UNIVERSIDADE FEDERAL DE MINAS GERAIS
Instituto de Ciência Exatas
Programa de Pós-graduação em Física

Mário Foganholi Fernandes

Quantum Ptychography

BELO HORIZONTE
2022

Mário Foganholi Fernandes

Quantum Ptychography

Tese apresentada ao Programa de Pós-Graduação em Física do Instituto de Ciências Exatas da Universidade Federal de Minas Gerais como requisito parcial para obtenção do título de Doutor em Ciências.

Orientador: Leonardo Teixeira Neves

Coorientador: Reinaldo Oliveira Vianna

Belo Horizonte

2022

UNIVERSIDADE FEDERAL DE MINAS GERAIS
INSTITUTO DE CIÊNCIAS EXATAS
PROGRAMA DE PÓS-GRADUAÇÃO EM FÍSICA

ATA DE DEFESA DE TESE

ATA DA SESSÃO DE ARGUIÇÃO DA 394ª TESE DO PROGRAMA DE PÓS-GRADUAÇÃO EM FÍSICA, DEFENDIDA POR MÁRIO FOGANHOLI FERNANDES orientado pelo professor Leonardo Teixeira Neves e coorientador pelo professor Reinaldo Oliveira Vianna, para obtenção do grau de **DOUTOR EM CIÊNCIAS, área de concentração física**. Às 09:00 horas de dezoito de fevereiro de dois mil e vinte e dois reuniu-se, por videoconferência, a Comissão Examinadora, composta pelos professores **Leonardo Teixeira Neves** (Orientador - Departamento de Física/UFMG), **Reinaldo Oliveira Vianna** (Coorientador - Departamento de Física/UFMG), **Sebastião Nascimento de Pádua** (Departamento de Física/UFMG), **Raphael Campos Drumond** (Departamento de Matemática/UFMG), **Miguel Angel Solis Prosser** (Universidad de la Frontera) e **Dardo Miguel Goyeneche** (Universidad de Antofagasta) para dar cumprimento ao Artigo 37 do Regimento Geral da UFMG, submetendo o Mestre **MÁRIO FOGANHOLI FERNANDES** à arguição de seu trabalho de Tese de Doutorado, que recebeu o título de **"Quantum Ptychography"**. O candidato fez uma exposição oral de seu trabalho durante aproximadamente 50 minutos. Após esta, os membros da comissão prosseguiram com a sua arguição, e apresentaram seus pareceres individuais sobre o trabalho, concluindo pela aprovação do candidato.

Belo Horizonte, 18 de fevereiro de 2022.

Prof. Leonardo Teixeira Neves
Orientador do estudante
Departamento de Física/UFMG

Prof. Raphael Campos Drumond
Departamento de Matemática/UFMG

Prof. Reinaldo Oliveira Vianna
Coorientador do estudante
Departamento de Física/UFMG

Prof. Miguel Angel Solis Prosser
Universidad de la Frontera

Prof. Sebastião Nascimento de Pádua
Departamento de Física/UFMG

Prof. Dardo Miguel Goyeneche
Universidad de Antofagasta

Candidato: Mário Foganholi Fernandes

Abstract

We have proposed a new method for quantum state reconstruction, called *quantum ptychography*. This method is based on an analogy with ptychography, a method for phase retrieval in classical optics.

The main advantage of quantum ptychography, when compared to other methods for quantum state reconstruction, is its experimental simplicity: it uses measurements in a single basis, while using intermediate projectors to achieve data diversity; on the other hand, usual methods require measurements in several different bases. By not using many bases, quantum ptychography allows simpler experimental setups. This greater simplicity might be especially invaluable in a context of material scarcity, such as brazilian science has been facing in the recent years. Quantum ptychography is also flexible regarding the intermediate projectors, which allows for the experimenter to customize it to their needs.

We have studied the method extensively in the context of pure state reconstruction. We have shown its equivalence to an alternating gradient search, allowing a mathematical understanding of how the method works. We have made numerical simulations of its reconstructions, using realistic amounts of noise, and showed that its reconstructions are excellent. We have also made an experiment with slit qudits to test its performance, and again attained very satisfactory results.

As it was with classical ptychography, there are many improvements that quantum ptychography could see. We have also generalized it for mixed states and obtained encouraging preliminary results. With this and other advances, we hope that quantum ptychography can become a useful tool for quantum information.

Keywords: quantum state reconstruction, ptychography, phase retrieval.

Resumo

Nós propusemos um novo método para reconstruir estados quânticos, chamado *pticografia quântica*. Esse método se baseia em uma analogia com a pticografia, um método de reconstrução de fase na óptica clássica.

A principal vantagem da pticografia quântica, quando comparada a outros métodos de reconstrução de estado, é sua simplicidade experimental: ela usa medidas em uma única base, ao passo que projetores intermediários são usados para se conseguir a diversidade de dados necessária; por outro lado, os métodos usuais requerem medidas em diversas bases diferentes. Ao não usar diversas bases, a pticografia quântica permite que as montagens experimentais sejam mais simples. Essa maior simplicidade pode ser especialmente valiosa no contexto de escassez material, tal qual a ciência brasileira vem enfrentando nos últimos anos. A pticografia quântica também é flexível com respeito aos projetores intermediários, o que permite ao experimentador ajustá-la às suas necessidades.

Nós estudamos esse método a fundo no contexto da reconstrução de estados puros. Nós mostramos a sua equivalência a uma busca por gradientes alternados, o que permite um entendimento matemático do seu funcionamento. Fizemos simulações numéricas de suas reconstruções, usando níveis realistas de ruído, e mostramos que suas reconstruções são excelentes. Também fizemos um experimento com qudits de fenda para averiguar a sua performance, e novamente obtivemos resultados muito satisfatórios.

Assim como a pticografia clássica, a pticografia quântica é passível de diversos avanços. Nós também a generalizamos para estados mistos e obtivemos resultados preliminares encorajadores. Com esse e outros desenvolvimentos, nós esperamos que a pticografia quântica possa se tornar uma ferramenta útil na informação quântica.

Palavras-chave: reconstrução de estados quânticos, pticografia, reconstrução de fase.

Contents

Chapter 1

Introduction and Motivation

The purpose of this thesis is to explain a method we proposed to solve the task of quantum state reconstruction. It has an interesting analogy with ptychography, a method devised to reconstruct wavefronts. But before diving into the details, we would like to motivate the task that it performs. And, to that end, nothing could be better than a short story involving quantum computers and the fate of humankind.

1.1 Colonizing exoplanets with the aid of quantum computers: a motivational story

Imagine that, in a near future, humanity is planning to populate a planet that happens to be about the size of Mars. Because of its gravity and temperature, its atmosphere is mainly constituted by carbon dioxide. Consequently, the colonies will need to use plants to generate oxygen and make breathable air. However, its atmospheric conditions and orbiting-star spectrum result in very little blue and red light reaching its surface: only green light is appreciably present. Consequently, photosynthesis would be hindered to the point where sustaining the colony would be impossible.

A plan is set to research molecules similar to chlorophyll, but optimized to harbor energy from the light at the planet's surface. The colonies could then use genetically-modified plants with that molecule to sustain themselves. Since simulating big molecules is a hard task for classical computers, a quantum computer is set to simulate the behavior of candidate molecules. It simulates the dynamics of each molecule when illuminated with the spectral composition found in the planet. By the end of each simulation, its

quantum memory holds the final state of the molecules, and enables the assessment of its photosynthetic efficiency. In order to read that information, however, it is necessary to reconstruct the state of its quantum memory.

By carrying these numerical simulations, several interesting molecules are found. Further studies and experiments finally arrive at an ideal molecule, and genetically-modified plants are prepared and sent together with the human crew to the colony. The stories of the settlers, their discoveries and hurdles are subject for another day, though.

This story illustrates how the task of reconstructing quantum states is fundamental for quantum technologies. Quantum computers can help us solve problems that ordinary computers cannot. But we need to read their results by the end of each computation, after all.

Because it is a costly task, there are quantum algorithms that are designed to avoid state reconstruction by encoding the solutions into the output basis. This is only feasible in problems where the output is discrete, though, and other problems, such as quantum simulation, cannot benefit from this trick[1]. Furthermore, irrespective of the problem, during the setup of a quantum computer, state reconstruction is a valuable tool to check if each part is performing as it should.

The costly aspect of quantum state reconstruction comes from the fact that it requires the probability distribution of several observables, and for that we need to accumulate the results of many different measurements. We must therefore prepare, evolve and measure the system several times. This is not strictly forbidding though, and some systems can carry these processes quite fast (for example, Google's Sycamore quantum processor can sample a 53-qubit circuit 5000 times per second [1]). Most systems are not as fast, however, and reconstructing states remains a significant burden.

As could be expected, being such a fundamental problem, quantum state reconstruction has already been solved. In fact, there are several methods that have been proposed. Usually, the approach is to devise an informationally-complete set of operators, and then implement them in the laboratory. We feel that this method, though natural from the logical point of view, might not always arrive at the simplest operators, from the experimental point of view. There is, therefore, opportunity for improvement, and it is in this sense that our work brings a contribution. The method we devised has the advantage of demanding simpler experimental setups, as we will see.

---

[1]If the quantum system is too large, reconstructing it can become a hopeless task. This is only feasible in smaller computers.

1.2   Thesis overview

This thesis is organized in the following manner:

- **chapter 2** will review the quantum state reconstruction problem, along with methods which have already been proposed in the literature to solve it;

- **chapter 3** will cover the analogy between the reconstructions of quantum states and classical wavefronts, and also present two methods to solve the latter, which served as starting points to our work;

- **chapter 4** presents the method we proposed for the pure state case, its key results and compares it to other methods;

- **chapter 5** shows the experimental results we obtained for our method;

- **chapter 6** discusses future perspectives for the method we proposed; we highlight the preliminary results we obtained regarding its generalization to mixed quantum states;

- and finally, **chapter 7** presents the conclusions of this thesis.

1.3   Prerequisites

We assume the reader to be familiar with the subjects of quantum mechanics and, to a lesser extent, classical wave optics. Quantum states (pure and mixed) and observables will be pervasive throughout the thesis, and we point the reader to reference [2] in case they are not familiar with them. Regarding classical optics, we recommend at least some degree of familiarity with Fourier optics and diffraction, and leave reference [3] as suggestion for those who need to get acquainted with this subject. Finally, the reader should also be familiar with gradient descent optimization, the Fourier transform, and elementary linear algebra, which can be found in many standard calculus textbooks.

Chapter 2

Quantum state reconstruction

Before we delve into the quantum analogue of ptychography that we proposed, we need to formulate quantum state reconstruction more precisely. This will be done in section 2.1. Then we shall briefly review some methods that have been proposed in the literature in section 2.2. Subsection 2.2.1 will focus in methods for general/mixed states, while subsection 2.2.2 will focus in pure states.

2.1   Problem formulation

Reconstructing a quantum state is a challenge that we have to tackle indirectly. This is because detectors cannot measure phases, only amplitudes. The process of reconstruction therefore demands several observables to be measured, in a process that might resemble the inference of the shape of an object from the shadows it casts from different angles of illumination (as illustrated in figure 2.1).

Just as different angles of illumination can make shadows that capture distinct features of an object, different observables can help us capture different features of a quantum state. For a general state, $\hat{\rho}$, the expected value $\langle \hat{O} \rangle$ of an observable $\hat{O}$ is

$$\langle \hat{O} \rangle = \mathrm{Tr}\left(\hat{O}\hat{\rho} \right) \quad . \tag{2.1}$$

The quantum reconstruction problem can hence be posed as following: given a set of known observables $\{\hat{O}_k\}_{k \in K}$ (where K is just a set of indexes), and their expected values
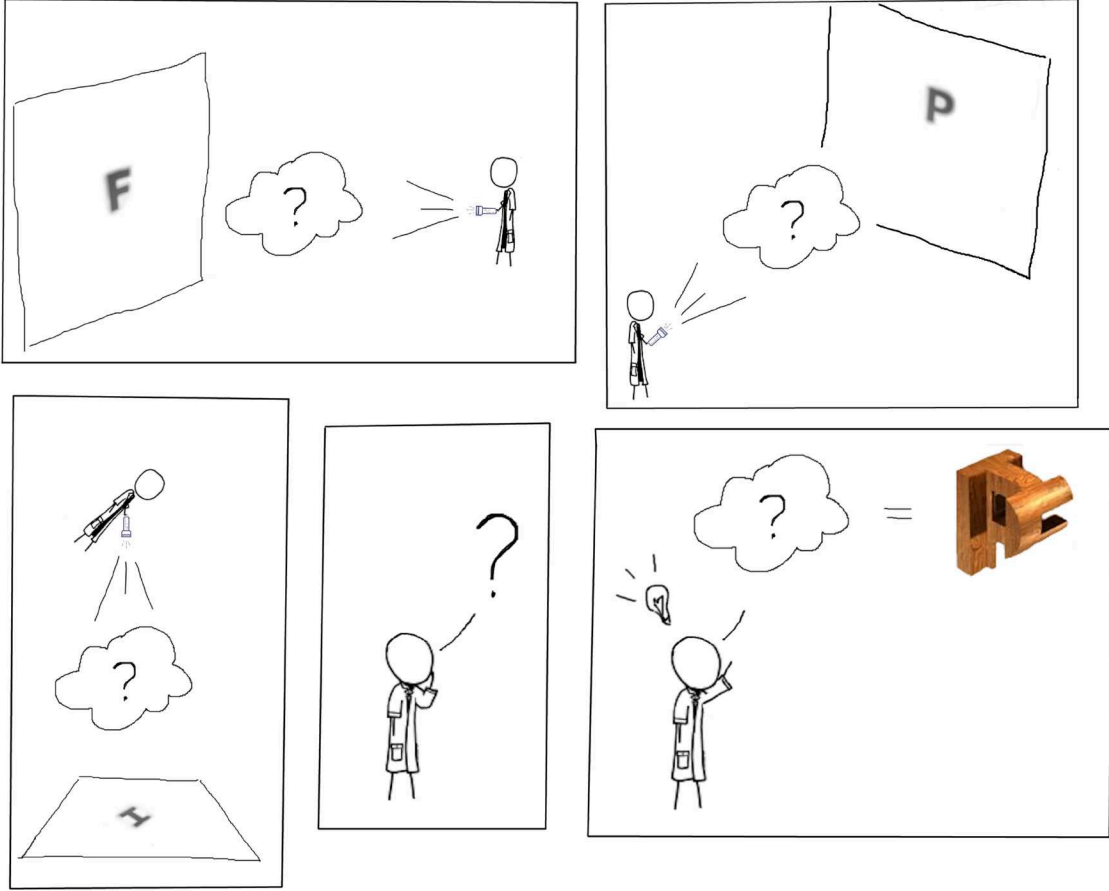
Figure 2.1: **An analogy for quantum state reconstruction.** An experimenter illuminates an unknown, "cloudy" object, from several angles and observes its shadows. After some pondering, they infer the object's three-dimensional shape. In this analogy with quantum state reconstruction, the cloudy object corresponds to the quantum state, the shadows to distributions of physical observables, and casting light to performing measurements. The object is "cloudy" in the sense that it is not directly observable. This figure used excerpts from reference [4]

$\{\langle \hat{O}_k \rangle\}_{k \in K}$, find the state $\hat{\rho}$ such that

$$\mathrm{Tr}\left(\hat{O}_k \hat{\rho}\right) = \langle \hat{O}_k \rangle \quad (\forall k \in K) \quad . \tag{2.2}$$

Since we need to measure several expected values, we will need to prepare many copies of the state.

