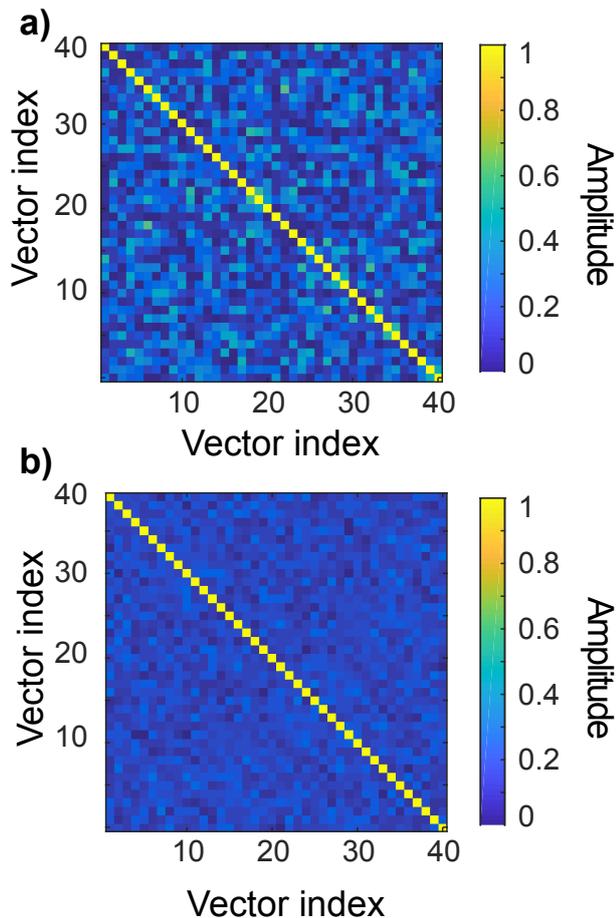


FIGURE 4.6: **Absolute square of the Gram matrices $|\mathcal{A}_p \mathcal{A}_p^\dagger|^2$, with the square operation applied element-wise** a) the initial vector set b) the final vector set, i.e. after running an optimization algorithm for minimizing Eq. 4.42. The dimensionality is $d = 8$ and the number of vectors in a measurement matrix is $N = 40$, such that $\mu = 0.10$. The maximum fidelity between any pair of different vectors is 0.56 in the initial set and 0.18 in the final set (closer to the predicted pairwise overlap μ).

4.4 Discussion

In our study of both existing and novel PGD algorithms and benchmark algorithms, we find that the the new algorithms converge faster than PGDB and the benchmark algorithms, and that they scale more favourably than SDPT3. Our conclusions corroborate and further extend that of a recent parallel work [31], in which a number of routines, including momentum and backtracking, are combined in a hybrid PGD algorithm. We also find in our work that PGD techniques present a versatility that is one of their key advantage. The PGD techniques prove able to quickly and successfully converge to the maximum likelihood state for a wide range of cases, in a variety of noise types, whether the measurement matrix is well-conditioned (ideal) or ill-conditioned (realistic conditions), and regardless of the desired accuracy. Our two new PGD algorithms, PGDM and PFISTA, are found to be especially well suited for ill-conditioned problems.

We have addressed in our work the problem of quantum state tomography, which is one of the three main subfields of quantum tomography, the other two being detector and process tomography. The algorithms developed here for state tomography are straightforwardly transferable to process tomography. However, transferring to the algorithms to detector tomography with success is not a trivial endeavour. In detector tomography, one strives to characterise an unknown detector POVM using an informationally complete set of known states. When using coherent states for probing the detector, the detector tomography problem is ill-conditioned and, similar to the density matrix, the POVM elements to recover must be positive-semidefinite [86]. The current state-of-the-art for solving this problem is to use semidefinite program solvers such as SeDuMi [85]. As we find our new algorithms, PGDM and PFISTA, to perform very well in the case of ill-conditioned measurement matrix, this suggests that they hold great promise for application to the field of optical detector tomography. The characterisation of detector POVMs in high dimensions is thus an interesting avenue for future work.

4.5 Conclusion

In summary, the PGD techniques we devised and implemented have proven to perform well and to converge towards the maximum likelihood density matrix reliably. The different PGD techniques we explore find themselves to be complement of each other in several aspects. We find that the running times of our three PGD variants, when evaluated with measurements of varying condition number, present crossover: for a specific condition number, there is therefore an optimal variant to be used, that is not the same whether the measurement is well or ill-conditioned. A judicious choice of PGD variant can offer speedups up to a 50x factor. We also find interesting behaviours with respect to the scaling with dimensionality: we find that PGDB is fastest in low dimension and PGDM, according to our numerical results, is fastest beyond five-qubit systems. We also find that all PGD techniques converges to ρ_{ML} significantly faster than the benchmark algorithms DIA and SDPT3 in the vast majority of scenarios, thus pushing the limits of state-of-the-art techniques for assumption-free quantum state tomography.

Chapter 5 – Entangled photons coincidence detection with an EMCCD camera

This body of work was initiated by my PhD advisor, Dr Jonathan Leach. It consists of a follow-up to the paper by Matthew *et al.*, “Imaging high-dimensional spatial entanglement with a camera” [97], where correlations between entangled photons are recorded with an EMCCD camera. We extend this work by measuring correlations in a wider range of conditions. Specifically, I formalised the theory for the detected number of photon coincidences and the measurement of sub-shot-noise statistics. I equally participated to the experiments together with Dr Jonathan Leach, who also performed a major part of the data analysis. This work is published in Journal of Optics [98].

This chapter is an outlier in that the content does not describe an inverse problem. Here, we measure photon coincidences generated through spontaneous parametric down-conversion (SPDC) on a camera using spatial information. SPDC photons have opposite momenta and coincidences can be measured from the signals on opposite pixels on a camera, building coincidence statistics over many frames. The formalism that we construct here, including a calibration independent figure of merit, can be used to characterise sub-shot-noise photon sources, such as SPDC sources.

5.1 Introduction

Methods for characterising quantum states of light and their photon number statistics are traditionally based on single-photon detectors. However, to alleviate the constraint of performing only point measurements, cameras and multi-pixel devices with high sensitivity are now being used to measure quantum correlations and characterise spatially entangled photons. In particular, sub-shot-noise imaging can be performed using CCD cameras to improve upon standard imaging methods when imaging in low photon flux conditions [99, 100]. Intensified CCD (ICCD) cameras have proven to be capable of performing real-time imaging of an entangled source [101]. Recent experiments on Einstein-Podolsky-Rosen type correlations have also included an electron multiplying CCD (EMCCD) camera with the aim of showing in a direct fashion the non-locality of quantum physics [3]. The state dimensionality that can be measured and analysed with cameras is potentially very high as the hundreds of thousands of pixels of cameras can be exploited, which makes this type of work also of interest in the field of quantum communications.

Cameras have also been exploited specifically in experiments aiming at measuring spatial correlations of photon pairs generated through spontaneous parametric down conversion (SPDC), a common source of photon with sub-shot-noise behavior [97, 102–110]. The detection of multiple correlated events per frame has been demonstrated in various studies using efficient CCD cameras [102, 104–108]. In this work, we demonstrate the capabilities of an EMCCD camera to perform coincidence measurements on a SPDC source for a mean photon number per pixel much greater than unity. We work in regimes where counts range from 1 to 10'000 pairs per exposure, thus making the bridge between the regime of a few events to a large number of events per exposure.

5.2 Theoretical development

We develop here the formalism for the coincidence detection of photon-pairs generated through SPDC and detected using an EMCCD camera.

5.2.1 Ideal photon number statistics

Traditional coherent monochromatic light sources are characterised by photon statistics in which the number of generated photons in a given time window follows a Poissonian distribution, in which the variance of the photon number in the time window $\sigma_{\text{Poisson}}^2$ is equal to the mean generated photon number in the same time window $\langle N_{\text{Poisson}} \rangle$:

$$\sigma_{\text{Poisson}}^2 = \langle N_{\text{Poisson}} \rangle. \quad (5.1)$$

The variations in photon number is fundamentally due to the uncertainty relation between photon number and phase [111], and is often referred to as “shot noise”. A source for which the photon number distribution has a lower variance than the mean is said to have a sub-shot-noise behavior:

$$\sigma_{\text{sub-shot}}^2 < \langle N_{\text{sub-shot}} \rangle. \quad (5.2)$$

The higher photon number accuracy of these sources reveals advantageous in a variety of situations where the accuracy on the phase is of less importance. In particular, sources with sub-shot-noise behavior can be used to reduce noise in an interferometer [112], and were used to increase the sensitivity levels of gravitational wave detectors [113–117]. They can also be exploited for accurate imaging methods at low light levels (a few thousands photons/camera exposure) [99, 100, 118] and high precision spectroscopy [119].

A common source presenting sub-shot-noise behaviour are photon pairs generated via SPDC. This behavior has been characterised by Jedrkiewicz *et al.* and Blanchet *et al.* , who performed sub-shot-noise correlation measurements of the difference of conjugate

photons in a SPDC field [105, 120]. The two groups performed their experiments in the medium average-photon-number regime (about 10 photons per pixel) and the low-gain regime (about 0.15 photons per pixel), respectively.

The SPDC process occurs inside a $\chi^{(2)}$ crystal through a non-linear effect, in which one pump photon can turn into two photons that are correlated in both the energy and the momentum degrees of freedom. The momentum of the pump photon is conserved in the process, such that the two photons of a given pair leave the crystal with opposite transverse momentum, creating a ring in the far field (as shown later in Fig. 5.3b). At the far field, the two photons of a pair therefore arrive at opposite locations of the SPDC ring: $\mathbf{p} = (x, y)$ and $\mathbf{q} = (-x, -y)$. The coordinates x and y have their origin in the center of the SPDC ring. The probability distribution characterising the difference k between the number of photons arriving at the two locations \mathbf{p} and \mathbf{q} in a given time window is expressed as $P(n_p - n_q = k)$. In ideal experimental conditions, the two photons of every pair are always detected at these locations, thus leading to an ideal distribution characterised by $P(k = 0) = 1$ (equal number of photons measured at the two locations).

Because of unavoidable losses and noise in experimental conditions, the probability distribution $P(k)$ widens, resulting in an increased variance $\sigma^2(n_p - n_q)$. The variance being defined as

$$\sigma^2(n_p - n_q) = \langle [(n_p - n_q) - \langle n_p - n_q \rangle]^2 \rangle, \quad (5.3)$$

Given a spreading of the distribution $P(k)$ around its null mean $\langle n_p - n_q \rangle = 0$, the variance of $P(k)$ is given by

$$\sigma^2(n_p - n_q) = \langle (n_p - n_q)^2 \rangle. \quad (5.4)$$

In two-photon experiments highlighting sub-shot-noise statistics of the source, a variance $\sigma^2(n_p - n_q)$ lower than the mean photon number on the two pixels must be measured [105]:

$$\sigma^2(n_p - n_q) < \langle n_p + n_q \rangle. \quad (5.5)$$

5.2.2 Measured photon number statistics

To measure coincidence and correlations using an EMCCD camera, we image the SPDC ring in the far field, such that we can measure photons arriving at opposite locations $\mathbf{p} = (x, y)$ and $\mathbf{q} = (-x, -y)$ on the ring, with x and y being discrete pixel numbers. Here, we calculate the variance of the photon number statistics $P(n_p - n_q = k)$ by not only considering that statistics of the source, but also both the efficiency of the EMCCD camera (losses) and its excess noise. The excess noise is a source of error induced by the EMCCD camera that is typically at least equal to the shot noise. It arises in the measurement of detected photon number and is due to the stochastic process of multiplying charges [121].

Another known type of error in an EMCCD measurement is due to clock induced charges, but as they are always much lower than the signal in our experiment, we neglect these. Our model is therefore summarised as follow:

$$\text{Source} \xrightarrow[\text{Poisson}]{\text{Pairs}} N \xrightarrow[\text{Binomial}]{\text{Losses}} N_p, N_q \xrightarrow[\text{Poisson}]{\text{Excess noise}} n_p, n_q.$$

Here, it is important that the crystal be used in the low gain regime, such that the probability of generating high order pairs ($|2\rangle|2\rangle, |3\rangle|3\rangle$, *etc*) is negligible compared to the probability of generating first order pairs ($|1\rangle|1\rangle$). In the low gain regime, one can produce entangled photons following the statistics of the pump laser, that is Poissonian in most cases.

We first model the photon statistics of the generated photons at the source by a Poisson process, as we work in the low-gain SPDC regime. The probability that N photon pairs are incident on pixels q and p is therefore given by

$$P(N) = \frac{\exp(-\langle N \rangle) \langle N \rangle^N}{N!}, \quad (5.6)$$

where $\langle N \rangle$ is the average number of photons generated by the source. For a number N of incident photons on pixels p and q , only N_p and N_q are detected, in accordance with the efficiency η of the camera. We model these losses as a binomial distribution, as each photon can either be detected (with probability η) or lost (with probability $1 - \eta$). The probability that N_p photons are detected at pixel p given that N pairs were generated is given by

$$P(N_p|N) = \frac{N!}{N_p!(N - N_p)!} \eta^{N_p} (1 - \eta)^{(N - N_p)}, \quad (5.7)$$

The excess noise then introduces an additional Poisson-type error on the measurement. For a number N_p of detected photons at pixel p , the probability that the camera reads out n_p photons is given by

$$P(n_p|N_p) = \frac{\exp(-N_p) N_p^{n_p}}{n_p!}. \quad (5.8)$$

As this process is independent at pixel p and q , the probabilities that N_q photons are detected and the camera reads out n_q photon at pixel q are given by Eq. 5.7 and 5.8, with the subscripts p replaced by q .