An interesting example of quantum state reconstruction is that of qubit states with the Pauli operators [5]. By measuring the expected values of the operators $\{\hat{\sigma}_z, \hat{\sigma}_x, \hat{\sigma}_y\}$,

we can reconstruct the state of a qubit system with the expression [6]

$$\hat{\rho} = \frac{1}{2}\hat{1} + \frac{1}{2}\big[\langle\hat{\sigma}_x\rangle\,\hat{\sigma}_x + \langle\hat{\sigma}_y\rangle\,\hat{\sigma}_y + \langle\hat{\sigma}_z\rangle\,\hat{\sigma}_z\big] \quad . \tag{2.3}$$

In this example, the operators $\{\hat{O}_k\}$ – in this case, $\{\hat{\sigma}_x, \hat{\sigma}_y, \hat{\sigma}_z\}$ – correspond to physical quantities. It is also possible, however, to use other operators, such as the projectors on the eigenstates of those quantities. In fact, in an apparatus that uses projective measurements, this is what usually takes place. The first procedure leads to assessing the entire probability distribution of the physical quantity being measured, not just the average. The second approach is equivalent to measuring the amplitudes of the state in a given basis.

It is possible to recast the state $\hat{\rho}$ and the operators $\hat{O}_k$ as column vectors $\vec{\rho}$ and $\vec{O}_k$, and by doing so it can be shown that the trace of the product of $\hat{\rho}$ and $\hat{O}_k$ corresponds to the usual inner product of their vector counterparts. By doing this, the set of equations (2.2) can be condensed in a single matrix equation:

$$\left[\begin{pmatrix}\vec{O}_1\end{pmatrix} \cdots \begin{pmatrix}\vec{O}_n\end{pmatrix}\right]\vec{\rho} = \begin{bmatrix}\langle\hat{O}_1\rangle \\ \vdots \\ \langle\hat{O}_n\rangle\end{bmatrix} \quad \longrightarrow \quad M\vec{\rho} = \vec{o} \quad , \tag{2.4}$$

with the matrix M aggregating the measurement operators in its columns and the vector $\vec{o}$ aggregating the expected values. When M is invertible, the solution $\vec{\rho}$ is unique. In this case, we call the measurement set *informationally complete*[1]. Although this is a desirable property, it does not guarantee that the measurement set will yield good reconstructions for all states[2].

---

[1]Informational-completeness and invertibility are not equivalent. For example, we could have an overcomplete set of measurements that yields a non-square matrix M, which is therefore non-invertible, but is still informationally complete. The interested reader can look into the topic of generalized inverse matrices.

[2]Another desirable property is that the eigenvalues of the measurement matrix M are not very small. If we further change equations (2.4) to the basis of eigenvectors of M, we get

$$\vec{\rho} = \begin{bmatrix}\frac{o_1}{\lambda_1} \\ \vdots \\ \frac{o_n}{\lambda_n}\end{bmatrix} \quad , \tag{2.5}$$

where λ_i are the eigenvalues of M and o_i are the components of the vector of expected values, in the basis of eigenvectors of M. Because of the change of basis, the components o_i will not equal the expected values $\langle\hat{O}_i\rangle$. Therefore, if one of the eigenvalues is small, even a small error in the corresponding measurement would be amplified and result in a bad reconstruction.

The first method that comes to mind for solving equation (2.4) is to find the inverse of M and simply multiplying it on both sides. This method is called *linear inversion*, and though simple, it usually finds non-physical solutions. Because of experimental errors in the measurements of the expected values $\langle \hat{O}_k \rangle$ (due to statistical fluctuations, systematic errors in the setup or errors of any other kind), this method will usually give as result a solution $\hat{\rho}^*$ which is close to $\hat{\rho}$, but not exactly equal. And since the set of trace-one, positive matrices is zero-measured in the space of complex matrices, $\hat{\rho}^*$ will have non-physical features such as eigenvalues with non-vanishing imaginary parts.

Another important aspect to consider is the number of measurement outcomes that is necessary to reconstruct the state. Resorting again to the vector picture of the operators (2.4), we can argue that if the measurement set that we chose can determine any state, then the rows of the measurement matrix M must span the entire operator space. Therefore M must have at least d^2 rows, which means that we would need at least d^2 measurement outcomes.

Besides linear inversion, many methods have been proposed to solve the state reconstruction problem, each having their own advantages and disadvantages. Let us now see some of them.

2.2 Overview of reconstruction methods

In this section we shall briefly overview several methods that have been proposed in the literature. We will not focus so much in the details of each method, as the interested reader would be much better served with the actual papers in which each method was proposed (and which can be found in the references). Our primary interest is in describing the requirements of each method, so we can later compare quantum ptychography with them.

2.2.1 General state reconstruction

Maximum likelihood estimation

The Maximum Likelihood approach [6] resorts to modelling the probability that any state would result in the measured values, and picking the most probable. Let us assume that the experiment employs projective measurements – that is, the observables $\hat{O}_k$ it uses are

projectors on eigenstates of several observables – so that, by the end of the experiment, we have a collection of outcome counts. The number of counts n_l that we will have for an outcome $\hat{O}_l$ follows a Poisson distribution. If the expected number of counts were N_l, and if we further approximate the Poisson distribution by a Gaussian distribution (which is a good approximation if we are in a large-number regime), then the probability that there will be n_l counts is

$$p(n_l|N_l) = Ce^{-\frac{(n_l-N_l)^2}{2N_l}} \quad , \tag{2.6}$$

where C is a normalization constant. Furthermore, if we were to perform the measurement $\mathcal{N}$ times, a state $\hat{\rho}$ would yield the outcome $\hat{O}_l$ a number N_l of times given by

$$N_l = \mathcal{N}\mathrm{Tr}\left(\hat{\rho}\hat{O}_l\right) \quad . \tag{2.7}$$

If we admit that the measurements of the different quantities are statistically independent, then the probability that the experiment will result in the collection of counts $(n_1, n_2, \ldots, n_L) = \vec{n}$ is

$$p(\vec{n}|\hat{\rho}) = \prod_l p(n_l|N_l) = e^{-\frac{1}{2}\sum_l \frac{(n_l-N_l)^2}{N_l}} \quad . \tag{2.8}$$

The probability (2.8) can be used to find the state that would most probably result in the collection of counts $\vec{n}$. Instead of maximizing the probability (2.8), it is computationally easier (and equivalent) to minimize the likelihood function

$$\mathcal{L}(\hat{\rho}) = \frac{1}{2}\sum_l \frac{\left[n_l - \mathcal{N}\mathrm{Tr}\left(\hat{\rho}\hat{O}_l\right)\right]^2}{\mathcal{N}\mathrm{Tr}\left(\hat{\rho}\hat{O}_l\right)} \quad , \tag{2.9}$$

which is just the argument to the exponential function. One approach to solve this optimization problem is to parametrize the set of operators that satisfy the conditions to represent a quantum state (positive-semidefiniteness, unit trace), and optimize (2.9) with respect to those parameters.

This method is used often and yields good results, but it has the disadvantages of being computationally demanding and not scaling efficiently with the state-space dimension[3].

---

[3]In reference [7], in which the authors perform the tomography of a system with 8 trapped-ion qubits, the computational step took weeks [8].

Simple correction and forced purity

Perhaps the simplest method to treat the solution found via linear inversion is to correct its eigenvalues, so that they satisfy the desired properties. We could drop their imaginary parts, and set their real parts to zero in case they are negative. If one knows to be working with pure states (which are rank-one operators), it is possible to further replace the greatest eigenvalue with 1 and set all others to zero.

This approach is fast and simple, but does not yield good results for all states. In particular, reference [9] shows that these methods give good results for pure states, but for other states the reconstructions have low fidelity. The advantage of these methods is, of course, the smaller computational effort.

Compressed sensing

In general, if one has prior information about the state that is being reconstructed (for example, that it is a pure state), it should be possible to narrow down the search space, and thus make for an easier search.

In reference [8], the authors show that, for rank-r, d-dimensional states, it is possible to achieve good reconstructions (with a high success probability) using only $O(rd \log^2 d)$ n-qubit Pauli operators. These operators have the form

$$\hat{A}_i = \bigotimes_{j=1}^{n} \hat{\sigma}_j \quad , \tag{2.10}$$

where each $\hat{\sigma}_j$ is one of the Pauli operators, $\{\hat{\mathbf{1}}, \hat{\sigma}_x, \hat{\sigma}_y, \hat{\sigma}_z\}$. There are $4^n = (2^n)^2 = d^2$ such operators, but only $m = crd \log^2 d$ of them (where c is a sampling ratio that can be chosen) must have their average value determined. This brings a relevant advantage if the reconstructed state is expected to be of low rank.

2.2.2 Pure state reconstruction

For some applications of quantum information, it suffices to use pure states. For example, there are quantum algorithms that are originally formulated with pure states [10–13], and there are models of quantum computation, such as quantum circuits [14] and cluster states [15], that use pure states. Therefore, it makes sense to search for reconstruction

methods that are specific to this context, as they will generally be faster and simpler, though not general.

Minimum requirement on the number of measurements

In reference [16], the authors show that a *Positive Operator-Valued Measure* (POVM[4]) with $2d - 1$ measurement outcomes is not sufficient to reconstruct pure states of dimension d, in the sense that it does not single out a unique state. They also construct a general POVM with $2d$ measurement outcomes that is sufficient, though it would not be good in realistic (and noisy) conditions, and thus only serves as a theoretical construct to show that POVMs with $2d$ measurement outcomes are informationally complete for pure states[5]. This sets a lower bound on the number of measurement outcomes needed to reconstruct pure states, and we can expect the required number to grow linearly with d (as opposed to quadratically for mixed states).

Adaptive measurements with Gell-Mann bases

Reference [18] shows a method that uses up to $5d$ measurement outcomes. It consists of projective measurements in four different bases, which are built upon the $SU(d)$ group

---

[4]POVMs are the most general kind of measurement in quantum mechanics. They are needed to represent, for example, the effect on a subsystem of a projective measurement on the entire system. Mathematically, a POVM is a set of operators $\{\hat{O}_1, \ldots, \hat{O}_k\}$ that satisfy the following conditions:

1. each operator $\hat{O}_k$ is positive-semidefinite,

2. they all sum to the identity operator: $\sum_j \hat{O}_j = \hat{\mathbf{1}}$.

Sometimes, the POVM elements (that is, the operators $\hat{O}_j$) are called POVMs. Though this is not really correct, the context usually makes clear which the author is referring to (the set or the operators).

While the POVM represents a measurement, each of its elements represents a possible outcome. To illustrate what a POVM is, lets us say we were to measure some physical quantity, say the spin of a particle in a given direction. We could use the spectral theorem and decompose the quantity into its eigenvectors and eigenvalues, and we also know that the eigenvectors would form a complete basis on the state space (since the quantity is a hermitian operator). The projectors onto the eigenstates would each correspond to a possible outcome of our measurement, and their sum would amount to the identity operator (since they form a complete basis). Therefore, they constitute a POVM.

[5]For higher dimensions, the lower bound is even higher. In reference [17] (proposition 1), the authors show that $3d$ measurement outcomes are still not enough for $d \geq 9$.

generators previously studied by Gell-Mann [19]. Namely, the four bases are

$$\mathcal{B}_1 = \{|2\nu\rangle \pm |2\nu + 1\rangle\} \quad , \tag{2.11}$$

$$\mathcal{B}_2 = \{|2\nu\rangle \pm i\,|2\nu + 1\rangle\} \quad , \tag{2.12}$$

$$\mathcal{B}_3 = \{|2\nu + 1\rangle \pm |2\nu + 2\rangle\} \quad , \tag{2.13}$$

$$\mathcal{B}_4 = \{|2\nu + 1\rangle \pm i\,|2\nu + 2\rangle\} \quad , \tag{2.14}$$

where $\nu \in \left[0, \left\lfloor \frac{d-2}{2} \right\rfloor \right]$ and $d \geq 3$. If d is odd, then $|d\rangle$ is also included in all the bases. These measurements then lead to a set of equations that can be solved recursively, yielding the components of the reconstructed state.

The recursive equations do have a mathematical shortcoming, though: they do become indeterminate if the reconstructed state has non-consecutive vanishing components. However, the set of states for which this happens has a measure of zero in the set of all states[6], so that this situation would not occur in practice. The authors did formulate a work-around, though, which is to make an initial measurement in the computational basis, and in case the state is verified to fall under this unfortunate category, apply the method to the subspace of the non-vanishing components. In this approach, the measurement bases are not strictly defined prior to the measurements, though.

Therefore, this method usually requires $4d$ measurement outcomes to reconstruct the state, but might require $5d$ if the adaptive approach is to be used (i.e., the one using the aforementioned workaround).

Phase-stepping interferometry

In reference [20], the authors propose a method based in another optical phase-retrieval technique, called phase-stepping interferometry. It uses a phase reference that is interfered with the other parts of the signal at phase steps of $\frac{\pi}{2}$. From the interference amplitudes, it is possible to determine the phase of the unknown part of the signal.

The authors show that the quantum analogue of phase-stepping interferometry can be achieved with projections on the states

$$|\Psi_\ell^{(k)}\rangle = \frac{1}{\sqrt{2}} \left[|0\rangle + e^{i\,\pi(l-1/2)/2}\,|k\rangle \right] \quad , \tag{2.15}$$

---

[6]The measure of a set can be though as its *n-dimensional volume* (if we picture it as embedded in an *n*-dimensional space). A zero-measured set is thus similar to a curve or a hyper-surface (or, perhaps, a countable union of them) with dimension smaller than the space it is embedded in.

where $\ell = 1, 2, 3$ and $k = 0, 1, \ldots, d - 1$. The projections of the state $|\psi\rangle = \sum_{j=0}^{d-1} c_j |j\rangle$ to be reconstructed are

$$| \langle \Psi_\ell^{(k)} | \psi \rangle |^2 = p_\ell^{(k)} = \frac{1}{2} \left[|c_0|^2 + |c_k|^2 \right] + \Re \left(c_0 c_k^* e^{i \pi (l - 1/2)/2} \right) \quad . \tag{2.16}$$

From these equations, we can arrive at

$$\sqrt{2} c_0 c_k^* = \left(p_1^{(k)} - p_2^{(k)} \right) + i \left(p_3^{(k)} - p_2^{(k)} \right) \quad , \tag{2.17}$$

and with the additional projection $\langle 0 | \psi \rangle$, we are able to assess c_0 and solve the set of equations and reconstruct $|\psi\rangle$. The authors mention that it is benefical to include measurements on the whole computational basis in order to certify the purity of the state. The set of equations becomes indefinite for states where $c_0 = 0$, and the method would not work for them. These states also form a zero-measured set, though.