We can now express the conditional probability $P(n_q, n_p)$ of measuring both n_p at pixel p and n_q at pixel q , given the above processes:

$$P(n_p, n_q) = \sum_{N_p, N_q, N} P(n_p, n_q | N_p, N_q) P(N_p, N_q | N) P(N), \quad (5.9)$$

where all sums are performed from 0 to ∞ , and $P(N_p, N_q|N) = 0$ when $N < \max(N_p, N_q)$. Since the processes are independent at pixels p and q , we can write

$$P(n_p, n_q|N_p, N_q) = P(n_p|N_p)P(n_q|N_q) \quad (5.10)$$

$$P(N_p, N_q|N) = P(N_p|N)P(N_q|N). \quad (5.11)$$

We now have the tools to express the variance of the measured photon statistics:

$$\sigma^2(n_p - n_q) = \sum_{n_p, n_q} \langle [(n_p - n_q) - \langle n_p - n_q \rangle]^2 \rangle P(n_p, n_q), \quad (5.12)$$

with $\langle n_p - n_q \rangle = 0$. Using $P(n_p, n_q)$ expressed with Eq. 5.9 - 5.11, we obtain

$$\begin{aligned} \sigma^2(n_p - n_q) &= \sum_N P(N) \sum_{N_p, N_q} P(N_p|N)P(N_q|N) \\ &\quad \sum_{n_p, n_q} P(n_p|N_p)P(n_q|N_q)(n_p - n_q)^2. \end{aligned} \quad (5.13)$$

Using Mathematica for solving Eq. 5.13, we first find the result for the sum on (n_p, n_q) :

$$\sigma^2(n_p - n_q) = \sum_N P(N) \sum_{N_p, N_q} P(N_p|N)P(N_q|N)[(N_p + N_q) + (N_p - N_q)^2]. \quad (5.14)$$

We evaluate the second sum on (N_p, N_q) and find

$$\sigma^2(n_p - n_q) = \sum_N P(N)[2N\eta(2 - \eta)]. \quad (5.15)$$

With $\langle N \rangle = \sum_N P(N)N$, we finally find that

$$\sigma^2(n_p - n_q) = 2\eta \langle N \rangle (2 - \eta). \quad (5.16)$$

On the other hand, the average photon counts is given by $\eta \langle N \rangle$, so that

$$\langle n_p + n_q \rangle = 2\eta \langle N \rangle. \quad (5.17)$$

We therefore find that the value of the variance is always above the shot-noise-limit for $\eta < 1$:

$$2\eta \langle N \rangle (2 - \eta) > 2\eta \langle N \rangle \quad (5.18)$$

$$\sigma^2(n_p - n_q) > \langle n_p + n_q \rangle. \quad (5.19)$$

This result indicates that a direct measurement of the photon statistics with the EM-CCD camera cannot yield sub-shot-noise behavior. To observe sub-shot-noise statistics, one needs to devise ways to take into account the excess noise, which we discuss in the following subsection.

5.2.3 Sub-shot-noise figure of merit

There are particular conditions where the excess noise can be suppressed through thresholding, when the light levels are very low (~ 0.15 photon per pixel) [120]. It is otherwise inevitable in higher flux regime (> 1 photon/pixel), so it needs to be taken into account by subtracting its effect from the variance $\sigma^2(n_p - n_q)$ [105]. Here, we propose a metric that is device-independent and can be used to demonstrate the capabilities of an EMCCD camera in both the low and high flux regime.

We define a figure of merit κ that compares the variance $\sigma^2(n_p - n_q)$ of two correlated pixels, with expected sub-shot-noise behavior, to the variance $\sigma_{\text{uncorrelated}}^2$ of two uncorrelated pixels, with expected Poissonian statistics:

$$\kappa = \left\langle \frac{\sigma^2(n_p - n_q)}{\sigma_{\text{uncorrelated}}^2} \right\rangle. \quad (5.20)$$

This figure of merit is similar to the Q_M and Q_B parameters in [122] in that they all measure statistical properties of light at the quantum mechanical level.

Since $\sigma_{\text{uncorrelated}}^2$ is the variance from a shot-noise-limited process, a source characterised by sub-shot-noise fluctuations measured with the same device will lead to a measured value of $\kappa(\text{sub-shot-noise}) < 1$, which we demonstrate in the following development.

To obtain a measurement from uncorrelated pixels, we measure the photon numbers at the same opposite pixels p and q that we did for the correlations measurement, but now in two separate, independent frames. We therefore measure the variance $\sigma^2(n_p - m_q)$ of the difference between the number of photons n_p at pixel p in a first frame and the number of photons m_q at pixel q in a second, independent frame. Our figure of merit can then be expressed as:

$$\kappa = \left\langle \frac{\sigma^2(n_p - n_q)}{\sigma^2(n_p - m_q)} \right\rangle, \quad (5.21)$$

with the $\langle \rangle$ operation indicating an average over many frames.

To calculate the shot-noise limited variance $\sigma^2(n_p - m_q)$, similarly to the development of $\sigma^2(n_p - n_q)$ above, we attribute a different number of generated photon pairs (N and

M with the same average $\langle N \rangle = \langle M \rangle$) to the two different frames:

$$\begin{aligned} \sigma^2(n_p - m_q) &= \sum_N P(N) \sum_M P(M) \sum_{N_p, M_q} P(N_p|N)P(M_q|M) \\ &\quad \sum_{n_p, m_q} P(n_p|N_p)P(m_q|M_q)(n_p - m_q)^2. \end{aligned} \quad (5.22)$$

We find from simplifying Eq. 5.22, with the help of Mathematica, and assuming the average photon number is equal in the two frames ($M = N$), that

$$\sigma^2(n_p - m_q) = 4\eta \langle N \rangle. \quad (5.23)$$

In the case of photon pairs generated through SPDC and measured with an EMCCD camera, we find that the figure a merit is

$$\kappa = \frac{\sigma^2(n_p - n_q)}{\sigma^2(n_p - m_q)} = 1 - \frac{\eta}{2}, \quad (5.24)$$

where η is the camera efficiency and represents the losses ($\eta < 1$), and the factor 2 arises from the excess noise. This figure of merit is therefore device independent in the sense that a sub-shot-noise source will always yield a value of $\kappa < 1$, independently of the camera calibration and the presence of detector noise.

5.2.4 Coincidence detection

We show here a method to measure the number of coincidences using the data recorded with the EMCCD camera. The coincidence counts for a source generating an average photon pair number $\langle N \rangle$ where each photon of the pair is detected with efficiency η [123] is given by

$$c = \eta^2 \langle N \rangle. \quad (5.25)$$

Since we cannot directly measure coincidences using time tags with an EMCCD camera, as it cannot be gated, we employ an indirect approach to the coincidence counts measurement based on the covariance between pixels p and q . The covariance of the two pixels is expressed as:

$$\text{cov}(n_p, n_q) = \langle n_p n_q \rangle - \langle n_p \rangle \langle n_q \rangle. \quad (5.26)$$

We first calculate the cross-correlations between the two pixels, or the average product of the counts in the two correlated pixels:

$$\langle n_p n_q \rangle = \sum_{n_p, n_q} n_p n_q P(n_p, n_q) \quad (5.27)$$

With a similar development as detailed in equations 5.9 to 5.17, we find that

$$\langle n_p n_q \rangle = \eta^2 \langle N^2 \rangle, \quad (5.28)$$

We can now rewrite Eq. 5.26 using Eq. 5.28 and the fact that the average number of detected photons is $\langle n \rangle = \eta \langle N \rangle$:

$$\text{cov}(n_p, n_q) = \eta^2 \langle N^2 \rangle - \eta^2 \langle N \rangle^2 = \eta^2 (\langle N^2 \rangle - \langle N \rangle^2). \quad (5.29)$$

Since the photon pair number follows a Poissonian statistics, the variance of the photon pair number is equal to its mean:

$$\langle N^2 \rangle - \langle N \rangle^2 = \langle N \rangle. \quad (5.30)$$

We thus find that the covariance of two correlated pixels gives us a measurement of the coincidence counts measured on these two pixels:

$$c = \text{cov}(n_p, n_q) = \eta^2 \langle N \rangle, \quad (5.31)$$

where the averages are performed over many frames. Crucially, we find that the presence of excess noise does affect this measurement.

5.3 Experimental results

We present here the experimental implementation of the coincidence counts and photon number statistics measurement and discuss the obtained results.

5.3.1 Experimental setup

To generate entangled photon pairs through parametric downconversion, we use a 7 mW 406 nm diode laser (Cobolt) to pump a 3-mm-long type-I beta-barium borate (BBO) crystal. The choice of a low pump power ensures that we operate in the low gain regime and obtain a photon pair source characterised by Poissonian photon statistics. The phase mismatch parameter, which determines the solid angle of the SPDC ring, is set to approximately $\phi = -4$. The EMCCD camera (Andor iXon 897) is positioned in the far field of the center of the crystal to image the SPDC ring, where the photons of a pair are

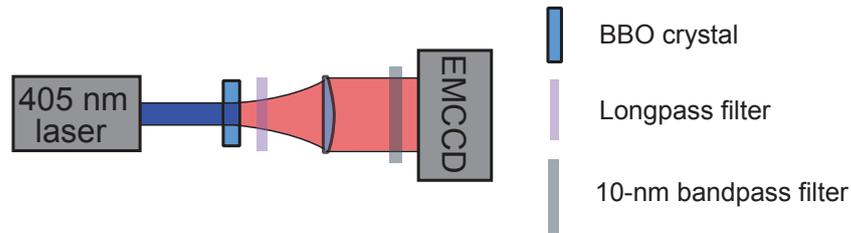


FIGURE 5.1: **Entangled photon pair generation.** Experimental setup used to generate photon pairs through SPDC and measure sub-shot-noise correlations with an EMCCD camera in the far-field of the crystal.

anti-correlated in their momenta, using a 100 mm lens. The entangled photon source is filtered from the pump and ambient light using a 10 nm bandpass filter centered at 812 nm. Fig. 6.2 illustrates the experimental setup.

This experimental procedure is similar to those outlined in references [97, 105, 109, 120], but the tools for measuring coincidence rates and sub-shot-noise statistics developed in section 5.2 allow us to perform measurements even when the average photon number per pixel is much greater than unity.

To measure spatial correlations and sub-shot-noise fluctuations for a wide range of average photon number per pixel, we control the number of photons incident on one pixel by changing the exposure time of the camera. Perhaps unsurprisingly, we find a linear relationship between exposure time and mean photon number per pixel. We set the exposure time between 100 μ sec to 1 sec, to span a range of coincidences per pixel from about 0.001 to 100. The relationship between exposure time and coincidence number is shown in Fig. 5.3 in the following subsection. The fact that this relation turns out to be non-linear is more surprising, but can be explained. At the low exposure times, the camera is dominated by noise events, and at high exposures, some pixels are saturated.

5.3.2 Sub-shot-noise statistics and coincidence number measurements

We record the SPDC ring at varying exposure times on the EMCCD, as detailed above, and measure the value of our figure of merit κ for correlated pixels both inside and outside the SPDC ring, to compare the photon number statistics of the SPDC source to the noise level captured on the EMCCD. The results are given in Fig. 5.2a, where each data point on the graph corresponds to a value of κ averaged over 200 frames, for the signal to surpass the background noise (Fig. 5.2b). We observe that the figure of merit shows that the signal outside the SPDC ring is shot-noise-limited, as expected. The measured statistics inside the SPDC ring show close to shot-noise-limited behavior for exposure times lower than 1 ms, because the strength of the measured correlations is reduced due to the importance of background noise and blooming (charges spilling over neighboring pixels) at those levels. At higher exposure time, we measure sub-shot-noise behavior from the SPDC