The authors also show an experimental implementation and collect many reconstruction results for $d = 14$. The donwside of their implementation was that they used an extra dimension as the phase reference. This effectively embedded the state in a $d + 1$ dimensional state, and this resource might not be feasible for some experimental contexts.

Taking into account the measurements in the computational basis, this method amounts for a total of $4d - 3$ measurement outcomes.

Semidefinite programming with orthogonal-polynomials bases

Finally, reference [21] shows how to construct five informationally-complete bases from any sequence of orthogonal polynomials. Their work is especially interesting as it brings general theorems on reconstruction schemes, besides their own method.

From a sequence $(p_n)_{n=0}^\infty$ of orthogonal polynomials, where each p_n is a polynomial of degree n, the authors pick the first d polynomials, and also the roots of p_d and p_{d-1}, which we will denote by x_j and y_k, respectively:

$$p_d(x_j) = 0, \quad j = 0, 1, \ldots, d - 1 \tag{2.18}$$

$$p_{d-1}(y_k) = 0, \quad k = 0, 1, \ldots, d - 2 \quad . \tag{2.19}$$

They take advantage of the fact that the roots of p_i are not roots of p_{i-1}, and choose a

real α such that $e^{ij\alpha} \notin \mathbb{R}$ to construct the family of $4d$ vectors

$$v_j^1 = [p_0(x_j), p_1(x_j), \ldots, p_{d-1}(x_j)] \tag{2.20}$$

$$v_j^2 = \left[p_0(x_j), e^{i\,\alpha}p_1(x_j), \ldots, e^{i\,(d-1)\alpha}p_{d-1}(x_j)\right] \tag{2.21}$$

$$v_j^3 = [p_0(y_j), p_1(y_j), \ldots, p_{d-2}(y_j), 0] \tag{2.22}$$

$$v_j^4 = \left[p_0(y_j), e^{i\,\alpha}p_1(y_j), \ldots, e^{i\,(d-2)\alpha}p_{d-2}(y_j), 0\right] \tag{2.23}$$

$$\tag{2.24}$$

with $j = 0, 1, \ldots, d-1$. These families are then normalized and aggregated in four measurement bases,

$$\mathcal{B}_k = \left\{ \frac{v_j^k}{||v_j^k||} \right\}, \quad k = 1, 2, 3, 4 \quad, \tag{2.25}$$

and the bases are finally completed with the computational basis. In their simulations, the authors chose the Chebychev polynomials of the second kind, and $\alpha = \pi/d$.

The authors then use the measurements on these bases to reconstruct the unknown state via *semi-definite programming*[7]. Their method entails, therefore, $5d$ measurement outcomes to reconstruct a state

2.3 Chapter summary

We have seen the formulation and several methods for quantum state reconstruction, both in the general and pure state case. General state reconstruction needs d^2 measurement outcomes, while the pure state case requires at least $2d$ measurement outcomes. In practice, current pure state methods in the literature have been using around $4d$ or $5d$ measurement outcomes. And, in general, different methods resort to different bases on which projective measurements are to be made.

---

[7] Semi-definite programming problems are a class of convex optimization problems, where the optimized function is linear and the search space is an intersection of the space of positive semidefinite matrices with affine spaces. Because of its convex nature, there are many efficient algorithms to solve it that are commonly implemented in numerical libraries [22–26]

Chapter 3

Classical wavefront reconstruction

In this chapter, we will show how to formulate pure quantum state reconstruction as a phase retrieval problem (section 3.1), and then we will study how this problem was first solved in the context of classical optics (section 3.2). We will then see the shortcomings of this approach, and then present a more modern approach called ptychography (section 3.3), which we have adapted to the quantum reconstruction context.

3.1 Pure quantum state reconstruction as a phase retrieval problem

Let us consider the reconstruction of states in a discrete, d-dimensional Hilbert space $\mathcal{H}_d$. A pure quantum state $|\psi\rangle$ in $\mathcal{H}_d$ can be written in the computational basis $\{|0\rangle, \ldots, |d-1\rangle\}$ as

$$|\psi\rangle = \sum_{k=0}^{d-1} c_k |k\rangle \quad , \tag{3.1}$$

and reconstructing it amounts to assessing all of its complex coefficients c_k. In special, if we were to make measurements in the computational basis, they would give us the amplitudes of the coefficients. In order to reconstruct $|\psi\rangle$, we would only be lacking knowledge of their phases.

It is not easy, however, to formulate a hermitian operator that would measure the phases of an arbitrary state directly. In fact, for any observable that we come up with, the probability distribution of its outcomes would be given by the squared amplitudes of
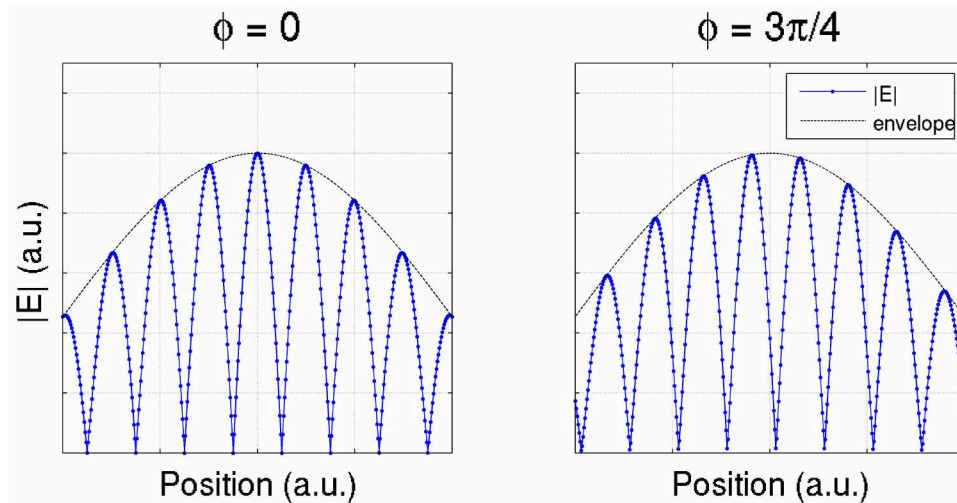
Figure 3.1: **Double-slit diffraction pattern.** The vertical axis shows the amplitude of the electrical field in the double-slit diffraction. The relative phase ϕ between the slits displaces the oscillatory pattern inside the amplitude envelope.

ψ's coefficients in the basis of its eigenstates. Therefore, the measurements would not give us the desired phases, at least not directly. How, then, could we infer the phases of the components of $|\psi\rangle$?

We must somehow cast the phase information into amplitude information. The reader might recall, at this point, that the amplitudes of the Fourier transform of a function depend on its phases. To make this fact more concrete, let us remember the double slit diffraction: when illuminated with the same intensity by a coherent source, the two slits produce a diffraction pattern consisting of a sinc envelope[1] and a sinusoidal oscillation (which is the Fourier transform of the double slit aperture function). By giving a phase to one of the slits (relative to the other), we can translate the oscillations inside the envelope, as depicted in figure 3.1. Conversely, it should be possible to infer the relative phase between the slits from the displacement of the oscillations in the diffraction pattern.

This fact leads to the following formulation of the *phase retrieval problem* [27], which can be stated both for continuous and discrete functions:

---

[1]The sinc function, which stands for *sine cardinal*, is usually defined as

$$\text{sinc}(x) = \frac{\sin(x)}{x} \quad ,$$

though there are other conventions that can be found in the literature.

> ### Phase-retrieval problem:
>
> Given:
>
> - the modulus of a (discrete) function $|f(x)|$
>
> - the modulus of its (discrete) Fourier transform $|F(u)|$
>
> determine the function $f(x)$.

Although we could use other transforms for this purpose, the Fourier transform is the most prevalent in the literature, perhaps because of the relative ease to implement it in optical experiments.

It is not simple, however, to directly derive the desired phases from these amplitude measurements. As we will see later, it is even possible that the measurements are consistent with more than one solution. Still, it is possible to find the phases corresponding to a given set of measurements numerically. The method to do so was first developed by Gerchberg and Saxton [27], and we shall study it next.

3.2 Standard phase retrieval: the Gerchberg-Saxton algorithm

The Gerchberg-Saxton algorithm [27] (GSA) was originally proposed in the context of electron microscopy, and soon found application in other imaging problems. Their approach was, perhaps, the simplest thing one could try.

But before we delve into the algorithm itself, let us quickly describe the physical situation during a microscopy. Imagine that there is a small, microscopic sample which we wish to image, such as some cells in a slide or a thin film. To this end, we can impinge a plane wave on it. The object will transmit the plane wave, but will also alter its amplitude and phase profiles. This will alter its diffraction pattern as well. If we could then reconstruct the wavefront immediately after the sample, we could image the object itself.

Imagine then that we had intensity measurements of the transmitted wave and of its diffraction pattern, giving us $|f|$ and $|F|$, respectively. If we were to assign a random phase profile to f, chances are that we would not guess them correctly, since there are so

many possible profiles. Moreover, we would be able to *tell* that we had a wrong guess, as the diffraction pattern of this "estimate" would not match the measured intensities $|F|$. We could, however, correct these amplitudes by imposing the measured intensities. By doing so, we are creating a new estimate which has a different diffraction pattern, compared to the previous, and consequently different phases – we have corrected the phases at the object domain[2]. Now, this new estimate probably does not have the correct amplitudes in the object domain, but we can make a similar correction using the measurements of $|f|$. Iterating over this procedure, we will hopefully arrive at the correct function f. This is the procedure that Gerchberg and Saxton proposed:

Gerchberg-Saxton algorithm:

Given the measurements of $|F|$ and $|f|$:

(i) make an estimate $\phi(x)$ of the phases in the object domain, then build an estimate g of the function f:

$$g(x) = |f(x)|e^{i\phi(x)};$$

(ii) take the Fourier transform G of the estimate g:

$$G(u) = \mathcal{F}[g] = |G(u)|e^{i\psi(u)};$$

(iii) correct the amplitudes of the Fourier transform G, keeping its phases:

$$G'(u) = |F(u)|e^{i\psi(u)};$$

(iv) take the inverse Fourier transform of G' to arrive at a new estimate g' of f:

$$g'(x) = \mathcal{F}^{-1}[G] = |g'(x)|e^{i\phi'(x)}.$$

The phases $\phi'(x)$ of $g'(x)$ from step (iv) can be used as the phase estimate $\phi(x)$ in step (i). The steps are repeated until some stop criterion is met.

The GSA is also illustrated in figure 3.2. In their article [27], Gerchberg and

---

[2]Following the nomenclature in optics, we will refer to the domains of f and F as object and Fourier domains, respectively.

Initial phase estimate

$$g(x) = |f(x)|e^{i\phi(x)} \xrightarrow{\text{FT}} G(u) = |G(u)|e^{i\psi(u)}$$

impose $|f|$ impose $|F|$

$$g'(x) = |g'(x)|e^{i\phi'(x)} \xleftarrow{\text{iFT}} G'(u) = |F(u)|e^{i\psi(u)}$$
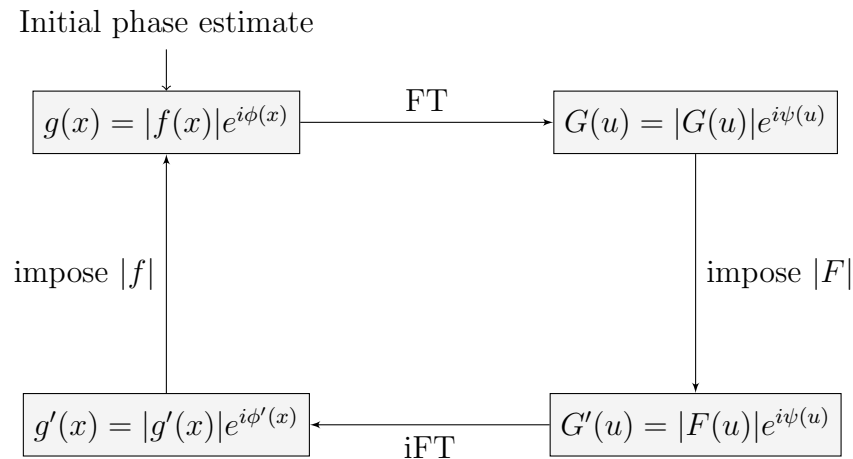
Figure 3.2: **Gerchberg-Saxton iterative algorithm.** This diagram shows the steps and the looping flowchart of this method.

Saxton have proved that this algorithm has a weak convergence property: each iteration of the algorithm will not increase the overall errors[3] of the estimate. Their proof relies on Parseval's theorem, and on the fact that correcting only the amplitudes of the estimate while keeping their phases is the smallest possible correction, in the Euclidean sense.

Now there are two interesting points about the weak convergence of the algorithm. First, it shows that the proposed corrections are in fact sensible, and do indeed correct the estimate in general (or at least, they never worsen it). Second, it hints that the GSA is related to the gradient descent method. This fact was shown later in 1982, by James R. Fienup [28].

It is not surprising, therefore, that this algorithm can suffer with local minima, i.e., the estimate can get trapped at a local minimum of the error metric E, and the algorithm would not be able to move away from that point. Even worse, it is entirely possible that more than one function f is compatible with the intensity measurements $|f|$ and $|F|$. It happens often that reconstructions face some setback related to this non-uniqueness,

---

[3]We can define two overall errors of our estimate, one in the object domain,

$$E_O = \frac{\sum_x \left[|g'(x)| - |f(x)|\right]^2}{\sum_x |f(x)|^2} \quad ,$$

and the other in the Fourier domain,

$$E_F = \frac{\sum_u \left[|G(u)| - |F(u)|\right]^2}{\sum_u |F(u)|^2} \quad .$$

which often translates to local stagnation problems [29]. In practice, though, this has not stopped the technique to give fruitful reconstruction results [30].

There are a number of ways to avoid these shortcomings. One of them is to use more data, so as to rule out ambiguities and possibly help the error metric to not have as many local minima. This is the approach that Rodenburg and Faulkner [31] took when proposing another method called *ptychography*, which we will discuss next.

3.3   Classical ptychography

In their 2004 article [31], Rodenburg and Faulkner propose a technique that uses more data – or a greater data diversity – to solve the phase retrieval problem. To illustrate their proposal, suppose we are in an ambiguous (or non-unique) scenario: we have ran the GSA several times, collected its results, and saw that there are a few different estimates for the function that generated the measured amplitude profiles, all consistent with the measurements. One manner to overcome this problem would be to use another measured quantity, so that all estimates but one would be inconsistent with it and could thus be discarded.