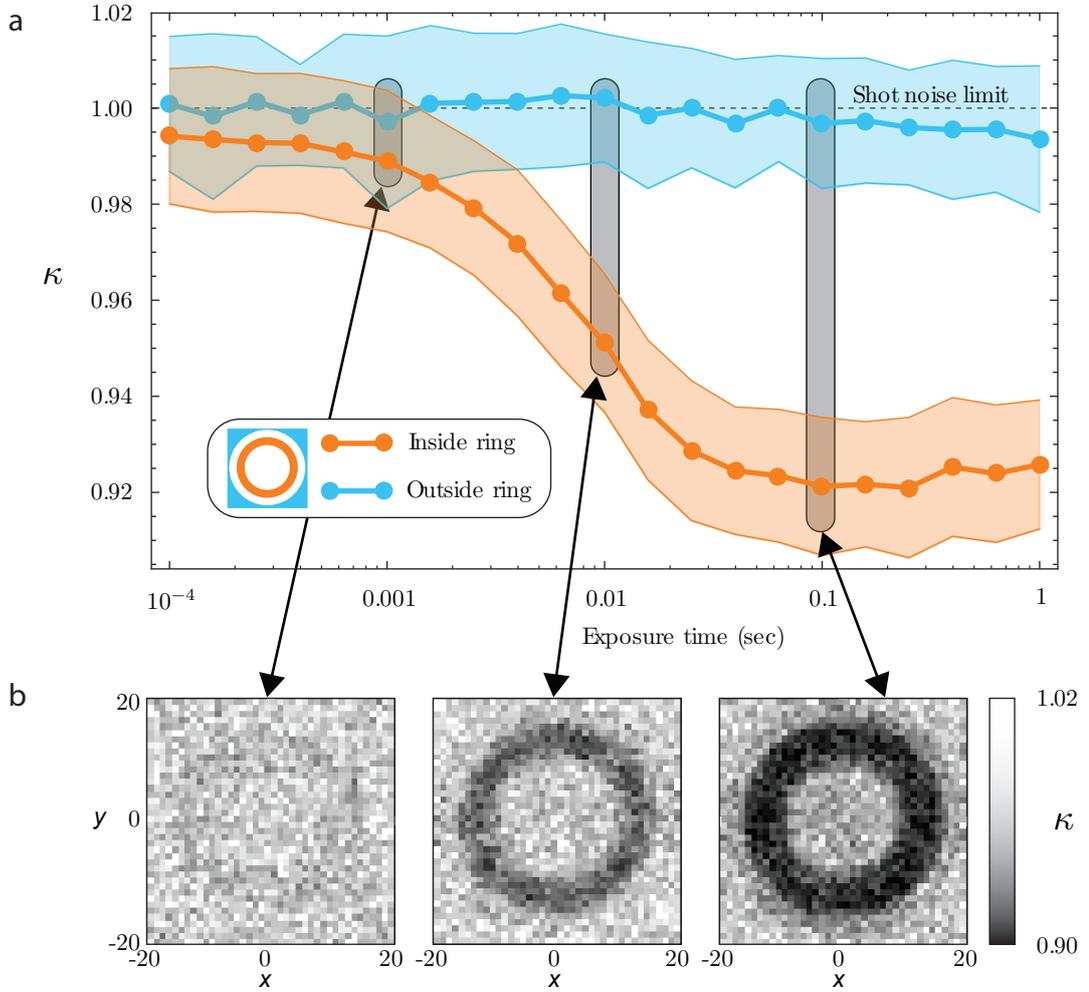


FIGURE 5.2: Spatial sub-shot-noise correlations. **a** The shot-noise limit is expressed by a value of $\kappa = 1$ and is indicated by a dashed line. The blue curve shows that there are no correlations outside the SPDC ring. The orange curve shows that the correlations inside the ring are characterised by sub-shot-noise signal. The light-colored regions around both solid lines represent the standard deviations on the results. **b** The figure of merit is calculated over the whole image and is shown for three different exposure times. The sub-shot-noise signal gradually increases with exposure time and makes its way above the background noise.

source, with κ reaching values of ~ 0.92 . At exposure times higher than one second, the saturation of the camera prevents us from detecting sub-shot-noise fluctuations accurately. Fig. 5.2b clearly illustrates that the measured sub-shot-noise behavior is concentrated inside the SPDC ring.

We also measure the mean number of coincidences per frame $C = \sum_{x,y} c(x,y)$ for varying exposure time, or varying mean photon number per pixel, by measuring the covariance of every correlated pixel (see Eq. 5.31). We sum the coincidences $c(x,y)$ only over the pixels that are illuminated by the SPDC ring. We show in Fig. 5.3 the mean number of coincidences per frame as a function of exposure time, which is also linear with the mean photon number per pixel $\langle n \rangle = \sum_{x,y} n(x,y) / \sum_{x,y} 1$.

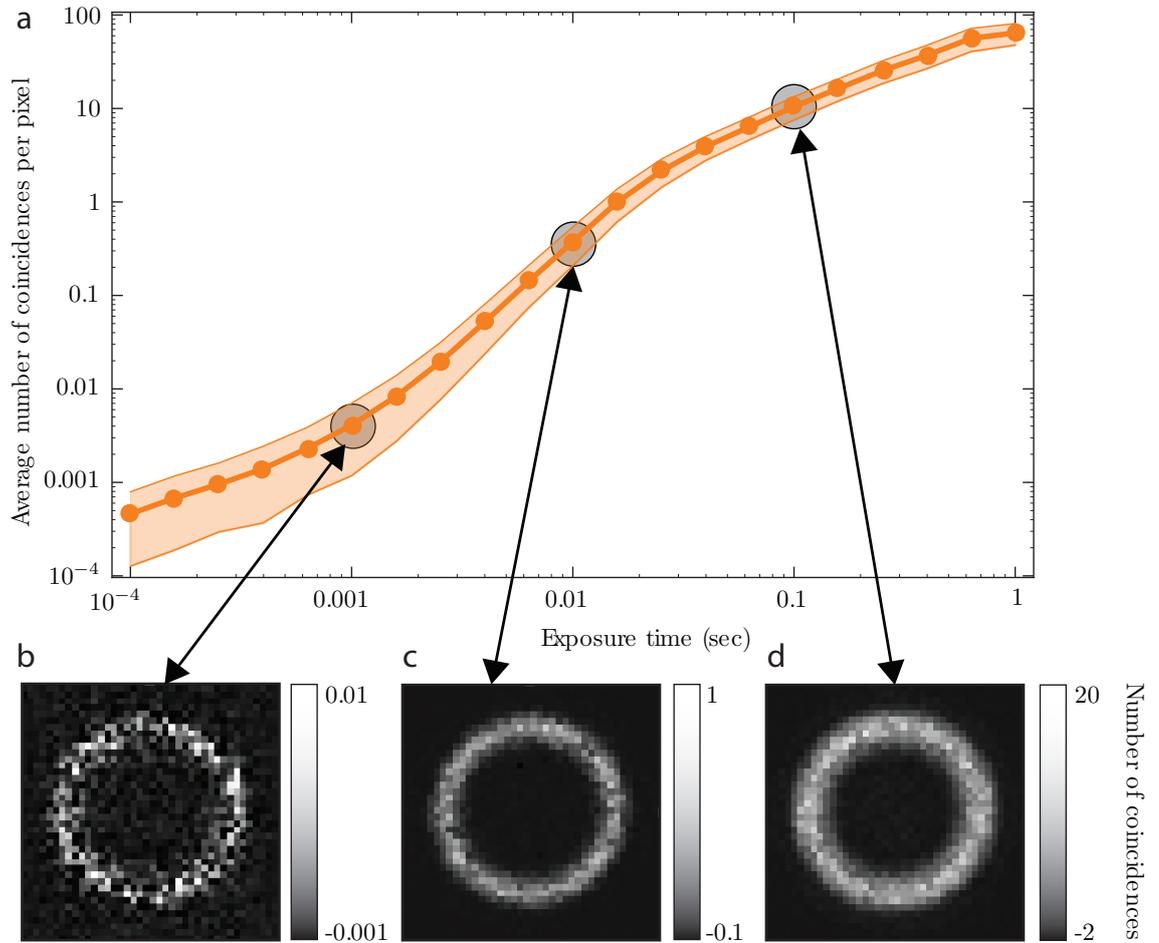


FIGURE 5.3: Coincidence measurement **a** For a given exposure, we sum the number of coincidences over every correlated pixels illuminated by the SPDC ring. At exposures below 1 ms, the measurement is dominated by blooming noise, whereby saturated pixels spill charges onto its neighboring pixels. **b** At an exposure of 1 ms, we observe granularity in the image of the measured coincidences. In these conditions, we detect on the order of one coincidence per frame. **c** The image of the ring gets smoother as we increase the exposure time of the camera to 10 ms. This corresponds to the linear regime between number of coincidences per frame and exposure time. A two-fold increase in exposure time yields close to a two-fold increase in the number of coincidences per frame. **d** This one-to-one linear relationship disappears at exposure times greater than 100 ms; the ring approaches the limit where it is perfectly smooth due to saturation, which prevents an accurate measurement of the coincidences.

We observe a relationship between number of coincidences and mean photon number that is close to linear, which demonstrates the dynamic range of the EMCCD camera in the context of quantum correlation measurements. Thanks to the wide dynamic range of the EMCCD, we can detect coincidences spanning four orders of magnitude, from 1 to 10,000 per frame. As mentioned above, we reach a limitation in both the low photon number regime, due to the background noise and clock induced charges, and the high photon number regime, due to the saturation of the camera. These two phenomena explain the non-linear sections of the coincidence number vs exposure time curve as observed in Fig. 5.3.

5.3.3 Discussion

The above results illustrate that state-of-the-art EMCCD cameras can be used to measure quantum correlations in a wide range of conditions. Establishing the presence of quantum phenomena with such techniques is fundamental to quantum science.

Although we do not directly record sub-shot-noise correlations in the standard fashion, our method provides a way to characterise a sub-shot-noise source as no shot-noise limited source will yield a value of κ , our figure of merit, lower than one. This way, we can detect sub-shot-noise correlations from a source even considering that the detector adds additional noise. In our experiments, the lowest value of κ measured with the EMCCD is 0.92. In the theoretical formalism we developed, we found κ to be given by $1 - \eta/2$, with the factor 2 arising from the excess noise of the EMCCD camera. Using a single photon avalanche detector (SPAD), lower values of κ could be achieved as this type of detector does not introduce excess noise. However, a SPAD presents the limitation of being constrained to measuring local correlations, compared to the ability of the EMCCD camera to detect the complete SPDC field and spatial correlations. On another hand, our experimental results are also limited by the heralding efficiency η . The efficiency we measure in our experiment is consistent with efficiencies reported in prior work with similar technology [97], but maximising the heralding efficiency would be key to future experiments that use this method for quantum measurements.

An other important advantage of EMCCD cameras for quantum measurements is their high sensitivity and large dynamic range, which allowed us to measure more than 100 photons per pixel. This performance is greater than what can be achieved with photon-resolving single pixel detectors, which can resolve 1 to about 25 photons [124–127].

5.4 Conclusion

In this chapter, we have demonstrated that we can measure spatial correlations with an EMCCD camera in regimes where the mean photon number per pixel is much greater than one. By modelling the generation and detection process, including efficiency and excess noise from the camera, we provided an analytical method to estimate, from the EMCCD measurements, the number of detected coincidences per frame. We have also demonstrated that we can measure spatial sub-shot-noise correlations without performing any background subtraction. The approach we developed allows us to measure photon number correlations for a wide range of input photon numbers. We foresee that this new approach will find uses in measuring and characterising quantum phenomena, such as sub-shot-noise fluctuations in a source, whether a low or high number of photons are generated [128]. It will also find applications in quantum imaging [99, 129].

Chapter 6 – Low-light-level image denoising

This chapter is about work that I initiated following discussions with Dr Jonathan Leach and Prof Miles Padgett about Monte-Carlo methods for denoising images acquired with single photon avalanche diode (SPAD) cameras. I improved upon Monte-Carlo methods by devising a new non-iterative algorithm for image denoising; I also benchmark this algorithm with a solution of the total-variation algorithm that uses the gradient descent tool detailed in Chapter 4. Most of the work presented in this chapter is theoretical, and, in addition to numerical results, we present a demonstration of the algorithm on experimental images recorded by Susan Chan and Dr Jonathan Leach.

Like the problems discussed in Chapters 3 and 4, the problem discussed here is an inverse problem. We possess statistical data about a scene, and we want to reconstruct this scene with the highest accuracy possible in as little time as possible. In particular, we have information about the intensity on each pixel of an image, but since the image is acquired at low light level, this information is plagued with shot noise. However, we know that the distribution of shot noise is Poissonian and we use this knowledge to our advantage through maximum likelihood estimation. In this chapter, we devise a denoising method and compare its performance with two existing techniques. We also apply the methods to images recorded experimentally using a single photon avalanche diode (SPAD) camera in a low-light-level settings. The manuscript that reports on this work is under preparation.

6.1 Introduction

State-of-the-art camera technology enables the detection of single photons. A few examples of such technology include, but are not limited to, ICCD, EMCCD and SPAD cameras [129–131]. It is always the case that the result of acquiring an image composed of single photon events is affected by shot noise. One way of mitigating the effect of shot noise is through squeezing, a quantum effect that is experimentally challenging to implement [130]. A more common way to improve over shot-noise-limited images is through implementing prior knowledge about natural images into a post-processing algorithm [129, 132–140]. Notably, low-light-level image denoising algorithms prove to be useful in astronomy [141, 142], ghost imaging [129], video surveillance [143], CT scans [144] and a wide range of fluorescent microscopy applications.

Image denoising involves finding the image most consistent with the data given some prior knowledge about both the image and the data. Prior knowledge about a natural

image can be, for example, the fact that it has a sparse representation in the discrete cosine basis (DCB) [145], that the variations from one pixel to its neighbors, the sum of which is known as total variation (TV), are rather small [146–148] or that it was acquired under blurring conditions [132, 141]. As for the data, we consider the case where, having the scene under a constant illumination level, the statistics of the measured number of photons per pixel follow a Poisson distribution, which models accurately imaging regimes with low-photon-number outcomes.

We seek to solve the following inverse problem: given the observations $N_{i,j}$ at each pixel position (i, j) , what is the underlying average intensity $M_{i,j}$? If we have no prior information about the image itself, then the trivial solution is $M_{i,j} = N_{i,j}$. However, when we add prior information to the model, the most likely image given this information can be found through a maximum likelihood analysis (see Chapter 4 for details on maximum likelihood analysis).

In this work, special care is taken to optimize the algorithms’ speed, such that the algorithms detailed here could eventually be running in real time in future applications. We also show the specific tradeoff between recovered image quality and the computational time spent on post-processing.

6.2 Theoretical development

Here, we show the theoretical development for TV-minimisation, which is an existing denoising technique that we use as benchmark. We then tackle the main theoretical result: a non-iterative algorithm for denoising using concepts of discrete cosine basis (DCB) sparsity.

The acquired image N is a $r \times c$ -dimensional matrix of integers corresponding to photon number, where r and c is the number of rows and columns, respectively. Given a matrix M that represents an estimate of true image, the probability distribution of the photon number $N_{i,j}$ at pixel (i, j) is given by [129, 137]

$$P(M_{i,j}; N_{i,j}) = \frac{(M_{i,j} + \epsilon)^{N_{i,j}} \exp^{-(M_{i,j} + \epsilon)}}{N_{i,j}!}, \quad (6.1)$$

where ϵ is the average background to the image. We build the likelihood function

$$\prod_j P(M_{i,j}; N_{i,j}), \quad (6.2)$$

such that the most likely image maximises this function. We employ the standard

strategy of minimising the negative log-likelihood in the pixel basis (PB)

$$\mathcal{L}_{\text{PB}}(M) = -\log \prod_j P(M_{i,j}; N_{i,j}) \quad (6.3)$$

$$= \sum_{i,j} M_{i,j} + \epsilon - N_{i,j} \log(M_{i,j} + \epsilon) + \log N_{i,j}! \quad (6.4)$$

This function is an integral part of many Poisson denoising algorithms; it provides the cost function for finding a guess M that is different from the raw data N . The derivative of Eq. 6.4 with respect to $M_{i,j}$ is given by

$$\frac{d\mathcal{L}_{\text{PB}}}{dM_{i,j}} = 1 - \frac{N_{i,j}}{M_{i,j} + \epsilon}. \quad (6.5)$$

It is important that the background level ϵ is non-zero if one wants to avoid divisions by zero; this parameter also accounts for dark noise on the detectors. For compactness, we will henceforth omit the indices i, j , such that Eq. 6.5 is written

$$\frac{d\mathcal{L}_{\text{PB}}}{dM} = 1 - \frac{N}{M + \epsilon}, \quad (6.6)$$

where both the derivative and the division are interpreted as element-wise.

6.2.1 TV-minimisation

Natural images typically vary very little from one pixel to the next. This is the underlying assumption behind the anisotropic total variation [146–148]

$$TV(M) = \sum_{i,j} |M_{i+1,j} - M_{i,j}| + |M_{i,j+1} - M_{i,j}|, \quad (6.7)$$

which is smooth everywhere except when the argument of any absolute value is zero. Assuming the absolute value arguments are greater than some constant μ , the derivative of Eq. 6.7 is given by

$$\begin{aligned} \frac{dTV}{dM_{i,j}} = & \text{sign}(M_{i,j} - M_{i-1,j}) + \text{sign}(M_{i,j} - M_{i,j-1}) - \\ & \text{sign}(M_{i+1,j} - M_{i,j}) - \text{sign}(M_{i,j+1} - M_{i,j}), \end{aligned} \quad (6.8)$$

where the sign function yields unity when its argument is positive, -1 when its argument is negative and 0 when the argument is zero. When an absolute value argument is below μ , we render Eq. 6.7 differentiable using Huber smoothing, whereby the low part of the absolute value function (around argument zero) is approximated by a quadratic function [149].

We define the matrices M_h^+ (M_h^-) and M_v^+ (M_v^-) as the shifted matrix M by one positive (negative) unit horizontally and vertically, respectively. The gradient of Eq. 6.7 is thus given by

$$\frac{dTV}{dM} = \text{sign}(M - M_v^+) + \text{sign}(M - M_h^+) + \text{sign}(M - M_v^-) + \text{sign}(M - M_h^-). \quad (6.9)$$

The cost function associated with anisotropic total variation is given by

$$f_{TV}(M, N) = \mathcal{L}_{PB} + \beta TV, \quad (6.10)$$

where β is a variable that is optimised through trial and error. Given the Huber smoothing, this cost function is differentiable and descending along its gradient is a numerically stable way of minimising it. The gradient of the total cost function is

$$\frac{df_{TV}(M, N)}{dM} = \frac{d\mathcal{L}_{PB}}{dM} + \beta \frac{dTV}{dM}. \quad (6.11)$$

To find the minimum of the total cost function, we use accelerated gradient descent with momentum. To do so, we define a momentum matrix $V^{(k)}$, which is initialised with the null matrix: $V^{(0)} = 0$. The update rule to $V^{(k)}$ is given by [91, 150, 151]

$$V^{(k+1)} = \alpha V^{(k)} - \gamma \frac{df_{TV}(M^{(k)} + \alpha V^{(k)}, N)}{dM^{(k)}}, \quad (6.12)$$

where α is the momentum term and γ is a learning rate. The update rule to the guess image is simply the sum of the momentum and the current guess image

$$M^{(k+1)} = V^{(k)} + M^{(k)}. \quad (6.13)$$

The momentum term adds inertia, i.e. memory of previous update directions to the algorithm; the closer α is to unity, the longer term the memory is. The greater the learning rate is, the faster the algorithm should reach the saturation regime, but the more likely it is to produce numerical instabilities. The argument of the gradient, $M^{(k)} + \alpha V^{(k)}$, is such that a step in the momentum direction is anticipated. This gradient descent approach is very similar to the projected gradient descent (PGD) algorithm detailed in Chapter 4.

6.2.2 Shrinking values in the discrete cosine basis

We provide the main theoretical result of the Chapter in this subsection. Since the following derivation is rather involved, we separate all the steps into subsection with bolded titles.

Theoretical background

Compared to random images, natural images are known to have more nearly-null components in the discrete cosine basis (DCB). To exploit this prior knowledge, one would ideally attempt to minimise the number of non-zero components in the DCB, but this strategy leads to a non-differentiable cost function. Typically, one relaxes the problem by minimizing the 1-norm of the image in the discrete cosine basis.

The image M expressed in the DCB is denoted A . There exists a pair of square unitary transform matrices U_r and U_c that can be used to perform a discrete cosine transform (DCT), such that $A = U_r M U_c^\dagger$ and $M = U_r^\dagger A U_c$. The elements (i, j) of a transform matrix of dimension d are given by

$$U_{d,(i,j)} = \sqrt{\frac{2}{d}} K \cos \frac{(2j+1)i\pi}{2d}, \quad (6.14)$$

where $K = 1/\sqrt{2}$ if $i = 0$ and $K = 1$ otherwise. To avoid overcrowding the notation, all operations are to be interpreted as element-wise, except when a basis change is implied, e.g. $A = U_r M U_c^\dagger$.

As mentioned above, a sensible measure of sparsity is given by the 1-norm

$$\|A\|_1 = \sum_{w,v} |A_{w,v}|, \quad (6.15)$$

which has a reputation for being hard to optimize because of its lack of smoothness [152, 153]. One can tackle this problem either by using sufficiently small step sizes in the steepest descent algorithm, by smoothing the cost function using a proximal method [132] or by finding an analytical solution. As we detail further, we take the approach of finding an analytical solution to the minimisation of the cost function. A minimisation in the pixel basis requires the derivative of Eq. 6.15 with respect to the pixel intensities, which is given by

$$\frac{d\|A\|_1}{dM} = U_r^\dagger \text{sign}(A) U_c, \quad (6.16)$$

where the derivative of a scalar with respect to a matrix is defined in the usual way; details of the derivative are given in the following subsection.

Eq. 6.16 involves a DCT and an inverse DCT. The computation of this derivative requires a time $\mathcal{O}(r^2c + rc^2)$ because of the matrix multiplications. It is likely that a fast way of performing the basis changes exists – similar to a Fast Fourier Transform – which would yield a computational complexity of $\mathcal{O}(rc \ln(rc))$. In any case, implementing Eq. 6.16 into the gradient descent algorithm would inevitably slow down the process, but this implementation is likely to give better results than TV-minimization alone. This

is an avenue of interest for future work.

DCB derivative

We detail here the steps to compute the derivative of Eq. 6.16. Let us define $Q = U_r \otimes U_c$ and the vectors $\mathbf{m} = \text{unwrap}(M)$ and $\mathbf{a} = \text{unwrap}(A)$, which corresponds to reshaped versions of the guess image in the pixel basis (M) and the DCB (A), respectively. Note that through the Kronecker product \otimes , one can replace a unitary transformation based on matrix multiplications by a matrix-vector operation: $\mathbf{a} = Q\mathbf{m}$. If \mathbf{q}_i^\dagger refers to line i of matrix Q , the 1-norm is

$$\sum_i |\mathbf{a}_i| = \sum_i |\mathbf{q}_i^\dagger \mathbf{m}|, \quad (6.17)$$

and its derivative with respect to an element j of \mathbf{m}

$$\frac{d \|\mathbf{a}\|_1}{d\mathbf{m}_j} = \sum_i \mathbf{q}_{ij} \text{sign}(\mathbf{q}_i^\dagger \mathbf{m}). \quad (6.18)$$

Replacing vector elements by entire vectors yields

$$\frac{d \|\mathbf{a}\|_1}{d\mathbf{m}} = Q^\dagger \text{sign}(Q\mathbf{m}). \quad (6.19)$$

Finally, using matrix notation again, we recover Eq. 6.16.

Fast algorithm for maximising sparsity in the DCB basis

We now suggest a fast algorithm for minimising the image 1-norm in the DCB rather than in the pixel basis by transferring the negative log-likelihood function to the DCB. This method requires a single shrinking step of the DCB image coefficients. This way of processing noisy images can easily be applied in real time. Variations of this approach are studied in References [133, 154], where images are decomposed into wavelets through a Haar transform, and in Reference [155], where the noise is assumed to be Gaussian. Here, we consider the cost function

$$f_{\text{DCB}}(A) = \mathcal{L}_{\text{DCB}} + \lambda \|A\|_1, \quad (6.20)$$

where \mathcal{L}_{DCB} is the negative log-likelihood of the guess matrix transferred to the DCB, and λ is optimised through trial and error. Working directly in the DCB is convenient because the derivative of the 1-norm is simply given by $\text{sign}(A)$. The challenge now lies in evaluating the likelihood function and its derivative.

Let us consider for simplicity a square image ($r = c$), and define $U_{\text{abs}} = |U|$, where the absolute value is applied element-wise. We define the following matrices in the discrete cosine basis:

$$S = U_{\text{abs}} M U_{\text{abs}}^\dagger \quad (6.21)$$

$$T = U N U^\dagger, \quad (6.22)$$

with T being the measured image transferred to the DCB. Since the matrix S is calculated using only positive numbers, its elements are generally much larger than those of T or A . The exception is $T_{1,1} = S_{1,1}$, which corresponds to the level of the background in the acquired image; we will ignore this discrepancy since it only affects one element.

We transfer the likelihood function $\mathcal{L}_{\text{PB}}(M)$ from the pixel basis to the DCB and get the negative log-likelihood cost function

$$\mathcal{L}_{\text{DCB}}(A) = - \sum_{w,v} \frac{T}{2} \ln \left(\frac{S+A}{S-A} \right) + \ln I_T(\sqrt{S^2 - A^2}), \quad (6.23)$$

where the indices w and v have been omitted on every matrix for compactness. Here, $I_{T_{w,v}}(\cdot)$ is the modified Bessel function of order $T_{w,v}$. This basis transfer involves calculating sums and differences of Poisson distributions according to the DCT. After simplifications and a first order Taylor expansion, the gradient of the likelihood cost in the DCB is given by

$$\frac{d\mathcal{L}_{\text{DCB}}(A_{v \ln, w})}{dA_{v,w}} \approx \frac{A_{v,w} - T_{v,w}}{S_{v,w}}, \quad (6.24)$$

where the division is performed element-wise. The details of this derivation are given in the following subsection, where we provide justification for using the above simple form for the derivative of \mathcal{L}_{DCT} . We also give a more accurate but more computationally intensive form for the above derivative. For natural images in general, the above approximation is sufficient.

Basis transfer

When transferring the data from the pixel basis to the DCB, the unitary matrix multiplies the data with positive and negative weights. Since one outcome has a Poisson distribution, the sum of positive outcomes thus yields a Poisson distribution. The same applies to the sum of negative outcomes. However, the weighted sum of multiple Poisson distributions does not yield a Poisson distribution unless the weights are solely taken from -1, 0 or 1. Unfortunately, the elements of the DCB matrix are real numbers between -1 and 1, which means that the resulting distribution is not really Poissonian. One choice of basis matrix that would indeed perfectly are Hadarmard matrices. The negative impact of the DCB matrix on the following derivation is not trivial, and is an avenue for future work.