Perhaps the additional quantity that could be most naturally measured with our apparatus was the amplitudes in an intermediate plane not corresponding to either object or Fourier domains[4]. The proposal of Rodenburg and Faulkner achieves this greater diversity in the data with a shifting illumination of the sample. The probing beam is focused to a small spot, so as to illuminate only a small area of the sample. The sample (or equivalently, the probe beam) is then translated so that many spots are illuminated. Each of the small areas produce a diffraction pattern that is recorded, and will be used in an iterative algorithm, similar to the GSA. It is crucial to keep some overlap between neighboring areas, else the problem becomes several disjoint phase retrievals, each prone to suffering with non-uniqueness problems. This method is called *ptychography* and is illustrated in figure 3.3.

The reconstruction algorithm in this scenario has the same underlying idea as the GSA. But first, we need to establish our notation. We shall consider that the object

---

[4]Diffraction at intermediate distances corresponds to fractional Fourier transforms [32], which is a family of operations that generalizes the standard Fourier transform. For the purposes of this section, it suffices to know that the fractional transform is linearly independent to the standard transform, and that could give new information to the phase-retrieval algorithm (i.e., information that is not redundant).
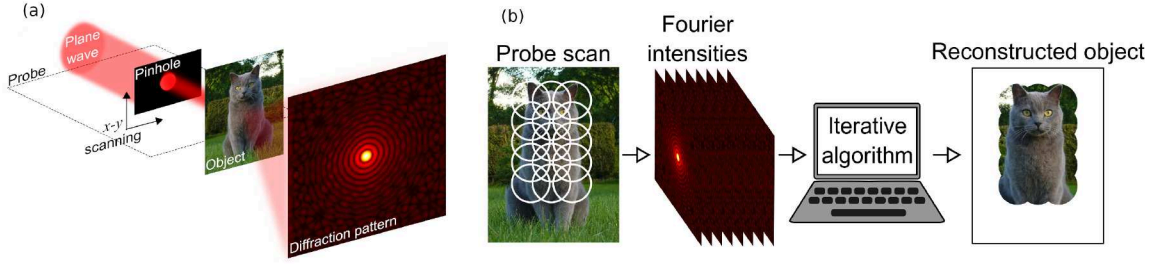
Figure 3.3: **Ptychographic scenario**. (a) A small region of the object is illuminated, and the resulting diffraction pattern is recorded. (b) The probe beam is shifted along the object in a structured pattern, and the many diffraction patterns are recorded. They are later fed to a computational algorithm that reconstructs its transfer function.

that we want to image is very thin, and can be considered as a two-dimensional material. We will denote the (two-dimensional) transverse position along it as $\vec{x}$. Admitting that the probe beam that illuminates the object is coherent and has a constant polarization across the entire setup, we can represent its wavefront immediately before the object as a complex function $P(\vec{x})$ and also leave the polarization aside. If we neglect transverse refraction, the interaction between beam and object can be represented by a complex transmission function $t(\vec{x})$, which encompasses the phase shifts and absorptions the object introduces to the beam, so that the beam that emerges from the object is

$$f(\vec{x}) = t(\vec{x})P(\vec{x}) \quad .$$

The shifting illumination can be achieved by translating the probe beam with respect to the object (in practice, it is simpler to translate the object, and leave the optical part of the setup still). If we let the set of transverse positions $\{\vec{r}_\ell\}_\ell$ encode these translations, then for each of them we will be generating a wavefront given by

$$o_\ell(\vec{x}) = t(\vec{x})P(\vec{x} - \vec{r}_\ell) \quad . \tag{3.2}$$

These wavefronts will then diffract, and the resulting amplitude profiles

$$|O_\ell| \propto |\mathcal{F}[o_\ell]| \tag{3.3}$$

will be measured in a further plane. The set $\{|O_\ell(\vec{u})|\}_\ell$ of amplitude profiles, together with the probe profile P and the illumination positions $\{\vec{r}_\ell\}_\ell$, constitutes the data that will be fed into to the reconstruction algorithm. Figure 3.3 illustrates the ptychographic scenario.

Similar to the GSA, the reconstruction algorithm for ptychography, known as *ptychographic iterative engine* (PIE) [31], proceeds by taking an estimate through the same steps as the sample (illumination, diffraction) and correcting the resulting amplitudes iteratively. Formally, the steps are the following:

Ptychographic iterative engine:

Given the set of diffraction amplitudes $\{|O_\ell(\vec{u})|\}_\ell$, the probe profile $P(\vec{x})$ and the set of centering positions $\{\vec{r}_\ell\}_\ell$:

(i) make an estimate $\phi(\vec{x})$ of the object;

(ii) apply the probe illumination, centered at position $\vec{r}_j$, on the estimate:

$$\psi_j(\vec{x}) = P(\vec{x} - \vec{r}_j)\phi(\vec{x});$$

(iii) apply the Fourier transform on the result:

$$\Psi_j(\vec{u}) = \mathcal{F}[\psi_j](\vec{u});$$

(iv) correct the amplitudes of Ψ_j, imposing the measured amplitudes $|O_j|$ but keeping its phases:

$$\Psi'_j(\vec{u}) = |O_j(\vec{u})|\, e^{i \arg[\Psi_j(\vec{u})]};$$

(v) apply the inverse Fourier transform:

$$\psi'_j(\vec{x}) = \mathcal{F}^{-1}[\Psi'_j](\vec{x});$$

(vi) correct the estimate ϕ, with greater weight to the points where its illumination was more intense:

$$\phi_c(\vec{x}) = \phi(\vec{x}) + \beta \frac{|P(\vec{x} - \vec{r}_j)|}{|P_{\max}|} \frac{P^*(\vec{x} - \vec{r}_j)}{(|P(\vec{x} - \vec{r}_j)|^2 + \delta)} \left[\psi'_j(\vec{x}) - \psi_j(\vec{x})\right],$$

where $|P_{\max}|$ is the maximum illumination intensity, β is the correction rate and δ is a regularization term; both β and δ can be tuned to enhance the performance of the algorithm;

(vii) repeat steps (ii)-(vi) using the new estimate ϕ_c and a new position $\vec{r_j'}$, until a stop criterion is met.

The PIE algorithm is illustrated in figure 3.4.
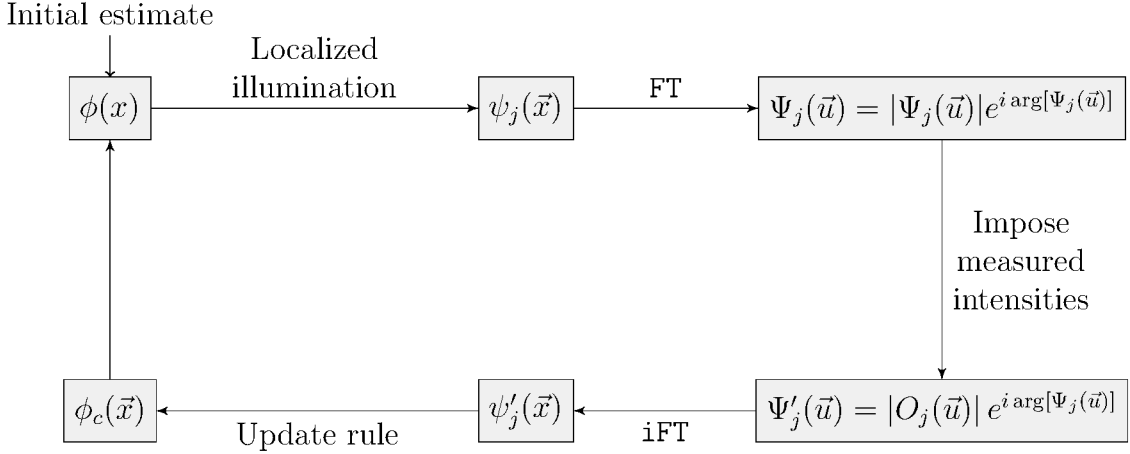
Figure 3.4: **Ptychographic Iterative Engine algorithm.** This diagram shows the steps and the flow in the ptychographic method.

Not surprisingly, the PIE is also related to gradient search methods. It is possible to show [29] that the gradient of the error metrics

$$E_\ell = \sum_{\vec{u}} \left(|\Psi_\ell(\vec{u})| - |O_\ell(\vec{u})| \right)^2 \qquad (3.4)$$

with respect to the estimate ψ_ℓ is

$$\left[\frac{\partial}{\partial \psi_\ell^R} + i \frac{\partial}{\partial \psi_\ell^I} \right] E_\ell = 2 \left(\psi_\ell' - \psi_\ell \right) \quad . \qquad (3.5)$$

Therefore, each update step in the PIE corrects the estimate in the same direction that a gradient-descent would[5], with respect to a metric that depends on which illumination profile is being used. By alternating the profiles, and thus the metric that is being minimized, the algorithm effectively avoids stagnating in local minima of a single error metric, as further iterations with other profiles will remove the estimate from such an eventual minimum. This is how information diversity avoids the non-uniqueness problem.

---

[5]The update step in the PIE actually uses a Wiener filter to limit the gradient-descent correction on regions that were not intensely illuminated. It is, therefore, a *Wiener-filtered* gradient-descent.

The ptychographic approach has been used with success in many applications of microscopy and crystallography. We refer the reader to references [33, 34] for some interesting examples.

Now that we understand the ptychographic approach, it is time to turn our attention back to the task of reconstructing quantum states. As we saw in the beginning of this chapter, reconstructing quantum states is analogous to reconstructing wavefronts. We are now in position to take advantage of this fact and translate the ptychographic approach to the quantum realm.

3.4   Chapter summary

We have seen how the reconstructions of a pure quantum state and of a coherent wavefront are similar. We discussed the Gerchberg-Saxton Algorithm, one of the first iterative algorithms developed to reconstruct wavefront phases, and the Ptychographic Iterative Engine, an enhancement of the GSA.

Chapter 4

Pure quantum state ptychography: theoretical results

When we designed the quantum ptychographic method [35], the main aspect we had in mind was experimental simplicity. While most pure state recontruction methods use measurements in several bases, as we saw in the last chapter, quantum ptychography uses a different strategy to achieve data diversity. In fact, it uses a single measurement basis, which simplifies experiments considerably. This method is the quantum analogue of classical ptychography we saw in section 3.3. Let us see how it works.

4.1 The ptychographic method

In classical ptychography (sec. 3.3), there were three main entities at play: an object whose phases we wanted to reconstruct, a shifting illumination that impinged on that object, and a measurement device that recorded amplitudes in the Fourier plane. In the quantum scenario, the object will be the quantum state itself, the shifting illumination will be performed by a set of projectors (which we will detail further), and the Fourier-plane measurements will be performed by projective measurements in the Fourier-transformed computational basis. Let us now formulate these ideas precisely.

As we saw earlier, an arbitrary pure quantum state in a d-dimensional Hilbert

space $\mathcal{H}_d$ may be written in the computational basis $\{|k\rangle\}_{k=0}^{d-1}$ as

$$|\psi\rangle = \sum_{k=0}^{d-1} c_k |k\rangle \quad , \tag{4.1}$$

where $\sum_k |c_k|^2 = 1$. This is the object we want to reconstruct, meaning that we want to determine the d complex coefficients $\{c_k\}_{k=0}^{d-1}$. In our ptychographic approach, the role of the localized, scanning, and partially overlapping illumination probe will be played by a set of n projectors, $\{\hat{P}_\ell\}_{\ell=0}^{n-1}$, onto r-dimensional subspaces ($1 < r < d$) of $\mathcal{H}_d$. For the sake of simplicity, let us admit that all the projectors have the same rank. Similar to the shifting illumination spots, these projectors have to satisfy two conditions:

(i) each projector in the set must have a partial overlap with at least one other partner, i.e., for any $\hat{P}_\ell$ there must be a $\hat{P}_{\ell'}$ such that

$$0 < \frac{\text{Tr}(\hat{P}_\ell \hat{P}_{\ell'})}{r} < 1 \quad ; \tag{4.2}$$

(ii) all dimensions in $\mathcal{H}_d$ must be addressed at least once (or, equivalently, the set of the states corresponding to the projectors must span $\mathcal{H}_d$).

We will see some suitable projector families later on. For now, though, let us focus on the general features of the protocol.

Given an ensemble of d-dimensional quantum systems described by the state $|\psi\rangle$, we proceed to collect ptychographic measurements. We first apply the ℓ-th probe projection on the input, generating an output sub-ensemble described by the (unnormalized) state

$$|\psi_\ell\rangle = \hat{P}_\ell |\psi\rangle \quad . \tag{4.3}$$

Next, we apply a quantum Fourier transform (QFT)[1] on this output, obtaining

$$|\widetilde{\psi_\ell}\rangle = \hat{\mathcal{F}}_d |\psi_\ell\rangle = \sum_{k=0}^{d-1} \tilde{c}_{k\ell} |k\rangle \quad , \tag{4.4}$$

---

[1] The QFT operator in a d-dimensional space, $\hat{\mathcal{F}}_d$, is defined as

$$\hat{\mathcal{F}}_d = \frac{1}{\sqrt{d}} \sum_{j,k=0}^{d-1} e^{i \frac{2\pi jk}{d}} |j\rangle \langle k| \quad .$$

where $\{\tilde{c}_{k\ell}\}_{k=0}^{d-1}$ is the set of Fourier-transformed amplitudes of $|\psi_\ell\rangle$. Finally, we perform projective measurements onto the computational basis. This procedure is repeated for each $\hat{P}_\ell$ and gives us a set of n count distributions,

$$\Pi_\ell = \left\{ \mathcal{N} \, |\tilde{c}_{k\ell}|^2 \right\}_{k=0}^{d-1} \quad , \tag{4.5}$$

where $\mathcal{N}$ is a constant dependent on particle flux and detector efficiencies. These distributions form our ptychographic data set as $\{\sqrt{\Pi_\ell}\}_{\ell=0}^{n-1}$, which, together with the set of probe projectors, will be the inputs to an adapted version of the PIE (see section 3.3). This process is illustrated in figure 4.1.
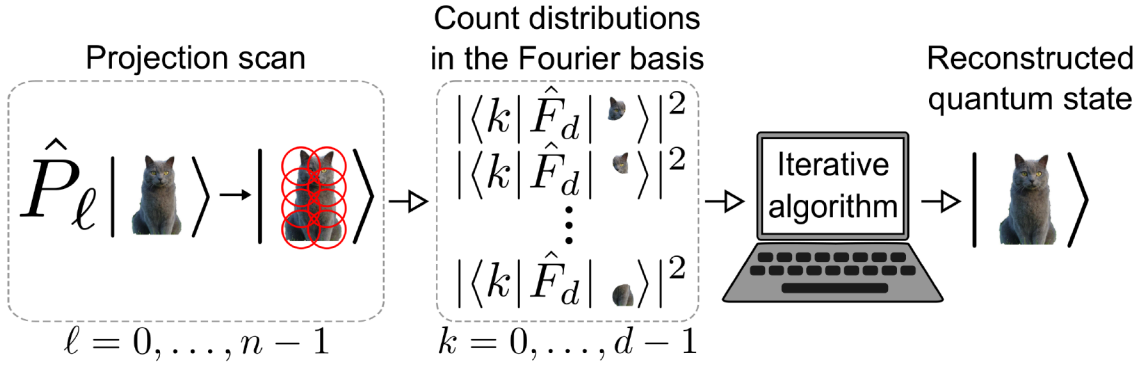
Figure 4.1: **Steps in the quantum ptychographic process.** We apply the n probe projectors $\hat{P}_\ell$, which play a role analogous to the shifting illumination, to the unknown state. We then accumulate the count distributions in the Fourier basis and acquire the ptychographic dataset $\{\sqrt{\Pi_\ell}\}_{\ell=0}^{n-1}$. This dataset, together with the probe projectors, is then fed to an iterative algorithm, which reconstructs the state.