We separate the outcomes with positive weights from those with negative weights by first defining the matrices of positive elements U^+ and negative elements U^- , such that

$$U_{v,w}^+ = \begin{cases} U_{v,w}, & \text{if } \text{sign}(U_{v,w}) = 1 \\ 0, & \text{if } \text{sign}(U_{v,w}) = -1 \text{ or } 0 \end{cases} \quad (6.25)$$

$$U_{v,w}^- = \begin{cases} U_{v,w}, & \text{if } \text{sign}(U_{v,w}) = -1 \\ 0, & \text{if } \text{sign}(U_{v,w}) = 1 \text{ or } 0. \end{cases}$$

We express the matrix A using the above positive and negative matrices:

$$A = U M U^\dagger \quad (6.26)$$

$$A = (U^+ + U^-) M (U^+ + U^-)^\dagger \quad (6.27)$$

$$A = U^{+\dagger} M U^+ + U^+ M U^{-\dagger} + U^- M U^{+\dagger} + U^- M U^{-\dagger}, \quad (6.28)$$

which can also be expressed as a combination of a matrix A^+ containing positive elements and a matrix A^- containing negative elements:

$$A = A^+ + A^- \quad (6.29)$$

$$\text{with} \quad (6.30)$$

$$A^+ = U^+ M U^{+\dagger} + U^- M U^{-\dagger} \text{ and} \quad (6.31)$$

$$A^- = U^+ M U^{-\dagger} + U^- M U^{+\dagger}.$$

Since the elements of both A^+ and A^- are Poisson distributions of opposite signs and the difference of two Poisson distributions is a Skellam distribution. The probability distribution associated with a measured image T and an estimate of the true image A is [156]:

$$P_{DCB}(T, A) = \text{Skellam}(T; A^+, A^-) = e^{-(A^+ - A^-)} \left(\frac{A^+}{-A^-} \right)^{\frac{T}{2}} I_T(2\sqrt{-A^+ A^-}), \quad (6.32)$$

where $T = U N U^\dagger$ is the directly acquired image transferred to the DCB. Again, we omitted writing the matrix subscripts $\{w, v\}$ for compactness. The elements of $A^+ - A^-$ follow a Poisson distribution. The estimates are all of the order of the total number of photons detected in the entire image, and thus change negligibly during an optimisation procedure. Replacing $A^+ - A^-$ with the constant matrix $S = U_{\text{abs}} N U_{\text{abs}}^\dagger$, defined in the main text, is highly accurate. Using this notation, we can rewrite Eq. 6.32 as [154]

$$\text{Skellam}(T, S; A) = e^{-S} \left(\frac{S - A}{S + A} \right)^{\frac{T}{2}} I_T(\sqrt{S^2 - A^2}). \quad (6.33)$$

The corresponding negative log-likelihood function ($-\ln \prod P_{DCB}$) is given by Eq. 6.23.

After simplification, the *exact* derivative of Eq. 6.23 is given by

$$\frac{df_{\mathcal{L}_2}}{dA} = -\frac{T}{S+A} + \frac{A}{\sqrt{S^2-A^2}} \frac{I_{T+1}(\sqrt{S^2-A^2})}{I_T(\sqrt{S^2-A^2})}. \quad (6.34)$$

Computing a Bessel function for each element can be computationally expensive. Hence, we use the following approximation [157]:

$$\frac{I_{\nu+1}(x)}{I_{\nu}(x)} \approx \frac{x}{\nu + 1/2 + \sqrt{x^2 + (\nu + 1/2)^2}}, \quad (6.35)$$

which is valid provided that $x + \nu > 0$. In our case, this criterion translates to

$$\sqrt{S_{w,v}^2 - A_{w,v}^2} + T_{w,v} > 0. \quad (6.36)$$

This is always true since, as explained above, the elements of S are much larger than those of A and T . From our observations, the approximation typically falls within 0.01% of the real value and the worst case occurs on element (1,1) for which the approximation is always within 1% of the real value. Using this approximation and a first order Taylor expansion, the derivative of Eq. 6.34 can be further simplified:

$$\frac{df_{\mathcal{L}_{\text{DCT}}}}{dA} = \frac{A-T}{S} + \mathcal{O}\left(\frac{T^2}{S^2}\right). \quad (6.37)$$

We thus retrieve Eq. 6.24.

Final equation

To minimize the cost function (Eq. 6.38):

$$f_{\text{DCB}}(A) = \mathcal{L}_{\text{DCB}} + \lambda \|A\|_1, \quad (6.38)$$

we analytically find the guess matrix A such that $df_{\text{DCB}}/dA = 0$, the solution of which is the shrinking operation

$$A_{v,w} = \begin{cases} T_{v,w} - \text{sign}(T_{v,w})S_{v,w}\lambda, & \text{if } |T_{v,w}| \geq S_{v,w}\lambda \\ 0, & \text{if } |T_{v,w}| < S_{v,w}\lambda. \end{cases} \quad (6.39)$$

We retrieve the final guess image through the inverse DCT:

$$M = U_r^\dagger A U_c. \quad (6.40)$$

This final result can contain pixels with negative values, but the amplitude of these negative values are typically very small because the problem is well-conditioned. An ill-conditioned problem such as denoising an image obtained under blurring conditions requires forcing positivity on the pixel values. This is what the Richardson-Lucy algorithm does [141].

The computational complexity of the shrinking operation itself is proportional to the number of pixels, but the bottleneck of the whole algorithm consists of calculating the S and T matrices in the first place with matrix multiplications, the computation of which scales as $O(r^2c + rc^2)$. However, only one shrinking step is required to maximize sparsity given Poisson-distributed data, which makes sparsity maximization about 10 times faster than TV-minimization via gradient descent for the considered image sizes. There is a threshold size for which we expect the latter to be faster than the former. Another difference between the two techniques is that whereas the TV algorithm requires an artificially high background ϵ , the shrink operation does not need any preprocessing of the data.

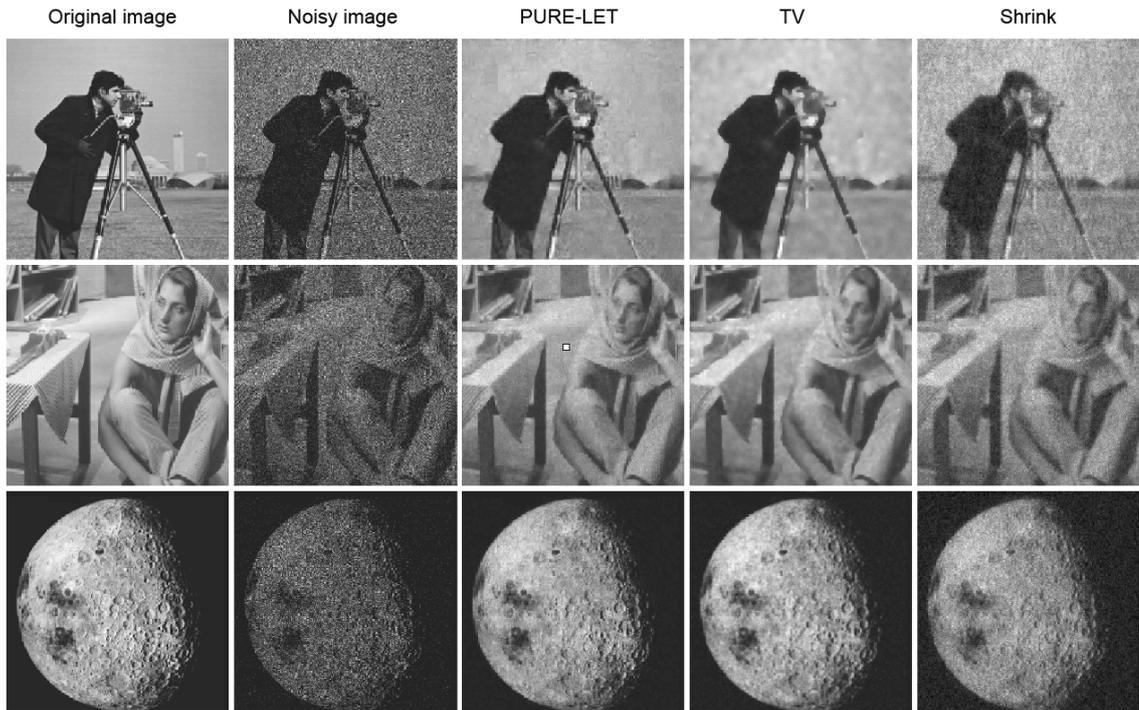


FIGURE 6.1: **Denoising simulations performed on images of different sizes: 256×256 (Cameraman), 512×512 (Barbara) and 700×700 (The Moon).** The images are rescaled to fit in this figure, but the simulations are performed on the original-size figures. The total number of photons in the noisy images of The Cameraman, Barbara and The Moon are 300 000, 1 200 000 and 780 000, respectively. The third, fourth and fifth columns correspond to the output of the following algorithms: PURE-LET [158], total variation (TV) minimization and DCB-sparsity maximization (Shrink). Refer to Table 6.1 for the PSNR and computing time of each result. In order of {Cameraman, Barbara, The Moon}, we use $\beta = \{0.4, 0.4, 0.7\}$ (TV) and $\lambda = \{0.003, 0.0015, 0.002\}$ (Shrink).

	Noisy image	PURE LET	TV	Shrink
Cameraman $r = c = 256$	26.9dB	57.6dB 0.288s	53.4 dB 0.053s	48.0dB 0.005s
Barbara $r = c = 512$	23.8dB	48.3dB 0.838s	47.3dB 0.230s	41.4dB 0.026s
The Moon $r = c = 700$	24.0dB	46.5dB 1.699s	47.0dB 0.473s	37.7dB 0.061s

TABLE 6.1: **PSNR and computing time for three Poisson denoising methods: PURE-LET from Ref. [158], total variations (TV) minimization and DCB-1-norm minimization (Shrink).** Simulations are performed in Matlab on an iMac i5 2.9GHz. When choosing an appropriate denoising method, there is a tradeoff between the quality of resulting denoised image and the speed at which it can be computed.

6.3 Numerical Results

We now perform numerical simulations to quantify the speed and performance of the two algorithms. The quality of the results are calculated via the peak signal-to-noise ratio

$$\text{PSNR}(M) = 20 \ln \frac{(2^b - 1)\sqrt{rc}}{\|M - M_{\text{actual}}\|_2}, \quad (6.41)$$

where b is the bit depth, $\|\cdot\|_2$ is the 2-norm and M_{actual} is the image without any noise.

We use the following values for the numerical simulation: $\alpha = 0.8$, $\gamma = 0.1$ and $k_{\text{max}} = 20$, where k_{max} is the maximum number of iterations. Higher values for γ can yield numerical instabilities. If the background level ϵ is small, artificially increasing it by adding a constant of about 3 to all pixels of the acquired image makes the algorithm more stable. The number of iterations required depends on the robustness of the derivatives to numerical instabilities.

The DCB is not necessarily the best basis to optimise in. Significant improvements can be made through trained dictionaries of basic images or learning the best basis for a given set of images [145]. All equations involving a discrete cosine transform with the U matrices could then be replaced by arbitrary transforms. Our work could then apply to any problem where the data is affected by Poisson noise and the signal is known to be sparse in a particular basis.

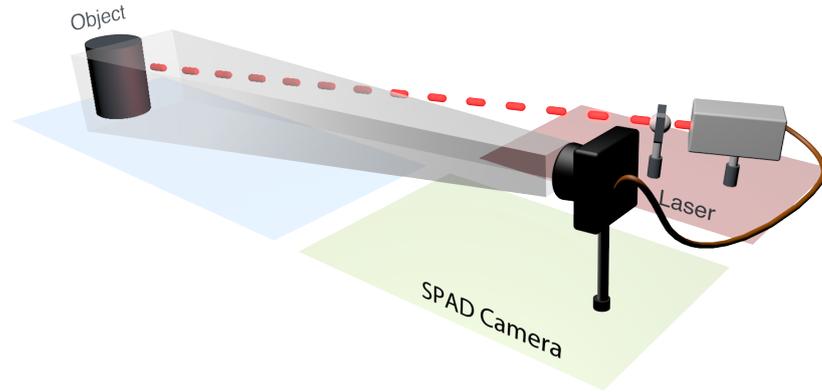


FIGURE 6.2: **Experimental setup consisting of a laser, a lens, the object (a mug) and a SPAD camera.** We varied the output power of the laser to control the level of Poisson-distributed shot noise on the images acquired with the SPAD camera. The object is illuminated uniformly using a lens. One raw image is obtained by adding 63 sequential frames together, and one frame is acquired over $5.5 \mu s$.



FIGURE 6.3: **Experimentally recorded reference image.** The image is recorded at high intensity. The noise in the image is caused by defective pixels.

	Noisy image	PURE LET	TV	Shrink
$15 \mu W$ $r = c = 256$	15.2dB	16.1dB 0.195s	18.1dB 0.063s	17.2 dB 0.004s
$5 \mu W$ $r = c = 256$	11.9dB	12.4dB 0.140s	14.6dB 0.058s	14.0dB 0.002s
$1 \mu W$ $r = c = 256$	11.27dB	11.3dB 0.140s	13.2dB 0.048s	12.8dB 0.003s

TABLE 6.2: **PSNR and computing time for experimentally recorded images at low light levels.** Computations are performed in Matlab on a single core of an i5 6600K 3.5GHz CPU. The reference image recorded with high-intensity illumination is shown in Fig. 6.3.