The iterative reconstruction proceeds by taking a random initial estimate through the same steps as the ensemble, imposing the measured data and applying an update rule to the estimate. The steps are shown in the next box.

Pure state quantum ptychography

(i) Start with a random estimate of the input state:

$$|\phi\rangle = \sum_{k=0}^{d-1} \gamma_k |k\rangle \quad . \tag{4.6}$$

(ii) Apply the ℓ-th probe projector to $|\phi\rangle$:

$$|\phi_\ell\rangle = \hat{P}_\ell|\phi\rangle = \sum_k \gamma_{k\ell}|k\rangle \quad . \tag{4.7}$$

(iii) Apply the QFT to $|\phi_\ell\rangle$:

$$|\tilde{\phi}_\ell\rangle = \hat{\mathcal{F}}_d\,|\phi_\ell\rangle = \sum_{k=0}^{d-1} \tilde{\gamma}_{k\ell}\,|k\rangle \quad . \tag{4.8}$$

(iv) Use the ℓ-th measured ptychographic data, $\sqrt{\Pi_\ell}$, to correct the moduli of the coefficients of $|\tilde{\phi}_\ell\rangle$, keeping their phases:

$$|\tilde{\phi}'_\ell\rangle = \sqrt{\mathcal{N}} \sum_{k=0}^{d-1} |\tilde{c}_{k\ell}| e^{i\arg\tilde{\gamma}_{k\ell}}\,|k\rangle \quad . \tag{4.9}$$

(v) Apply the inverse QFT to obtain an updated estimate for the output state:

$$|\phi'_\ell\rangle = \hat{\mathcal{F}}_d^{-1}\,|\tilde{\phi}'_\ell\rangle \quad . \tag{4.10}$$

(vi) Update the current estimate of the input state:

$$|\phi'\rangle = |\phi\rangle + \beta\hat{P}_\ell\big(\,|\phi'_\ell\rangle - |\phi_\ell\rangle\,\big) \quad , \tag{4.11}$$

where β is a feedback parameter, roughly within $(0,2]$, that controls the step-size of the update and can be adjusted to improve convergence (see appendix A.1).

(vii) Use this updated estimate as input to repeat the steps (ii)–(vi) with a new value of ℓ.

This sequence is summarized in the diagram of figure 4.2: a single PIE iteration consists of n iterations through the closed loop (steps (ii)–(vii)), where each probe projector and corresponding ptychographic data is used once to update the state estimate. At each iteration we calculate the relative distance between the current and updated estimates,

$$D = \frac{\||\phi'\rangle - |\phi\rangle\|^2}{\||\phi\rangle\|^2} \quad . \tag{4.12}$$

The algorithm terminates when it achieves either a sufficiently small value of D or a preset maximum number of PIE iterations. The final estimate is a pure state that must be normalized.
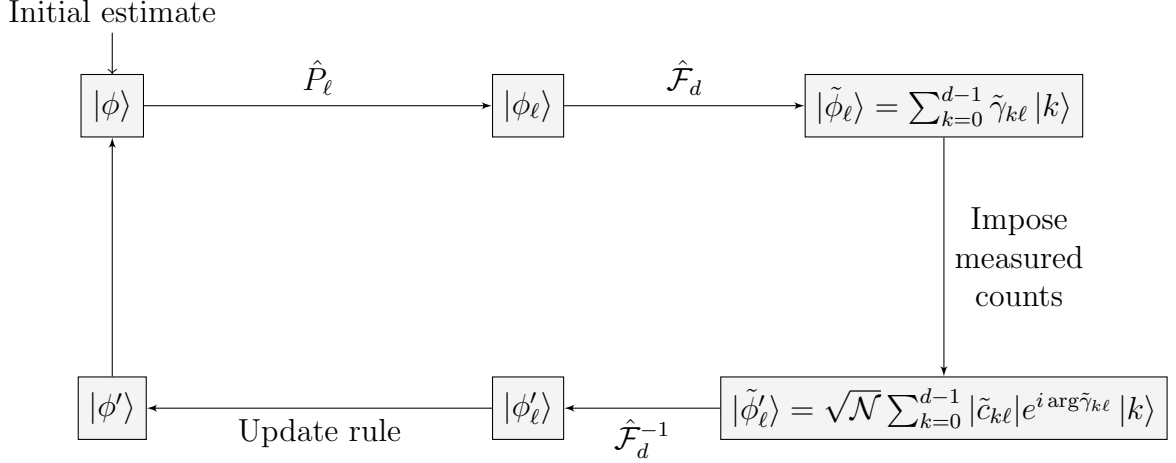
Figure 4.2: **Pure state, quantum ptychographic algorithm.** Sequence of steps in the iterative algorithm. The sequence is repeated until a stop criterion is met (see text).

This algorithm is simple and fast, but it is also possible to use other post-processing techniques. For example, one could use SDP (see section 2.2.2) to find the pure state that is most compatible with the ptychographic measurements.

4.2 Relation to alternating-gradient search

Before we go into the choices of the projector family, we will show that the update rule (4.11) is equivalent to a gradient-descent with some weighting factor. We will follow the same reasoning that has been applied in classical ptychography [28, 29]. More precisely, we will show that an iteration with the projector $\hat{P}_\ell$ corrects the estimate in the same manner as a gradient descent would, with respect to the error metric

$$E_\ell = \sum_k \left[|\tilde{\gamma}_{k\ell}| - |\tilde{c}_{k\ell}| \right]^2 \quad . \tag{4.13}$$

We will calculate the gradient of E_ℓ with respect to the coefficients $\gamma_{k\ell}$ of the estimate (equation (4.7)) to find the direction of fastest descent. To this end, we calculate the

complex gradient[2] of E_ℓ,

$$\left[\frac{\partial}{\partial \gamma_{j\ell}^{\mathrm{R}}} + i \frac{\partial}{\partial \gamma_{j\ell}^{\mathrm{I}}} \right] E_\ell \quad ,$$

where $\gamma_{j\ell}^{\mathrm{R}}$ and $\gamma_{j\ell}^{\mathrm{I}}$ denote the real and imaginary parts of $\gamma_{j\ell}$, respectively. The chain rule, applied to equation (4.13), yields

$$\frac{\partial E_\ell}{\partial \gamma_{j\ell}} = 2 \sum_k \left[|\tilde{\gamma}_{k\ell}| - |\tilde{c}_{k\ell}| \right] \frac{\partial |\tilde{\gamma}_{k\ell}|}{\partial \gamma_{j\ell}} \quad , \tag{4.14}$$

where $\partial / \partial \gamma_{j\ell}$ denotes differentiation with respect to either the real or imaginary parts of $\gamma_{j\ell}$. The chain rule gives further

$$\frac{\partial |\tilde{\gamma}_{k\ell}|}{\partial \gamma_{j\ell}} = \frac{\partial}{\partial \gamma_{j\ell}} \{ \tilde{\gamma}_{k\ell} \tilde{\gamma}_{k\ell}^* \}^{\frac{1}{2}} = \frac{1}{2|\tilde{\gamma}_{k\ell}|} \left\{ \tilde{\gamma}_{k\ell} \left[\frac{\partial \tilde{\gamma}_{k\ell}}{\partial \gamma_{j\ell}} \right]^* + \tilde{\gamma}_{k\ell}^* \frac{\partial \tilde{\gamma}_{k\ell}}{\partial \gamma_{j\ell}} \right\} \quad . \tag{4.15}$$

The derivatives of $\tilde{\gamma}_{k\ell}$ can be easily calculated if we use their Fourier-transform relation to the coefficients $\gamma_{j\ell}$:

$$\frac{\partial \tilde{\gamma}_{k\ell}}{\partial \gamma_{j\ell}^{\mathrm{R}}} = \frac{\partial}{\partial \gamma_{j\ell}^{\mathrm{R}}} \frac{1}{\sqrt{d}} \sum_m \left(\gamma_{m\ell}^{\mathrm{R}} + i \gamma_{m\ell}^{\mathrm{I}} \right) e^{-i2\pi mk/d} = \frac{1}{\sqrt{d}} e^{-i2\pi jk/d} \quad , \tag{4.16}$$

$$\frac{\partial \tilde{\gamma}_{k\ell}}{\partial \gamma_{j\ell}^{\mathrm{I}}} = \frac{\partial}{\partial \gamma_{j\ell}^{\mathrm{I}}} \frac{1}{\sqrt{d}} \sum_m \left(\gamma_{m\ell}^{\mathrm{R}} + i \gamma_{m\ell}^{\mathrm{I}} \right) e^{-i2\pi mk/d} = \frac{i}{\sqrt{d}} e^{-i2\pi jk/d} \quad . \tag{4.17}$$

Inserting equations (4.16) and (4.17) into (4.15), we arrive at

$$\frac{\partial |\tilde{\gamma}_{k\ell}|}{\partial \gamma_{j\ell}^{\mathrm{R}}} = \frac{1}{\sqrt{d}} \frac{\tilde{\gamma}_{k\ell} e^{i2\pi jk/d} + \tilde{\gamma}_{k\ell}^* e^{-i2\pi jk/d}}{2|\tilde{\gamma}_{k\ell}|} \quad , \tag{4.18}$$

$$\frac{\partial |\tilde{\gamma}_{k\ell}|}{\partial \gamma_{j\ell}^{\mathrm{I}}} = \frac{i}{\sqrt{d}} \frac{-\tilde{\gamma}_{k\ell} e^{i2\pi jk/d} + \tilde{\gamma}_{k\ell}^* e^{-i2\pi jk/d}}{2|\tilde{\gamma}_{k\ell}|} \quad , \tag{4.19}$$

---

[2]It is possible to show that a non-constant, real-valued function over the complex numbers $f(z)$ is not differentiable in the complex sense ($\mathbb{C}$-differentiable, or holomorphic), and this fact hinders the use of the complex derivative directly for gradient methods. However, it is possible to picture f as a function of two real variables (the real and imaginary parts of z), $f(z) = f(x, y)$, so that it is differentiable as a mapping $\mathbb{R}^2 \to \mathbb{R}^2$ ($\mathbb{R}$-differentiable). By doing this trick, one arrives at the complex gradient definition we use here, though it is not the only possible definition. The interested reader should see [36] for a deeper discussion.

and inserting (4.18) and (4.19) back into (4.14), we find

$$\frac{\partial E_\ell}{\partial \gamma_{j\ell}^{\mathrm{R}}} = \frac{1}{\sqrt{d}} \sum_k \left[\tilde{\gamma}_{k\ell} - \tilde{\gamma}'_{k\ell}\right] e^{i2\pi jk/d} + \left[\tilde{\gamma}_{j\ell}^* - \tilde{\gamma}_{j\ell}'^*\right] e^{-i2\pi jk/d}$$
$$= \left[\gamma_{j\ell} - \gamma'_{j\ell}\right] + \left[\gamma_{j\ell}^* - \gamma_{j\ell}'^*\right] \quad, \tag{4.20}$$

$$\frac{\partial E_\ell}{\partial \gamma_{j\ell}^{\mathrm{I}}} = \frac{i}{\sqrt{d}} \sum_k \left[-\tilde{\gamma}_{k\ell} + \tilde{\gamma}'_{k\ell}\right] e^{i2\pi jk/d} + \left[\tilde{\gamma}_{j\ell}^* - \tilde{\gamma}_{j\ell}'^*\right] e^{-i2\pi jk/d}$$
$$= -i \left[\gamma_{j\ell} - \gamma'_{j\ell}\right] + i \left[\gamma_{j\ell}^* - \gamma_{j\ell}'^*\right] \quad, \tag{4.21}$$

where we have used $\tilde{\gamma}'_{k\ell} = |\tilde{c}_{k\ell}| e^{i \arg \tilde{\gamma}_{k\ell}}$. These results finally lead to

$$\left[\frac{\partial}{\partial \gamma_{j\ell}^{\mathrm{R}}} + i \frac{\partial}{\partial \gamma_{j\ell}^{\mathrm{I}}}\right] E_\ell = \left[\gamma_{j\ell} - \gamma'_{j\ell}\right] + \left[\gamma_{j\ell}^* - \gamma_{j\ell}'^*\right] + \left[\gamma_{j\ell} - \gamma'_{j\ell}\right] - \left[\gamma_{j\ell}^* - \gamma_{j\ell}'^*\right]$$
$$= 2 \left[\gamma_{j\ell} - \gamma'_{j\ell}\right] \quad. \tag{4.22}$$

To better compare it to the update term in (4.11), we note that the latter's component at $|j\rangle$ is

$$\langle j|\phi'\rangle = \epsilon_{j\ell}\beta(\gamma'_{j\ell} - \gamma_{j\ell}),$$

where $\epsilon_{j\ell} = 1$ if $\hat{P}_\ell$ comprehends the dimension generated by $|j\rangle$, and $\epsilon_{j\ell} = 0$ otherwise[3]. Therefore, leaving the factor $\epsilon_{j\ell}$ aside, the update term makes a steepest descent correction (that is, in the direction opposite to the gradient). The factor β allows for controlling the size of the corrections, which can make the algorithm either converge faster (if we use greater values for β) or more slowly (for smaller values). And since each iteration uses a different value of ℓ, the algorithm effectively alternates between the gradients of the error functions (4.13).

The use of several different gradients is necessary to overcome stagnation problems, which could arise if we used a single gradient instead. If the algorithm became stagnated at a local (but not global) minimum of one error function, a later iteration using a different function would move the estimate, and the algorithm would be able to proceed further. Ideally, the only point which should be a common minimum to all the error functions is the correct solution, so the algorithm should only stop there.

---

[3]Or, more formally: $\epsilon_{j\ell} = \langle j| \hat{P}_\ell |j\rangle$.

The factor $\epsilon_{j\ell}$ can be pictured as a Wiener filter [31], which allows greater corrections at points which received greater illumination, and therefore have a better signal-to-noise ratio. Since our projectors will be similar to binary illuminations, where each dimension either is or is not "illuminated", this term reduces to $\epsilon_{j\ell}$.

Now, even though the use of alternating gradients is an interesting solution, it is not itself a guarantee that reconstructions will be successful. In fact, if we want the alternating gradients to work as expected, the algorithm must be fed with data that has enough diversity and redundancy. This is attained with good sets of probe projectors. This means their form, number, rank and even their overlap must be chosen with some criterion. In the next section, we will discuss a few choices we have studied, and later we will see the results of their reconstructions.