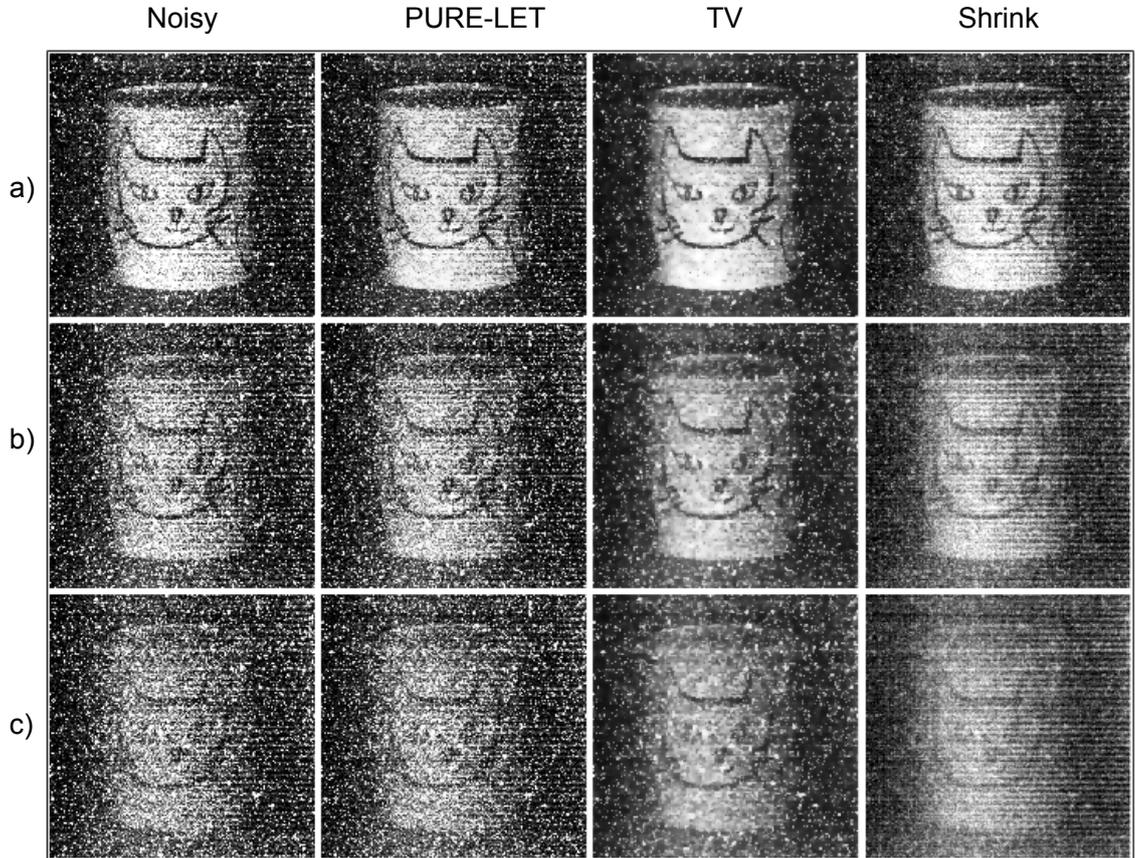


FIGURE 6.4: **Experimentally recorded and denoised images for a laser intensity of a) $15 \mu W$, b) $5 \mu W$ and c) $1 \mu W$.** The associated PSNRs and recovery times are shown in Table 6.2. After running a denoising algorithm, we rescale the noisy images to make sure that the number of photons is similar to that of the reference image recorded at high intensity (See appendix); this ensures that the PSNR calculated with respect to this reference image has a fair value. We also set every pixel with negative values to zero and every pixel above with values above 63 to 63 since the recorded images have a bit depth of 6. The parameters used for TV and Shrink are $\{k_{\max} = 20, \alpha = 0.9, \gamma = 0.2\}_{\text{TV}}$ for all rows and a) $\{\beta = 0.81, \lambda = 0.002\}$, b) $\{\beta = 1.03, \lambda = 0.0034\}$ and c) $\{\beta = 1.3, \lambda = 0.0042\}$ for the rest.

6.4 Experimental results

We put the Shrink algorithm to the test with a single-photon avalanche detector (SPAD) camera (240x320 pixels) in a low-light-illumination setting. This setting gives rise to significant shot noise, which is characterised by the Poisson-distributed photon numbers from Eq. 6.1. The experimental setup is shown in Fig. 6.2. Our Picoquant pulsed laser diode has a ~ 670 nm peak wavelength, a 5 MHz repetition frequency and illuminates a mug through a 1.45 mm focal length lens. The SPAD camera is triggered by the laser every time a pulse is produced. There is an internal delay to the SPAD camera to synchronise the incoming pulse with the electronic gate of the camera (11 ns). Our object (see Fig. 6.3) is imaged through a Tamron 70-200 mm lens at f/2.8. We build an image by acquiring 63 frames and summing them together. The bit depth is thus 8. Our goal is to obtain the highest PSNR between a reference image obtained in high illumination conditions and images obtained at lower light levels using the TV-minimisation algorithm.

Our setup produces an undesirable effect when lowering the light intensity: the average pixel value lowers with the laser power output. To be able to perform a fair comparison between the reference image and any other image acquired under lower light-level conditions, the average pixel values of the low-light-level image must be raised to that of the reference image. To this end, we scaled the low-intensity image up by multiplying it with a scaling factor κ . In addition, when rescaling an image, if the value of a pixel exceeds $2^6 - 1 = 63$, we reduce its value back to 63, that is we apply a hard maximum corresponding to the original 6-bit image. We simply try many values of κ and keep the one that maximises the PSNR.

We first record a reference image at high laser intensity, shown in Fig. 6.3. We then record noisy versions of the mug image for three laser power values, that is with different levels of shot noise, and recover denoised versions. These images are shown in Fig. 6.4. The PSNR of each image is shown in Table 6.2. The total variation minimization algorithm runs for an average time of 55 ms, which means that the denoising technique can be used at a rate of 20 images per second with the SPAD camera. This is enough for real-time imaging. Our own Shrink method take 3 ms to run on average, which is equivalent to about 300 frames per second. The Shrink algorithm could thus be applied to faster analytics, although the quality of the recovered image with the TV-minimisation algorithm is always better. Surprisingly, even after optimising the parameters of the PURE-LET method, the recovered image quality is not on par with the two other techniques. PURE-LET is better suited for low noise levels.

6.5 Conclusion

We have shown that there are principled ways to recover from noisy data an image that highly resembles the actual scene and that there is a clear tradeoff between processing time and image quality. The algorithm that we developed, Shrink, assumes sparsity of the image coefficients in the discrete cosine basis representation. This algorithm proved to work experimentally: using a SPAD camera in the dark we recovered images with significantly higher PSNR (or fidelity to the actual scene) than without total variation minimization. It proves to be significantly faster than the benchmark algorithm – TV-minimisation and PURE-LET – because it is non-iterative. Our technique coupled with the use of SPAD arrays could eventually be applied to denoising transient images in real-time [5].

Chapter 7 – Contributions to other published work

In addition to the work presented in Chapters 3 to 6, I performed during my PhD satellite work that lead to the contributions in the two following papers.

Symmetric states

I have a contribution to the following published paper:

M. Agnew, E. Bolduc, K.J. Resch, S. Franke-Arnold, J. Leach.

Discriminating single-photon states unambiguously in high dimensions.

Physical Review Letters **113**, 020501 (2014).

This work consists of an experimental demonstration of single-shot quantum state discrimination of non-orthogonal states in high dimensions. We used the orbital angular momentum degree of freedom of single photons to implement the protocol in the laboratory.

I was not involved in the experiment, but I was tasked to find $d - 1$ symmetric quantum states in a d -dimensional Hilbert space, which is a necessary element to the error-free discrimination protocol. This task is similar to that of finding a SIC-POVM, which consists of d^2 symmetric states in a d -dimensional space, although the former has a more straightforward solution. We presented in Sections 2.4.1 and 4.3.1 a more in-depth discussions of SIC-POVMs. To find the $d - 1$ states, I developed an iterative algorithm whereby the i^{th} state has i non-zero coefficients:

$$|\Psi\rangle_i = \sum_{j=0}^{i-1} c_j^{(i)} |j\rangle \quad (7.1)$$

We start by deriving the overlap of a state with any other state given the symmetry requirements: the solution leads to an overlap of $-1/(d - 1)$. We initialise the first state to be $|0\rangle$, which only has one non-zero coefficient. We can construct the second state using the overlap and the normalisation conditions. We thus know that the $c_0^{(1)}$ has to be $-1/(d - 1)$ and $c_1^{(1)}$ is deduced from normalisation. The third state must also have $c_0^{(2)} = -1/(d - 1)$ by virtue of the overlap condition with the first state. The $c_1^{(2)}$ component of the third state is determined by the overlap condition with the second state. The last component of the third state is calculated using normalisation. One can repeat this procedure until the the full set of $d - 1$ states is constructed.

Coherent absorption

I also worked on an experiment whose results are published in Nature Communications:

T. Roger, S. Vezzoli, E. Bolduc, J. Valente, J.J.F. Heitz, J. Jeffers, C. Soci, J. Leach, C. Couteau, N.I. Zheludev, D. Faccio. *Coherent perfect absorption in deeply sub-wavelength films in the single-photon regime*. Nature Communications **6** (2015).

The goal of this work was to experimentally demonstrate the concept of coherent absorption at the single-photon level. Coherent absorption occurs when light interferes at a layer that is semi-absorbant. In order for interference to even occur, light has to hit both sides of the absorbant layer. Depending on the phase difference between the two paths, the layer can be totally absorbant or not at all in the extremes.

My role in this project was two-fold: I designed and built the first of two sections of the experiment, and I measured the second order correlation function $g^{(2)}$ of this part of the experiment. The experimental setup that I built is the same as the one from Chapter 3, except that we replace the SLMs by mirrors and use a different magnification. The experimental setup from Chapter 3 is shown in Fig. 3.1.

The second order correlation function is useful when one wants to show quantum properties of a source. For example, an ideal source of single photons would produce photons on demand or at the press of a button. To show that the photons are really produced one by one, we measure the second order correlation functions, which involves coincidence detection. If there are no coincidences at all, then the photons are solitary, and they behave in a sub-shot-noise manner. Recall that this was the subject of Chapter 5. In our case, however, we have a source of heralded photons: one photon is detected by a SPAD detector and heralds the presence of its partner. In order to show that its partner behaves in a sub-shot-noise manner, we have to measure a value of $g^{(2)} < 0.5$. In our experiment, we measured $g^{(2)} = 0.19 \pm 0.02$, which confirms the quantum behaviour that is key to the whole experiment.

Chapter 8 – Conclusions

Quantum mechanics suffers from the curse of dimensionality: quantum states grow exponentially with the number of components in a system. To push the limit of research and development in quantum information, we need powerful quantum tomography solvers.

Quantum tomography solvers come in many flavours, some of which do not require complicated post-measurement processing. This is the case of our direct approach to quantum state reconstruction. The simplicity of the method allowed us to measure and reconstruct the quantum state with a dimension of 100 000. However, this approach comes at the price of recovering an information-restricted version of the quantum state, that is the state-vector. We developed an additional quantum tomographic method, this time for which the target can be any density matrix of any purity, and the measurements can be any complete set of Hermitian operators. The key part of this technique is to let the density matrix estimate wander outside of its comfort zone, that is positive semi-definite matrices. As soon as the estimate is out, we bring it back to the closest physically-allowed density matrix. This projected gradient descent algorithm shows great promise because it is applicable to the widest range of scenarios, from well- to ill-conditioned measurements.

In this thesis, we have also seen how an EMCCD can be used to measure quantum correlations. We derived a new figure of merit to measure sub-shot-noise statistics, which signifies that the variance of the particle number distribution is lower than its mean. This type of statistics can only be achieved with quantum effects such as entanglement. We were able to show that our entangled photon source was indeed of sub-shot-noise nature.

Finally, we have laid out a new fast algorithm for image denoising in the context of low light levels. The algorithm owes its speed to the fact that it is non-iterative. We have shown that there is a tradeoff to be had between the speed of the denoising algorithms and the accuracy of the final image estimate.

A natural continuation of the work presented in this thesis, would be to apply machine learning techniques to the quantum tomography problem. The skills that one learns from digging deep into quantum tomography are transferable to other algorithm-heavy fields, such as machine learning. One question that begs for an answer is: what can the field of machine learning bring to research in quantum mechanics? It would be interesting to teach a machine to perform tomography, as perhaps this would speedup the recovery process once the task is learned. In addition to the speedup, a learned representation could perhaps best an analytical one in modelling the subtle transformations involved in a real experiment affected by interactions with the environment.

Bibliography

- [1] E. H. Kennard. *Zur quantenmechanik einfacher bewegungstypen*. Zeitschrift für Physik A Hadrons and Nuclei **44**, 326–352 (1927).
- [2] H. Weyl. *Quantenmechanik und gruppentheorie*. Zeitschrift für Physik A Hadrons and Nuclei **46**, 1–46 (1927).
- [3] A. Einstein, B. Podolsky & N. Rosen. *Can Quantum-Mechanical Description of Physical Reality Be Considered Complete?* Phys Rev A **47**, 777–780 (1935).
- [4] F. Natterer. *The mathematics of computerized tomography* (SIAM, 2001).
- [5] G. Gariepy, F. Tonolini, R. Henderson, J. Leach & D. Faccio. *Detection and tracking of moving objects hidden from view*. Nature Phot (2015).
- [6] O. Katz, E. Small & Y. Silberberg. *Looking around corners and through thin turbid layers in real time with scattered incoherent light*. Nature Phot **6**, 549–553 (2012).
- [7] Y. LeCun *et al.* *Comparison of learning algorithms for handwritten digit recognition*. In *International conference on artificial neural networks*, 53–60 (Perth, Australia, 1995).
- [8] Y. LeCun, Y. Bengio & G. Hinton. *Deep learning*. Nature Phot **521**, 436–444 (2015).
- [9] G. H. Golub, P. C. Hansen & D. P. O’Leary. *Tikhonov regularization and total least squares*. SIAM Journal on Matrix Analysis and Applications **21**, 185–194 (1999).
- [10] A. N. Tikhonov. *On the stability of inverse problems*. In *Dokl. Akad. Nauk SSSR*, 195–198 (1943).
- [11] M. Ringbauer *et al.* *Measurements on the reality of the wavefunction*. Nature Phys. **11**, 249–254 (2015).
- [12] E. J. Candès, X. Li & M. Soltanolkotabi. *Phase retrieval via Wirtinger flow: Theory and algorithms*. Information Theory **61**, 1985–2007 (2015).
- [13] Y. Chen & E. Candes. *Solving random quadratic systems of equations is nearly as easy as solving linear systems*. In *Advances in Neural Information Processing Systems*, 739–747 (2015).
- [14] A. Peres. *Separability criterion for density matrices*. Phys Rev Lett **77**, 1413 (1996).
- [15] R. Horodecki, P. Horodecki, M. Horodecki & K. Horodecki. *Quantum entanglement*. Reviews of modern physics **81**, 865 (2009).
- [16] M. Bourennane *et al.* *Experimental detection of multipartite entanglement using witness operators*. Phys Rev Lett **92**, 087902 (2004).