4.3 Choice of projectors

We have constructed three families of probe projectors to study the ptychographic protocol outlined above. We considered, initially, the projectors given by

$$\hat{P}_\ell = \sum_{j=0}^{r-1} |j \oplus s_\ell\rangle\langle j \oplus s_\ell| \quad , \tag{4.23}$$

where r is its rank, $\oplus$ denotes addition modulo d and s_ℓ is a nonnegative integer that sets the skip between adjacent operators and may be arranged in a n-entry vector $\mathbf{s}^{(n)} = (s_0, \ldots, s_{n-1})$. This is perhaps the simplest choice of projectors, with all of them being diagonal in the computational basis and encompassing contiguous dimensions of the Hilbert space.

We derived two families of the form given above. In both cases, we used ranks around $d/2$, chosen from a numerical analysis seeking those that optimized PIE's convergence (see appendix A.2). For even dimensions we used $r = d/2$ whereas for the odd ones we alternated between $r = \lfloor d/2 \rfloor$ and $r = \lceil d/2 \rceil$, selecting whichever provided better reconstructions. In general, both ranks gave similar results.

The first family used $n = 4$ projectors about equally spaced along the dimensions of $\mathcal{H}_d$. This family requires $4d$ measurement outcomes to be carried, which is an amount comparable to other works in the literature of pure state reconstruction [18, 20, 21, 37]. In this case, the vector of skips is given by $\mathbf{s}^{(4)} = \left(0, \left\lceil \frac{d-r-2}{3} \right\rceil, 2\left\lceil \frac{d-r-2}{3} \right\rceil, \left\lceil \frac{d}{2} \right\rceil\right)$, which will

give an average overlap of $\mathcal{O} = 2/3$ (see equation (4.2)).

The second family was made of $n = d$ operators with a skip vector $\mathbf{s}^{(d)} = (0, 1, \ldots, d-1)$, which gives an overlap $\mathcal{O} = 1 - 1/r$. This family is overcomplete, as it requires d^2 measurement outcomes to be carried, much more than is required for pure state reconstruction. However, it is capable of reconstructing a class of states that the first family is not, but which forms a zero-measure set in $\mathcal{H}_d$, as we shall see in the next section. Figures 4.3(a) and 4.3(b) illustrate the action of these families of probe projectors on the input states considering $d = 8$.

The third and last family was conceived for the case where the quantum system is comprised of N qubits. In this case, it is desirable to avoid operations that involve more than one qubit at a time, and use only local, one-qubit operations instead. This is because the latter are easier to implement experimentally [38]. We considered, therefore, the set of $n = 6N$ probe projectors given by

$$\hat{P}_{\ell j} = \hat{\pi}_{\ell j} \otimes \hat{\mathbb{I}}^{\otimes N-1} \quad , \tag{4.24}$$

where $\hat{\pi}_{\ell j} = |\ell_j\rangle\langle\ell_j|$ are projectors onto the eigenstates of the Pauli operators $\hat{\sigma}_x$ ($\ell = +, -$), $\hat{\sigma}_y$ ($\ell = R, L$) and $\hat{\sigma}_z$ ($\ell = 0, 1$) of the j-th qubit, and $\hat{\mathbb{I}}$ is the identity in the qubit space. Therefore, $\hat{P}_{\ell j}$ projects the part of the state on the j-th qubit subspace while leaving the remainder unchanged. These probe projectors have rank $r = 2^{N-1}$ and an overlap

$$\mathcal{O} = \frac{1}{2} + \delta_{jj'} \left(|\langle\ell|\ell'\rangle|^2 - \frac{1}{2} \right) \quad ; \tag{4.25}$$

most of them ($\ell \neq 0, 1$), unlike those in equation (4.23), are not diagonal in the computational basis. Thus, they are better visualized by their matrix components in that basis. Figure 4.3(c) illustrates this for a two-qubit system.

It is worth noting that there is plenty of freedom in constructing the family of projectors. The only requirements are those established in the beginning of section 4.1, namely that they must have some overlap and address all levels in $\mathcal{H}_d$. The reader is encouraged to pursue the families that are most-easily implemented in their setup.
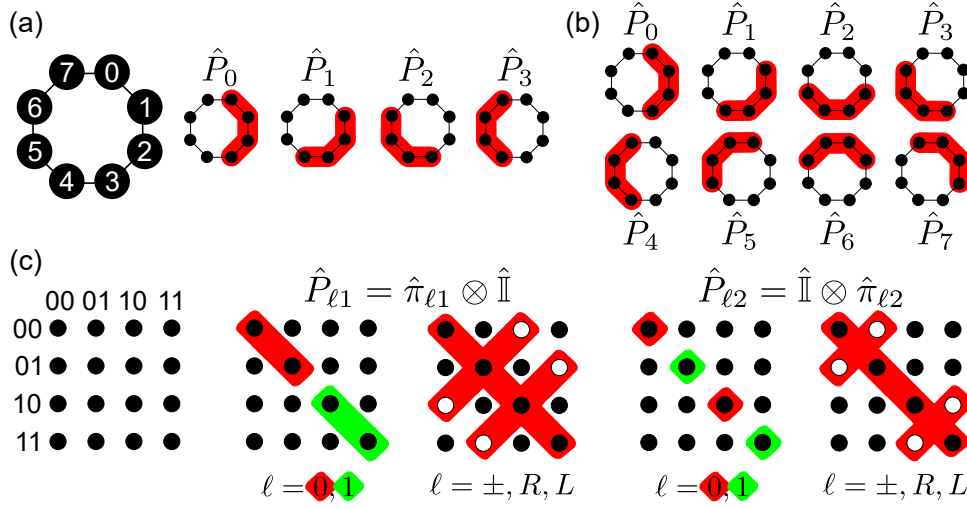
Figure 4.3: **Schematic representation of the probe projectors.** Let us illustrate the projectors of the three families. For the families of the form (4.23), we have: (a) $n = 4$ and (b) $n = 8$ rank-4 projectors, acting as filters in the dimensions of $\mathcal{H}_d$. For the multiqubit family given by equation (4.24), (c) highlights their nonzero matrix components in the computational basis for a two-qubit system; the empty circles indicate the components multiplied by a phase factor that depends on ℓ.

4.4 Blind reconstruction results

We have numerically studied the performance of the protocol when using the three sets of projectors defined in the previous section. We used blind reconstructions to asses its performance: we sampled many pure quantum states randomly, simulated the ptychographic data that it would generate, and fed it to the algorithm; we then calculated the fidelity between the state and its reconstruction, and repeated the procedure for many states. Of course, the algorithm did not have any further information about the quantum state itself (thus the "blind" reconstructions).

Our study comprised several dimensions up to $d = 100$. For each d, 10^4 input quantum states were randomly generated according to the Haar measure[4]. The ptychographic

---

[4]In the set of pure quantum states plus unitary transformations (which are the natural transformations in this set), the Haar measure tells us how to sample pure states "fairly". As a side note, let us clarify that Haar measures exist for other sets and other transformations as well.

To explain what a Haar measure is, we first need to remind ourselves of what a group is. A group is a set equipped with an operation (or several of them) that leaves the set unchanged (think of them as symmetry operations). For example, the sphere and the set of rotations around any axis containing the origin constitute a group.

A Haar measure on a group is a measure (think of it as a way of measuring the size of subsets of that set) that is invariant with respect to the group operation(s). If we wanted a measure that allowed us to fairly sample the elements of the set, it would make sense to demand that it were invariant with respect

data sets corresponding to these states were used to simulate experimental detections realistically (more details next), which were then fed into the PIE algorithm. We chose for our feedback parameter the value $\beta = 1.5$, optimized numerically (see appendix A.1). Our stop criteria consisted of two clauses: reaching $D < 10^{-5}$, or 100 PIE iterations, whichever happened first. If the algorithm was stopped by the second clause, we made it start again with a new random estimate. We allowed the algorithm to reinitiate up to 100 times. Finally, the quality of each reconstruction was quantified by computing the fidelity

$$F = |\langle \phi_{\text{PIE}} | \psi \rangle|^2 \tag{4.26}$$

and infidelity $I = 1 - F$ between the input state, $|\psi\rangle$, and the normalized estimate of the algorithm, $|\phi_{\text{PIE}}\rangle$.

To make the simulations physically realistic, we simulated the imperfect generation of states and the random nature of detections, which are unavoidable sources of uncertainty in actual experiments. These sources of error were modelled as depolarization and Poisson noise, respectively. The first can be modeled as a random fluctuation in the density matrix of the pure state ($|\psi\rangle\langle\psi|$), so that the generated state will be

$$\hat{\rho} = (1 - \eta)|\psi\rangle\langle\psi| + \eta\hat{\rho}_{\text{rand}} \quad , \tag{4.27}$$

where η is the noise level, and the random perturbation $\hat{\rho}_{\text{rand}}$ is drawn according to the Hilbert-Schmidt measure[5] in the mixed-state space [40]. The probe projection followed by

---

to the group operation(s): after all, these operations leave the set unchanged.

The reader might be surprised, but this requirement is enough to find a unique measure (up to a multiplicative constant). Therefore, the Haar measure on a group represents the probability distribution we have to follow if we want to make fair randomizations. This is a very important concern in simulations.

In the case of pure quantum states of a given dimension, the group is formed by the Hilbert space of states plus the unitary transformations.

In practice, random states can be drawn according to the Haar measure by using rescaled gaussian coefficients [39].

[5]The space of mixed quantum states of a given dimension, together with unitary transformations, does not form a group in the usual sense (after all, mixed states will turn into non-states if we apply a unitary transformation alone; we need something of the form $\hat{U}\rho\hat{U}^\dagger$ to keep it a state), and therefore we do not say that there is a Haar measure for them.

However, we could use an approach similar to quantum state purification and resort to a pure state in a higher-dimensional Hilbert space (which could be sampled according to its Haar measure), and take a partial trace to arrive back at the mixed space of interest. This is precisely the approach behind the Hilbert-Schmidt measure.

the QFT will respectively produce

$$\hat{\rho}_\ell = \hat{P}_\ell \hat{\rho} \hat{P}_\ell \tag{4.28}$$

and

$$\hat{\sigma}_\ell = \hat{F}_d \hat{\rho}_\ell \hat{F}_d^{-1} \quad . \tag{4.29}$$

The diagonal components of $\hat{\sigma}_\ell$, denoted by $\{|\widetilde{C}_{k\ell}|^2\}_{k=0}^{d-1}$, provided the simulated data to which we applied a Poisson distribution of average $\lambda|\widetilde{C}_{k\ell}|^2$, where λ is a count rate factor. In our simulations we used $\eta = 0.05$ and $\lambda = 10^3$, in agreement with values found in the literature [8,9], which we also checked to be realistic (see appendix A.3). Note that although the noise will introduce mixedness in the input states, the PIE algorithm will treat the noisy data as if they came from a pure state, and it will output a pure state as estimate.

Figure 4.4 shows the histograms of the infidelity obtained and figure 4.5 shows their averages; overall, the infidelities were lower than 0.01 ($I < 10^{-2}$), which attests that the reconstructions were excellent in general. Only a small fraction of them were not satisfactory. In $d = 100$, for instance, about 4% of input states have not been well characterized (i.e., had fidelities below 0.9) by the family with 4 projectors. A straightforward way to overcome this issue is to include one or a few more probe projectors in the measurement at the expense of increasing the experimental cost.

We observed that the d-projector family of projectors took reasonably longer times than the 4-projector family to reconstruct the states. For example, to reconstruct 10^4 states of dimension $d = 11$ with projectors of rank $r = 6$, the scheme with $n = 4$ projectors took about 800 s, while about 10000 s were needed for $n = d$. For $d = 100$ and $r = 50$, the 10^4 reconstructions took a total of 40 min and 6 h for $n = 4$ and $n = d$, respectively. This suggests that a greater diversity in the data enabled the algorithm to detect more imperfections in the estimate, and consequently more iterations are needed to reach the threshold for D. To make an analogy, picture a carpenter that refines the shape of a wooden piece; by looking at the piece from two orthogonal angles, she is able to correct its shape to some extent. Then, by looking from an angle intermediate to those two, she is able to perceive other imperfections that were not evident and correct them. Therefore, by using more diverse data, the algorithm is able to correct the estimate in a manner that would not be possible by just using more iterations of less diverse data. On the other hand, it could also make more efficient corrections to the estimate, as they will now happen on more subspaces of $\mathcal{H}_d$, so that it is not trivial to foresee how data diversity impacts
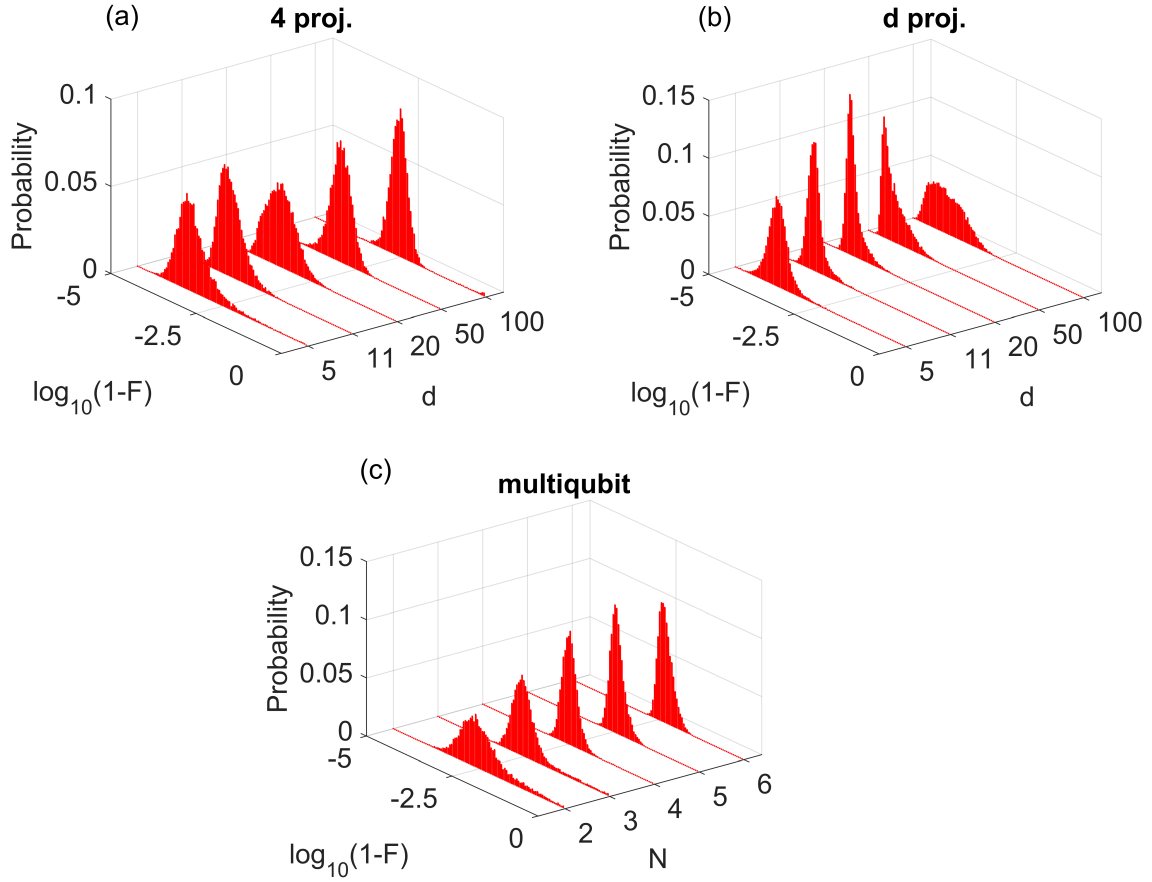
Figure 4.4: **Blind reconstruction results.** Histograms of infidelities for reconstructions using (a) $n = 4$, (b) $n = d$ and (c) multiqubit families of projectors. We used ranks $r = 3, 6, 10, 25, 50$ for dimensions $d = 5, 11, 20, 50, 100$, respectively. The reconstructions were very satisfactory, with infidelities frequently smaller than $10^{-2.5}$. Probabilities were estimated by rescaling the frequencies of the histogram.

convergence.

In general, the algorithm required a few thousands of iterations to converge. However, each iteration is extremely fast since the estimate is just a small complex vector, so the operations that need to be carried out are comparatively simple. Overall, with a modest laptop, the reconstruction of a single state took from fractions of a second, at smaller dimensions, up to a few seconds at greater ones, as can be inferred from the reconstruction times mentioned above.

Figure 4.5: **Average infidelities of reconstructions**. Average infidelities of the reconstructions shown in figure 4.4, using the $n = 4$ (red circles), $n = d$ (black triangles) and multiqubit (cyan squares) families of projectors. The standard deviations 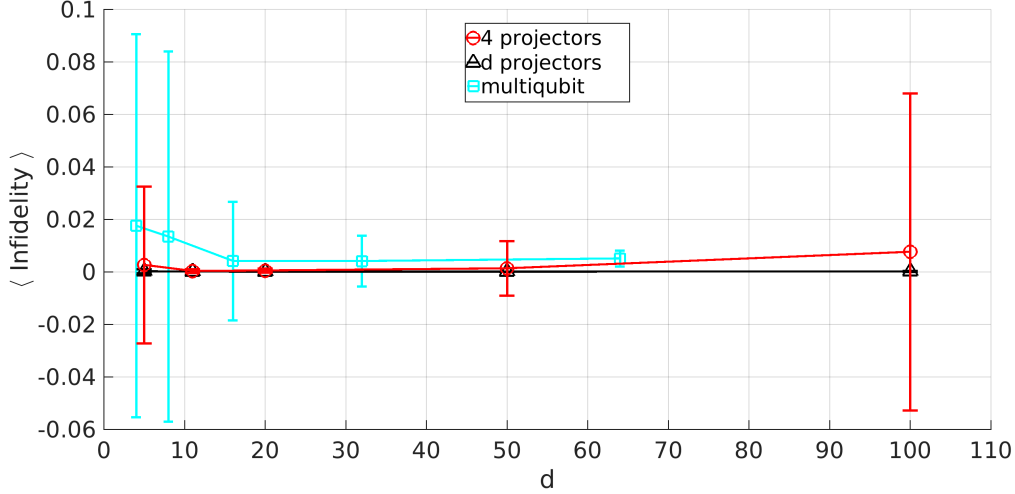are shown as error bars. Overall, the infidelities were very close to zero as desired, especially for the d-projector family.

4.5   Pathological states and non-overlaping projectors

In order to stress the importance of the overlap between probe projectors in the ptychographic measurements, we have also simulated reconstructions using non-overlapping $\hat{P}_\ell$'s. For $d = 20$, $n = 4$, $r = 5$, $\mathbf{s}^{(4)} = (0, 5, 10, 15)$ in equation (4.23), and 10^4 random states, we obtained fidelities ranging from 10^{-4} to 0.81 with an average of 0.15. These results are essentially random, and show that the multiple overlaps are crucial in the protocol: without them, the ptychographic problem becomes several disjoint standard phase retrieval problems [27], which are known to suffer from non-uniqueness and stagnation issues [29] as we have discussed in section 3.2.

To appreciate the difficulty that this would pose, the reader can picture the reconstruction of a four-dimensional state, in which we used the non-overlapping projectors $\hat{P}_0 = |0\rangle\langle 0| + |1\rangle\langle 1|$ and $\hat{P}_1 = |2\rangle\langle 2| + |3\rangle\langle 3|$. We would have, effectively, two disjoint two-dimensional retrieval problems. After reconstructing the relative phases between the components in each of the two two-dimensional subspaces, we would still have no information about their "global" phases (or, better said, their phases relative to components outside of the two-dimensional subspaces). Therefore, there would not be enough information to reconstruct the state. The disjoint nature of the sub-reconstructions

would hinder us from reconstructing the states.

The ptychographic method, with the three projector families used so far, will suffer difficulties in the reconstruction of sparse pure states, i.e., states in which most of the components are zero. It succeeds only if every nonzero component was addressed with at least one other nonzero component by some probe projector, so that they could interfere. This is hindered, in general, when the state is sparse. The result is that the algorithm falls into a disjoint group of phase retrieval problems, as discussed above. One way to avoid this is to use the already considered set of $n = d$ probe projectors (equation (4.23) with $\mathbf{s}^{(d)} = (0, 1, \ldots, d-1)$), but now with rank $r \geq \lfloor d/2 \rfloor + 1$, because when the levels of $\mathcal{H}_d$ are addressed cyclically, the biggest distance between nonzero components will be $\lfloor d/2 \rfloor$. A second way is to use an adaptive approach: first, one measures in the computational basis; if the state is verified to be sparse, then one applies the ptychographic method building the probe projectors according to the distribution of its nonzero components. Such an extra step is, in general, easy to carry out.

4.6   Comparison with other methods

4.6.1   Number of measurement outcomes

The ptychographic method introduced here requires a total number of $\mathcal{M} = nd$ measurement outcomes—d QFT state-amplitudes for each of the n probe projections. In this regard, its experimental cost will be determined by the number of $\hat{P}_\ell$'s adopted. Along with the specific form of the projectors, this number also defines the diversity of the ptychographic data set and its degree of redundancy arising from the partially overlapping projections. When choosing n, one should be aware that a high value, although experimentally more demanding, provides more diversity and redundancy in the data set. This has consequences on the quality of the reconstructions and on the convergence of the algorithm, and our results revealed that there can be appreciable differences between projector families.

We have investigated three specific families of projectors, which were described in section 4.3. The family with 4 projectors entails $\mathcal{M} = 4d$ measurement outcomes, which is slightly less than other methods in the literature (see section 2.2), but does not work for a zero-measured set of states, just as those methods. The family with d projectors, on the other hand, does not have such restriction, but requires $\mathcal{M} = d^2$ measurement

outcomes, which is more expensive than other methods in the literature. Finally, the qubit family uses $\mathcal{M} = 6d \log_2 d$ measurement outcomes, also more expensive than other reported methods[6]. We summarize a comparison of these three ptychographic families and the pure-state methods we have seen in table 4.1.

Table 4.1: **Comparison between pure-state reconstruction methods.** Here we summarize the main features and limitations of the pure state reconstruction methods we have seen so far, both from works in the literature (subsection 2.2.2) and the three ptychographic families we have proposed.

Method	Gell-Mann basis	Phase stepping	Orthogonal pols. SDP	Ptychography, 4 projectors	Ptychography, Multi-qubit projectors	Ptychography, d projectors
Number of outcomes	$5d$	$4d$	$5d$	$4d$	$6d \log_2 d$	d^2
Limitations	Does not work for a zero-measured set of states	Requires an extra dimension for encoding	–	Does not work for a zero-measured set of states	Only for multiqubit systems	–

Once again, we remind the reader that there is ample freedom to choose the ptychographic projectors. Perhaps by making smart choices, the reader can arrive at a family that beats all the others in the comparison table.

4.6.2 Experimental implementation of quantum state ptychography

To illustrate the simplicity of the quantum ptychographic scheme, consider d-dimensional states encoded in the propagation modes of single photons (or any other type of radiation). A multiport interferometer (MI), sketched in the right box of figure 4.6 for $d = 8$, can implement any unitary transformation on this encoding [41]. Under these circumstances, the probe projectors given by equation (4.23) would be realized by mode filters at the input ports of the interferometer, as shown in the left box of figure 4.6. By setting the MI to perform $\hat{F}_d$, the ptychographic data would be collected simply by shifting the mode filters n times at the input ports and recording the counts at the output ports. For comparison, to reconstruct these states by using the other methods in the literature [18, 20, 21], the mode filters would not be necessary, but one would have to reconfigure the whole MI for

---

[6]We created this projector family by considering all possible projectors in equation (4.24). It is possible that only a fraction of those projectors would be enough, which would lower the number of requiered measurment outcomes. This will remain a topic for further study, though.

each measurement basis employed. This shows a nice feature of the ptychographic method: the measurements are effectively performed in a single basis while the probe projectors are "shifted" through the Hilbert space.
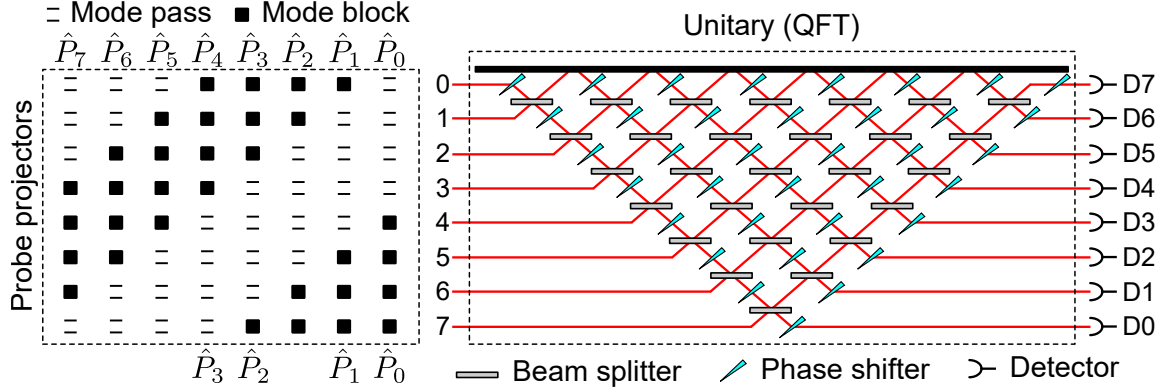

Figure 4.6: **Ptychographic reconstruction in a multiport interferometer.** Scheme for ptychographic reconstruction of 8-dimensional quantum states in a multiport interferometer. The rank-4 probe projectors given in equation (4.23) are implemented via mode filters: (top) $n = 8$; (bottom) $n = 4$.

When considering the reconstruction problem, it is natural to think that implementing several unitary transformations demands an experimental setup that changes as a whole. The ptychographic approach shows, however, that most of the setup can be left unchanged, and only the simplest parts (corresponding to the projectors) need to change. This is a considerable simplification.

From the experimental point of view, the implementation of quantum state ptychography should be straightforward. On one hand, the QFT has been experimentally realized in many different platforms including trapped ions [42,43], superconducting qubits [44], nuclear magnetic resonance [45], neutral molecules [46], and photonic systems [47–51], both for single-particle and multiqubit scenarios. On the other hand, the families of probe projectors proposed here and given by equations (4.23) and (4.24) are both simple to implement in these scenarios.

4.7   Chapter summary

In this chapter, we have presented the quantum version of ptychography. We showed how its iterative reconstruction algorithm can be formulated as an alternating-gradient search, which clarifies how the protocol overcomes stagnation and non-uniqueness problems that

standard phase retrieval methods have. We also showed three possible projectors sets, each with their interesting features, and studied their performance in blind reconstructions. Overall, their performance was excellent, even with substantial amounts of noise included in the simulations. Finally, we compared the ptychographic approach to other methods found in the literature and stressed its advantages, most importantly the simpler experimental setup it requires.

Chapter 5

Pure quantum state ptychography: experimental results

After investigating quantum ptychography numerically, it was time to prove it in the laboratory. We performed a blind reconstruction experiment [52] where we generated slit qudits, acquired ptychographic datasets, and compared the ptychographic reconstructions to the original states. We shall detail our apparatus, the experiment, and its results in this chapter.

5.1   Slit qudits

Mathematically, a pure qudit is a vector in a d-dimensional, complex Hilbert space. Experimentally, there are many physical quantities that are naturally represented as qudits, and thus can serve as experimental platforms for pure quantum states. In our laboratory, we use the transverse profile of a highly-coherent laser beam to encode d complex numbers, in the manner that we will detail next.

Our experimental apparatus is capable of controlling the transverse profile of a laser beam. We use this capability to give the beam a shape of multiple rectangular slits, its intensity and phase being constant within each slit, and independent between different slits. Therefore, if we give the beam the shape of d slits, it will encode d complex numbers. Figure 5.1 illustrates this kind of beam profile. In the next section, we will discuss how to control the number of slits, its amplitudes and phases.

If the slit regions have a height of Y and a width of X, and considering the laser
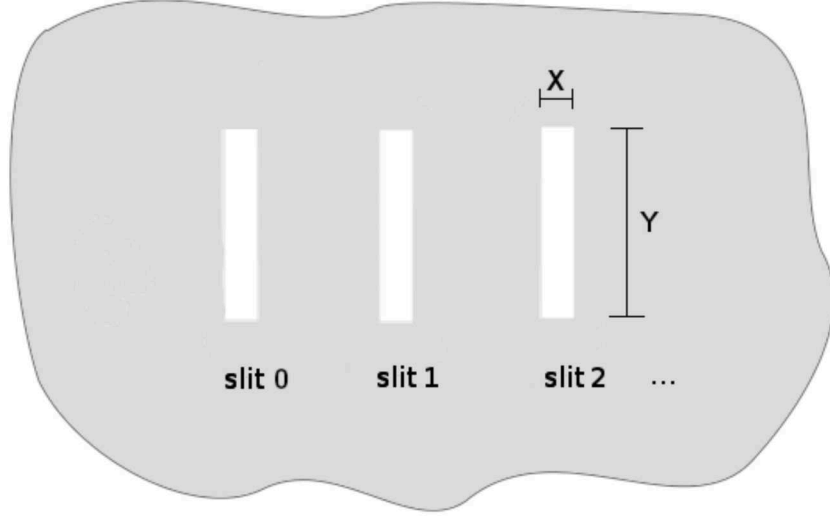
Figure 5.1: **Transverse profile of the beam in our apparatus.** Our laser beam had the profile of several slits. By controlling the phases and amplitudes of the beam in each slit, we are able to make it encode several complex numbers.

beam as a monochromatic plane-wave source, we can express the transverse profile of the j-th slit as

$$E_j(x, y, t) = E_0 e^{-i\omega t} \text{rect}\left(\frac{x - x_j}{X}\right) \text{rect}\left(\frac{y}{Y}\right) \quad , \tag{5.1}$$

where ω is its angular frequency, x_j the slit's central position along the x axis, E_0 the field's amplitude, and

$$\text{rect}(\eta) = \begin{cases} 1, & \text{if } \eta \in [-0.5, 0.5] \\ 0, & \text{otherwise} \end{cases} \tag{5.2}$$

is the rectangle function. A field with amplitudes a_j and relative phases ϕ_j at each slit can therefore be written as

$$E(x, y, t) = \sum_{j=0}^{d-1} a_j e^{i\phi_j} E_j(x, y, t) \quad . \tag{5.3}$$

The slit fields E_j therefore play the role of an orthogonal basis[1] for this class of fields.