- [17] M. Agnew, J. Leach & R. W. Boyd. *Observation of entanglement witnesses for orbital angular momentum states*. The European Physical Journal D-Atomic, Molecular, Optical and Plasma Physics **66**, 1–4 (2012).
- [18] J. S. Bell. *On the einstein podolsky rosen paradox*. Physics 195–200 (1964).
- [19] L. Allen, M. W. Beijersbergen, R. J. C. Spreeuw & J. P. Woerdman. *Orbital angular momentum of light and the transformation of Laguerre-Gaussian laser modes*. Phys Rev A **A45**, 8185–8189 (1992).
- [20] E. Bolduc, N. Bent, E. Santamato, E. Karimi & R. W. Boyd. *Exact solution to simultaneous intensity and phase encryption with a single phase-only hologram*. Opt. Lett. **38**, 3546–3549 (2013).
- [21] J. Leach, M. J. Padgett, S. M. Barnett, S. Franke-Arnold & J. Courtial. *Measuring the orbital angular momentum of a single photon*. Phys Rev Lett **88**, 257901 (2002).
- [22] M. P. Lavery *et al.* *Efficient measurement of an optical orbital-angular-momentum spectrum comprising more than 50 states*. New J. Phys. **15**, 013024 (2013).
- [23] W. Pauli. *General Principles of Quantum Mechanics* (Springer Berlin Heidelberg, Berlin, Heidelberg, 1980).
- [24] U. Fano. *Description of states in quantum mechanics by density matrix and operator techniques*. Reviews of modern physics **29**, 74 (1957).
- [25] S. Lloyd, M. Mohseni & P. Rebentrost. *Quantum principal component analysis*. Nature Phys. **10**, 631–633 (2014).
- [26] W. K. Wootters & B. D. Fields. *Optimal state-determination by mutually unbiased measurements*. Annals of Physics **191**, 363–381 (1989).
- [27] D. A. B. Miller. *Self-configuring universal linear optical component [Invited]*. Photon. Res. **1**, 1 (2013).
- [28] K. Ried *et al.* *A quantum advantage for inferring causal structure*. Nature Phys. **11**, 414 (2015).
- [29] K. Banaszek, G. M. D’Ariano, M. G. A. Paris & M. F. Sacchi. *Maximum-likelihood estimation of the density matrix*. Phys Rev A **61**, 10304 (2000).
- [30] D. F. James, P. G. Kwiat, W. J. Munro & A. G. White. *Measurement of qubits*. Phys Rev A **64**, 052312 (2001).
- [31] J. Shang, Z. Zhang & H. K. Ng. *Superfast maximum likelihood reconstruction for quantum tomography*. arXiv preprint arXiv:1609.07881 (2016).
- [32] J. Řeháček, Z. Hradil & M. Ježek. *Iterative algorithm for reconstruction of entangled states*. Phys Rev A **63**, 040303 (2001).
- [33] J. Řeháček, Z. Hradil, E. Knill & A. I. Lvovsky. *Diluted maximum-likelihood algorithm for quantum tomography*. Phys Rev A **75**, 042108 (2007).
- [34] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker & J. Eisert. *Quantum State Tomography via Compressed Sensing*. Phys Rev Lett **105**, 150401 (2010).

- [35] M. Grant, S. Boyd & Y. Ye. *CVX: Matlab software for disciplined convex programming* (2008).
- [36] E. J. Candès, X. Li & M. Soltanolkotabi. *Phase retrieval via Wirtinger flow: Theory and algorithms*. *Information Theory* **61**, 1985–2007 (2015).
- [37] S. Becker, V. Cevher & A. Kyrillidis. *Randomized singular value projection*. *Sampling Theory and Approximation (SampTA)* (2013).
- [38] E. Bolduc, G. Gariépy & J. Leach. *Direct measurement of large-scale quantum states via expectation values of non-Hermitian matrices*. *Nature Phot* **7**, 10439 (2016).
- [39] D. Collins, N. Gisin, N. Linden, S. Massar & S. Popescu. *Bell Inequalities for Arbitrarily High-Dimensional Systems*. *Phys Rev Lett* **88**, 040404 (2002).
- [40] A. C. Dada, J. Leach, G. S. Buller, M. J. Padgett & E. Andersson. *Experimental high-dimensional two-photon entanglement and violations of generalized Bell inequalities*. *Nature Phys.* **7**, 1–4 (2011).
- [41] A. C. Dada & E. Andersson. *ON BELL INEQUALITY VIOLATIONS WITH HIGH-DIMENSIONAL SYSTEMS*. *Int. J. Quantum Inform.* **9**, 1807–1823 (2011).
- [42] T. Monz *et al.* *14-Qubit Entanglement: Creation and Coherence*. *Phys Rev Lett* **106**, 130506 (2011).
- [43] X. C. Yao *et al.* *Observation of eight-photon entanglement*. *Nature Phot* **6**, 225–228 (2012).
- [44] S. Yokoyama *et al.* *Ultra-large-scale continuous-variable cluster states multiplexed in the time domain*. *Nature Phot* **7**, 982–986 (2013).
- [45] M. Krenn *et al.* *Generation and confirmation of a (100x100)-dimensional entangled quantum system*. *Proc. Natl Acad. Sci.* **111**, 6243–6247 (2014).
- [46] G. Tóth *et al.* *Permutationally Invariant Quantum Tomography*. *Phys Rev Lett* **105**, 250403 (2010).
- [47] M. Cramer *et al.* *Efficient quantum state tomography*. *Nature Phot* **1**, 149–147 (2010).
- [48] S. T. Flammia, A. Silberfarb & C. M. Caves. *Minimal Informationally Complete Measurements for Pure States*. *Found Phys* **35**, 1985–2006 (2005).
- [49] J. Finkelstein. *Pure-state informationally complete and “really” complete measurements*. *Phys Rev A* **70**, 052107 (2004).
- [50] J. S. Lundeen, B. Sutherland, A. Patel, C. Stewart & C. Bamber. *Direct measurement of the quantum wavefunction*. *Nature* **474**, 188–191 (2012).
- [51] T. Heinosaari, L. Mazzarella & M. M. Wolf. *Quantum Tomography under Prior Information*. *Commun. Math. Phys.* **318**, 355–374 (2013).
- [52] J. Z. Salvail, M. Agnew, A. S. Johnson, E. Bolduc & J. Leach. *Full characterization of polarization states of light via direct measurement*. *Nature Phot* (2013).

- [53] A. Shabani *et al.* *Efficient Measurement of Quantum Dynamics via Compressive Sensing*. Phys Rev Lett **106**, 100401 (2011).
- [54] D. Mogilevtsev *et al.* *Data pattern tomography: reconstruction with an unknown apparatus*. New J. Phys. **15**, 025038 (2013).
- [55] R. Storn & K. Price. *Differential evolution—a simple and efficient heuristic for global optimization over continuous spaces*. J. Global Optim. **11**, 341–359 (1997).
- [56] F. M. Miatto, A. M. Yao & S. M. Barnett. *Full characterization of the quantum spiral bandwidth of entangled biphotons*. Phys Rev A **83**, 033816 (2011).
- [57] M. Agnew, J. Leach, M. McLaren, F. S. Roux & R. W. Boyd. *Tomography of the quantum state of photons entangled in high dimensions*. Phys Rev A **84**, 062101 (2011).
- [58] J. Leach, E. Bolduc, D. J. Gauthier & R. W. Boyd. *Secure information capacity of photons entangled in many dimensions*. Phys Rev A **85**, 060304 (2012).
- [59] V. D. Salakhutdinov, E. R. Eliel & W. Löffler. *Full-Field Quantum Correlations of Spatially Entangled Photons*. Phys Rev Lett **108**, 173604 (2012).
- [60] D. S. Tasca *et al.* *Imaging high-dimensional spatial entanglement with a camera*. Nat Commun **3**, 984–6 (2012).
- [61] D. Geelen & W. Löffler. *Walsh modes and radial quantum correlations of spatially entangled photons*. Opt. Lett. **38**, 4108–4111 (2013).
- [62] D. Giovannini *et al.* *Characterization of high-dimensional entangled systems via mutually unbiased measurements*. Phys Rev Lett **110**, 143601 (2013).
- [63] J. Emerson, Y. S. Weinstein, M. Saraceno, S. Lloyd & D. G. Cory. *Pseudo-random unitary operators for quantum information processing*. Science **302**, 2098–2100 (2003).
- [64] A. Ekert & P. L. Knight. *Entangled quantum systems and the Schmidt decomposition*. Am. J. Phys. **63** (1995).
- [65] C. K. Law & J. H. Eberly. *Analysis and interpretation of high transverse entanglement in optical parametric down conversion*. Phys Rev Lett **92**, 127903 (2004).
- [66] J. S. Bell. *On the einstein-podolsky-rosen paradox*. Physics **1**, 195–200 (1964).
- [67] A. K. Ekert. *Quantum Cryptography Based on Bell’s Theorem*. Phys Rev Lett **67**, 1–3 (2011).
- [68] C. I. Osorio, A. Valencia & J. P. Torres. *Spatiotemporal correlations in entangled photons generated by spontaneous parametric down conversion*. New J. Phys. **10**, 113012 (2008).
- [69] C. I. Osorio. *Spatial Characterization Of Two-Photon States*. Ph.D. thesis, Barcelona (2014).
- [70] K. Bartkiewicz, K. Lemr & A. Miranowicz. *Direct method for measuring of purity, superfidelity, and subfidelity of photonic two-qubit mixed states*. Phys Rev A **88**, 052104 (2013).

- [71] M. Malik *et al.* *Direct measurement of a 27-dimensional orbital-angular-momentum state vector.* *Nature Phot* **5**, 3115 (2014).
- [72] Y. Aharonov, D. Z. Albert & L. Vaidman. *How the result of a measurement of a component of the spin of a spin-1/2 particle can turn out to be 100.* *Phys Rev Lett* **60**, 1351 (1988).
- [73] D. S. Gonçalves, M. A. Gomes-Ruggiero & C. Lavor. *A projected gradient method for optimization over density matrices.* *Optimization Methods and Software* **31**, 328–341 (2016).
- [74] W.-B. Gao *et al.* *Experimental demonstration of a hyper-entangled ten-qubit Schrödinger cat state.* *Nature Phys.* **6**, 331–335 (2010).
- [75] P. Schindler *et al.* *Experimental repetitive quantum error correction.* *Science* **332**, 1059–1061 (2011).
- [76] X.-C. Yao *et al.* *Observation of eight-photon entanglement.* *Nature Phot* **6**, 225–228 (2012).
- [77] G. B. Silva, S. Glancy & H. M. Vasconcelos. *Investigating Bias in Maximum Likelihood Quantum State Tomography.* arXiv preprint arXiv:1604.00321 (2016).
- [78] K.-C. Toh, M. J. Todd & R. H. Tütüncü. *SDPT3—a MATLAB software package for semidefinite programming, version 1.3.* *Optimization Methods and Software* **11**, 545–581 (1999).
- [79] J. F. Sturm. *Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones.* *Optimization Methods and Software* **11**, 625–653 (1999).
- [80] J. A. Smolin, J. M. Gambetta & G. Smith. *Efficient method for computing the maximum-likelihood quantum state from measurements with additive gaussian noise.* *Phys Rev Lett* **108**, 070502 (2012).
- [81] Z. H. Guo *et al.* *Full reconstruction of a 14-qubit state within four hours.* *New J. Phys.* **18**, 083036 (2016).
- [82] A. Beck & M. Teboulle. *A fast iterative shrinkage-thresholding algorithm for linear inverse problems.* *SIAM J Imaging Sci* **2**, 183–202 (2009).
- [83] K. Życzkowski, K. A. Penson, I. Nechita & B. Collins. *Generating random density matrices.* *Journal of Mathematical Physics* **52**, 062201 (2011).
- [84] A. Miranowicz *et al.* *Optimal two-qubit tomography based on local and global measurements: Maximal robustness against errors as described by condition numbers.* *Phys Rev A* **90**, 062123 (2014).
- [85] A. Feito *et al.* *Measuring measurement: theory and practice.* *New J. Phys.* **11**, 093038 (2009).
- [86] J. S. Lundeen *et al.* *Tomography of quantum detectors.* *Nature Phys.* **5**, 27–30 (2009).
- [87] R. Bianchetti *et al.* *Control and tomography of a three level superconducting artificial atom.* *Phys Rev Lett* **105**, 223601 (2010).

- [88] M. S. Kaznady & D. F. James. *Numerical strategies for quantum tomography: Alternatives to full optimization*. Phys Rev A **79**, 022109 (2009).
- [89] W. H. Press, S. A. Teukolsky, W. T. Vetterling & B. P. Flannery. *Numerical recipes in C*, vol. 2 (Cambridge university press Cambridge, 1996).
- [90] C. Michelot. *A finite algorithm for finding the projection of a point onto the canonical simplex of ∞ n*. Journal of Optimization Theory and Applications **50**, 195–200 (1986).
- [91] I. Sutskever, J. Martens, G. Dahl & G. Hinton. *On the importance of initialization and momentum in deep learning*. In *Proceedings of the 30th international conference on machine learning*, 1139–1147 (2013).
- [92] D. S. Gonçalves, M. A. Gomes-Ruggiero & C. Lavor. *Global convergence of diluted iterations in maximum-likelihood quantum tomography*. Quantum Information & Computation **14**, 966–980 (2014).
- [93] S. Boyd & L. Vandenberghe. *Convex optimization* (Cambridge university press, 2004).
- [94] S. Becker, V. Cevher, C. Koch & A. Kyrillidis. *Sparse projections onto the simplex*. ICML, to appear (2013).
- [95] T. Heinosaari, L. Mazzarella & M. M. Wolf. *Quantum tomography under prior information*. Commun. Math. Phys. **318**, 355–374 (2013).
- [96] G. Gour & A. Kalev. *Construction of all general symmetric informationally complete measurements*. Journal of Physics A: Mathematical and Theoretical **47**, 335302 (2014).
- [97] M. P. Edgar *et al.* *Imaging high-dimensional spatial entanglement with a camera*. Nature Phot **3**, 984 (2012).
- [98] E. Bolduc, D. Faccio & J. Leach. *Acquisition of multiple photon pairs with an EMCCD camera*. Journal of Optics **19**, 054006 (2017).
- [99] G. Brida, M. Genovese & I. R. Berchera. *Experimental realization of sub-shot-noise quantum imaging*. Nature Phot **4**, 227–230 (2010).
- [100] M. A. Taylor *et al.* *Biological measurement beyond the quantum limit*. Nature Phot **7**, 229–233 (2013).
- [101] R. Fickler, M. Krenn, R. Lapkiewicz, S. Ramelow & A. Zeilinger. *Real-time imaging of quantum entanglement*. Sci. Rep. **3** (2013).
- [102] B. Jost, A. Sergienko, A. Abouraddy, B. Saleh & M. Teich. *Spatial correlations of spontaneously down-converted photon pairs detected with a single-photon-sensitive CCD camera*. Optics express **3**, 81–88 (1998).
- [103] S. Oemrawsingh, W. J. Van Drunen, E. R. Eliel & J. P. Woerdman. *Two-dimensional wave-vector correlations in spontaneous parametric downconversion explored with an intensified CCD camera*. JOSA B **19**, 2391–2395 (2002).

- [104] E. Brambilla, A. Gatti, M. Bache & L. A. Lugiato. *Simultaneous near-field and far-field spatial quantum correlations in the high-gain regime of parametric down-conversion*. Phys Rev A **69**, 023802 (2004).
- [105] O. Jedrkiewicz *et al.* *Detection of sub-shot-noise spatial correlation in high-gain parametric down conversion*. Phys Rev Lett **93**, 243601 (2004).
- [106] L. A. Lugiato, A. Gatti & E. Brambilla. *Quantum imaging*. J. Opt. B: Quantum Semiclass. Opt **4** (2002).
- [107] A. Mosset, F. Devaux, G. Fanjoux & E. Lantz. *Direct experimental characterization of the Bose-Einstein distribution of spatial fluctuations of spontaneous parametric down-conversion*. The European Physical Journal D-Atomic, Molecular, Optical and Plasma Physics **28**, 447–451 (2004).
- [108] E. Brambilla, L. Caspani, O. Jedrkiewicz, L. A. Lugiato & A. Gatti. *High-sensitivity imaging with multi-mode twin beams*. Phys Rev A **77**, 053807 (2008).
- [109] P.-A. Moreau, J. Mougins-Sisini, F. Devaux & E. Lantz. *Realization of the purely spatial Einstein-Podolsky-Rosen paradox in full-field images of spontaneous parametric down-conversion*. Phys Rev A **86**, 010101 (2012).
- [110] M. Dabrowski, R. Chrapkiewicz & W. Wasilewski. *Spatially Multimode Holographic Quantum Memory for Single and Multiple Photons Generation*. In *CLEO: QELS Fundamental Science*, FM2E-1 (Optical Society of America, 2017).
- [111] S. M. Barnett & D. T. Pegg. *On the Hermitian optical phase operator*. J Mod Opt **36**, 7–19 (1989).
- [112] C. M. Caves. *Quantum-mechanical noise in an interferometer*. Physical Review D **23**, 1693 (1981).
- [113] J. Gea-Banacloche & G. Leuchs. *Squeezed states for interferometric gravitational-wave detectors*. J Mod Opt **34**, 793–811 (1987).
- [114] K. McKenzie, D. A. Shaddock, D. E. McClelland, B. C. Buchler & P. K. Lam. *Experimental demonstration of a squeezing-enhanced power-recycled Michelson interferometer for gravitational wave detection*. Phys Rev Lett **88**, 231102 (2002).
- [115] K. Goda *et al.* *A quantum-enhanced prototype gravitational-wave detector*. Nature Phys. **4**, 472–476 (2008).
- [116] L. S. Collaboration & others. *A gravitational wave observatory operating beyond the quantum shot-noise limit*. Nature Phys. **7**, 962–965 (2011).
- [117] B. P. Abbott *et al.* *Observation of gravitational waves from a binary black hole merger*. Phys Rev Lett **116**, 061102 (2016).
- [118] M. I. Kolobov & P. Kumar. *Sub-shot-noise microscopy: imaging of faint phase objects with squeezed light*. Opt. Lett. **18**, 849–851 (1993).
- [119] P. H. Ribeiro, C. Schwob, A. Maître & C. Fabre. *Sub-shot-noise high-sensitivity spectroscopy with optical parametric oscillator twin beams*. Opt. Lett. **22**, 1893–1895 (1997).

- [120] J.-L. Blanchet, F. Devaux, L. Furfaro, E. Lantz & others. *Measurement of sub-shot-noise correlations of spatial fluctuations in the photon-counting regime*. Phys Rev Lett **101**, 233604 (2008).
- [121] J. Hyneczek & T. Nishiwaki. *Excess noise and other important characteristics of low light level imaging using charge multiplying CCDs*. Electron Devices, IEEE Transactions on **50**, 239–245 (2003).
- [122] J. Sperling, W. Vogel & G. S. Agarwal. *Sub-binomial light*. Phys Rev Lett **109**, 093601 (2012).
- [123] J. Leach, E. Bolduc, D. J. Gauthier & R. W. Boyd. *Secure information capacity of photons entangled in many dimensions*. Phys Rev A **85**, 060304 (2012).
- [124] D. Rosenberg, A. E. Lita, A. J. Miller & S. W. Nam. *Noise-free high-efficiency photon-number-resolving detectors*. Phys Rev A **71**, 061803 (2005).
- [125] D. I. Schuster *et al.* *Resolving photon number states in a superconducting circuit*. Nature Phot **445**, 515–518 (2007).
- [126] D. Achilles *et al.* *Photon-number-resolving detection using time-multiplexing*. J Mod Opt **51**, 1499–1515 (2004).
- [127] F. E. Becerra, J. Fan & A. Migdall. *Photon number resolution enables quantum receiver for realistic coherent optical communications*. Nature Phot **9**, 48–53 (2015).
- [128] T. Iskhakov, I. N. Agafonov, M. V. Chekhova & G. Leuchs. *Polarization-entangled light pulses of 100'000 photons*. Phys Rev Lett **109**, 150502 (2012).
- [129] P. A. Morris, R. S. Aspden, J. E. Bell, R. W. Boyd & M. J. Padgett. *Imaging with a small number of photons*. Nature Phot **6** (2015).
- [130] G. Brida, M. Genovese & I. R. Berchera. *Experimental realization of sub-shot-noise quantum imaging*. Nature Phot **4**, 227–230 (2010).
- [131] N. A. Dutton, L. Parmesan, A. J. Holmes, L. A. Grant & R. K. Henderson. *320×240 oversampled digital single photon counting image sensor*. In *VLSI Circuits Digest of Technical Papers, 2014 Symposium on*, 1–2 (IEEE, 2014).
- [132] F.-X. Dupé, J. M. Fadili & J.-L. Starck. *A proximal iteration for deconvolving Poisson noisy images using sparse representations*. IEEE Trans. Image Process. **18**, 310–321 (2009).
- [133] K. Hirakawa & P. J. Wolfe. *Efficient multivariate Skellam Shrinkage for denoising photon-limited image data: An empirical bayes approach*. In *Image Processing (ICIP), 2009 16th IEEE International Conference on*, 2961–2964 (IEEE, 2009).
- [134] S. Lefkimmiatis, P. Maragos & G. Papandreou. *Bayesian inference on multiscale models for poisson intensity estimation: Applications to photon-limited image denoising*. IEEE Trans. Image Process. **18**, 1724–1741 (2009).
- [135] S. Setzer, G. Steidl & T. Teuber. *Deblurring Poissonian images by split Bregman techniques*. J Vis Commun Image Represent **21**, 193–199 (2010).

- [136] M. A. Figueiredo & J. M. Bioucas-Dias. *Restoration of Poissonian images using alternating direction optimization*. IEEE Trans. Image Process. **19**, 3133–3145 (2010).
- [137] S. Ono & I. Yamada. *Poisson image restoration with likelihood constraint via hybrid steepest descent method*. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, 5929–5933 (IEEE, 2013).
- [138] J. Salmon, Z. Harmany, C.-A. Deledalle & R. Willett. *Poisson noise reduction with non-local PCA*. J. Math. Imaging Vision **48**, 279–294 (2014).
- [139] R. Giryes & M. Elad. *Sparsity-based Poisson denoising with dictionary learning*. IEEE Trans. Image Process. **23**, 5057–5069 (2014).
- [140] M. Sonnleitner, J. Jeffers & S. M. Barnett. *Image retrodiction at low light levels*. Optica **2**, 950–957 (2015).
- [141] L. B. Lucy. *An iterative technique for the rectification of observed distributions*. The astronomical journal **79**, 745 (1974).
- [142] J. Schmitt, J. L. Starck, J. M. Casandjian, J. Fadili & I. Grenier. *Poisson denoising on the sphere: application to the Fermi gamma ray space telescope*. Astronomy & Astrophysics **517**, A26 (2010).
- [143] B. M. Onat *et al.* *Ultra-low dark current InGaAs technology for focal plane arrays for low-light level visible-shortwave infrared imaging*. In *Defense and Security Symposium*, 65420L–65420L (International Society for Optics and Photonics, 2007).
- [144] J. Leipsic, G. Nguyen, J. Brown, D. Sin & J. R. Mayo. *A prospective evaluation of dose reduction and image quality in chest CT using adaptive statistical iterative reconstruction*. American Journal of Roentgenology **195**, 1095–1099 (2010).
- [145] M. Elad & M. Aharon. *Image denoising via sparse and redundant representations over learned dictionaries*. IEEE Trans. Image Process. **15**, 3736–3745 (2006).
- [146] L. I. Rudin, S. Osher & E. Fatemi. *Nonlinear total variation based noise removal algorithms*. Physica D: Nonlinear Phenomena **60**, 259–268 (1992).
- [147] C. R. Vogel & M. E. Oman. *Iterative methods for total variation denoising*. SIAM J. Sci. Comput **17**, 227–238 (1996).
- [148] A. Chambolle. *An algorithm for total variation minimization and applications*. J. Math. Imaging Vision **20**, 89–97 (2004).
- [149] P. J. Huber. *Robust estimation of a location parameter*. The Annals of Mathematical Statistics **35**, 73–101 (1964).
- [150] Y. Nesterov. *A method of solving a convex programming problem with convergence rate $O(1/k^2)$* . In *Soviet Mathematics Doklady*, 372–376 (1983).
- [151] Y. Nesterov. *Introductory Lectures on Convex Optimization*, vol. 87 of *Applied Optimization* (Springer US, Boston, MA, 2004).
- [152] Y. Nesterov. *Gradient methods for minimizing composite objective function*. CORE Discussion Papers **76** (2007).

- [153] T. Goldstein & S. Osher. *The split Bregman method for L1-regularized problems*. SIAM J Imaging Sci **2**, 323–343 (2009).
- [154] K. Hirakawa & P. J. Wolfe. *Skellam shrinkage: Wavelet-based intensity estimation for inhomogeneous poisson data*. Information Theory **58**, 1080–1093 (2012).
- [155] A. Chambolle, R. A. De Vore, N.-Y. Lee & B. J. Lucier. *Nonlinear wavelet image processing: variational problems, compression, and noise removal through wavelet shrinkage*. IEEE Trans. Image Process. **7**, 319–335 (1998).
- [156] J. G. Skellam. *The frequency distribution of the difference between two Poisson variates belonging to different populations*. Journal of the Royal Statistical Society. Series A (General) **109**, 296 (1946).
- [157] D. E. Amos. *Computation of modified Bessel functions and their ratios*. Mathematics of Computation **28**, 239–251 (1974).
- [158] F. Luisier, C. Vonesch, T. Blu & M. Unser. *Fast interscale wavelet denoising of Poisson-corrupted images*. Signal Processing **90**, 415–427 (2010).