We can use slit fields to represent finite-dimensional, pure states of a quantum system[2]: to represent a state $|\psi\rangle = \sum_{j=0}^{d-1} c_j |j\rangle$, we use the the field $E(x, y, t) =$

---

[1]It is easy to check that their inner product is $\langle E_l, E_m \rangle = E_0^2 \delta_{lm}$, since distinct slits will have no intersection.

[2]More formally, it is possible to show that the set of slit fields constitutes a vectorial space with internal product, dimension equal to the number of slits, and which is complete. Therefore, if we use

$\sum_{j=0}^{d-1} |c_j| e^{i \arg c_j} E_j(x, y, t)$. Hence we can use classical optics to imitate quantum states.

This analogy, however, does have its downsides: if we want to generate the analogous of an entangled state while still using a single laser source (as we have, in our laboratory), we have to generate a representation (via slit fields) of the global state of the system[3]. If we wanted to represent a system composed of many parts (for example, many qubits), the number of required slits would grow exponentially with the number of parts.

Finally, it is important to note that this analogy between fields and states is only complete if we are able to apply operators to the first in the same manner we would to the second. In the next sections we will discuss how we can generate slit qudits in our experimental setup and how to implement projective measurements on them, which are the operations we will need for the ptychographic protocol.

5.2   Experimental implementation of slit qudits

The main component of our setup is the *spatial light modulator* (SLM). In our case, this component is a small *liquid crystal display* (LCD). To understand its functioning, we have to keep in mind that its liquid material is made of molecules which are elongated and cylindrical, with a shape resembling a cigar. The liquid is contained between two sheets of glass, and behind one of those is a silicon substrate with reflective coating[4]. The substrate is divided in several pixels that can act as capacitors and generate electrical fields inside the region filled with the liquid. The internal facets of the glass sheets are both polished in the same direction; this favors the liquid's molecules to align themselves, throughout the entire volume, along the polishing direction. The structure of this type of modulator is shown in figure 5.2.

When turned on, the substrate pixels generate an electrical field in the volume occupied by the liquid crystal, and is perpendicular to the glass sheets. Consequently, they stand under two competing tendencies that affect their alignment: the first being the tendency to align themselves parallel to polishing of the glass sheets, and the second to the perpendicular electric field. These two situations are illustrated in figure 5.2. The

---

d slits, the set of slit fields forms a d-dimensional Hilbert space. Finally, if we remember that any two Hilbert spaces of the same dimension are isomorphous, the analogy between slit fields and pure quantum states is formalized.

[3]This is the same problem that a classical computer has when trying to represent an entangled state.

[4]Our SLM is of the reflective type. There are, however, transmissive SLMs, which do not have the reflective coating at the backside.
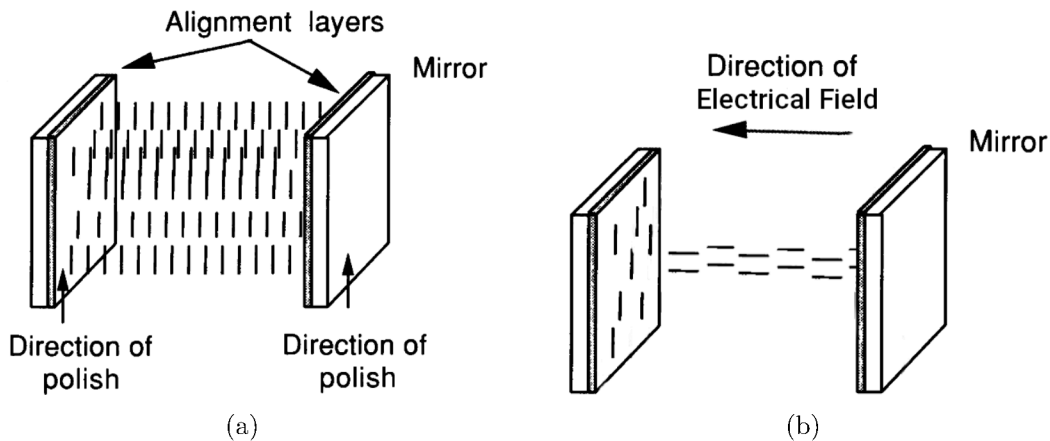
Figure 5.2: **Structure of the spatial light modulator.** When the electric field is off (a), the nematic molecules tend to align themselves with the polishing direction of the glass inner facets. When it is on (b), they tend to align themselves with the field. Figures taken and edited from [3].

net result is that the electric field deviates them from their resting alignment, making the molecules align obliquely to the glass sheets and the silicon substrate. Furthermore, the refraction index of the (aligned) liquid crystal is anisotropic, so light propagating along the alignment axis will suffer a delay distinct to light propagating transversely to it. Therefore, controlling the alignment of the molecule in the crystal, which can be done through the voltage at the pixels at their base, leads to interesting optical applications [3].

This type of SLM is projected to be illuminated by a (uniform and coherent) laser beam that propagates perpendicularly to the glass sheets, with polarization parallel to the polishing. The beam will propagate through the liquid crystal, then be reflected by the coating at the back glass sheet, then propagate through the liquid crystal once again and finally emerge from the SLM. Each part of its transverse profile will experience, however, different refraction indexes, depending on the voltage that is set at each pixel. We can impart, therefore, transverse phase profiles to the beam. In other words, the SLM acts as a programmable diffractive element. This has been used to generate slit qudits in [53]. Let us see briefly how the process works.

The SLM is operated with greyscale images that a computer sends to the SLM's controller, which performs the task of translating those images into pixel voltages applied to each substrate pixel. The grey level of each pixel correlates to a phase shift. These images are thus called *phase masks*. In our experiments, we use periodic patterns (which will act as diffraction gratings) in the regions of the beam profile where we want slits, and leave the remaining regions with a zero phase shift. A typical phase mask is shown

in figure 5.3(a). We can imagine, with a little simplification, that the part of the beam that impinges on the patterns is diffracted to several orders, while the other parts are just reflected back and mix with the zero-order diffraction, as shown in figure 5.3(b). Our setup filters, at the focal plane of a spherical lense, the light that is diffracted to the first order and blocks the others (figure 5.3(c)).
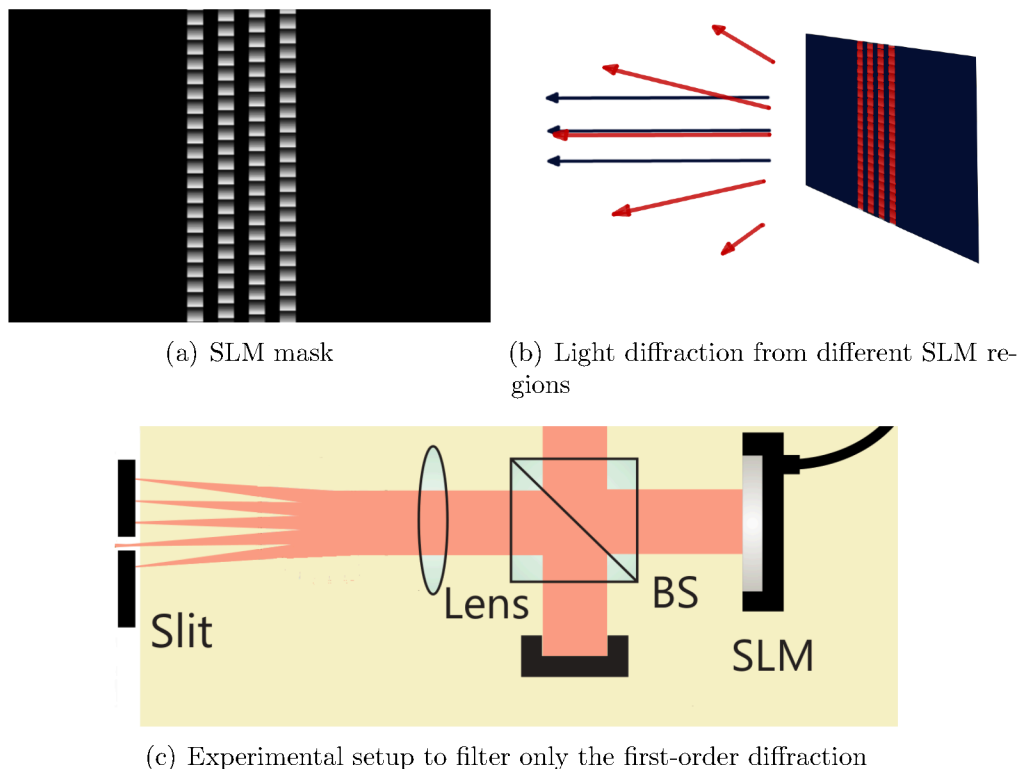

(a) SLM mask

(b) Light diffraction from different SLM regions

(c) Experimental setup to filter only the first-order diffraction

Figure 5.3: **Diffractive masks for our experiment.** (a) We use diffractive gratings in the regions we wish to create slits. (b) The technique consists of diffracting the light from these regions to several orders, while the light from the other regions is diffracted only to the zeroth order. (c) By using a slit to block all diffractive orders but the first, we are effectively allowing only the light from the slit regions to continue further in the optical apparatus.

After the spatial filter, we use a lens to make images of the slits (figure 5.4(a)). It is possible to control the number of slits, as well as their phases and amplitudes, by using the appropriate phase mask. The number of slits can be set with the number of rectangular regions with diffractive patterns. Their amplitude is controlled by the diffractive pattern profile, and their phase by a constant grey level that is added to the pattern[5]. Figure 5.4(b)

---

[5]Actually, our setup uses a translation of the diffraction gratings. While mathematically this is

shows the phase mask for a four-slit field.
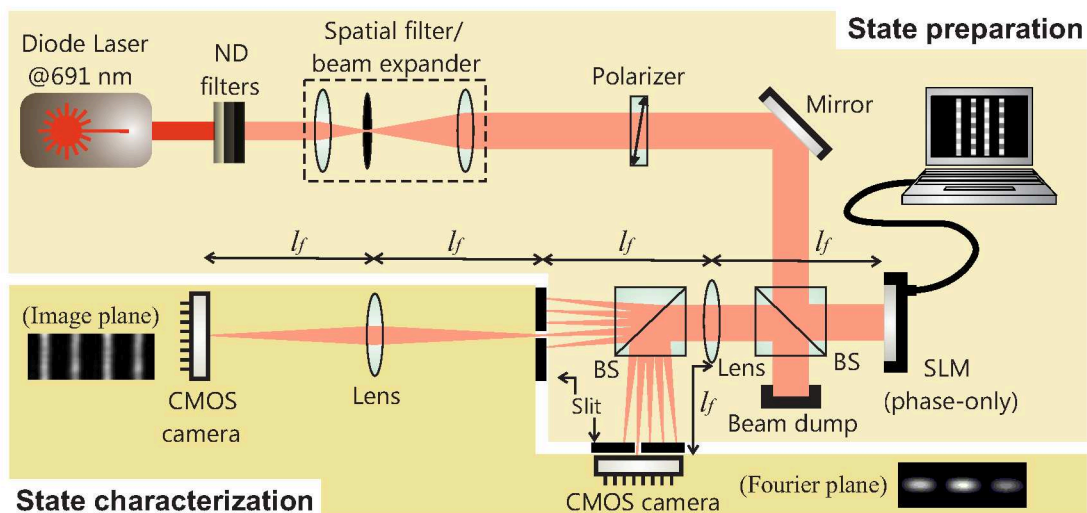
(a) Setup for our experiment
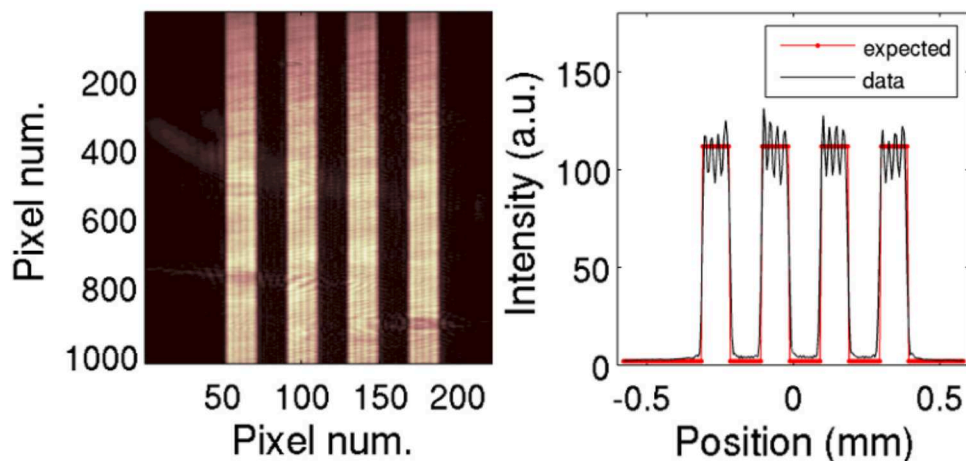
(b) Image and intensity profile of a slit field generated in our setup

Figure 5.4: **Full experimental setup.** These figures show (a) the full experimental setup we used for our experiment, and (b) a typical slit field that can be generated with it.

After preparing the slit qudits, we need to operate on them. We will see how in the next section.

equivalent to adding a constant phase, since the phase space is cyclic, it is more feasible experimentally.

5.3   Projective measurements

When using a laser beam in an optical system, the natural kind of measurement that comes to mind are intensity measurements. In this section, we shall formulate measurements at specific points of the setup as projective measurements on the slit qudits. We will follow the discussion in [54].

In our optical setup, the only relevant variables are the horizontal and longitudinal positions, which we will denote by x and z respectively. For the sake of simplicity, we will omit other variables from our notation (time and beam frequency, for example). In the formalism of quantum field theory [55], the electric field can be described by means of the electric-field operator $\hat{E}(x, z)$. It turns out that this operator can be separated into its positive- and negative-frequency parts,

$$\hat{E}(x, z) = \hat{E}^{(+)}(x, z) + \hat{E}^{(-)}(x, z) \quad , \tag{5.4}$$

and that these parts comprise only destruction and creation operators, respectively, that can be applied to the vacuum state. The state of the field at (x, z) is given by

$$|E(x, z)\rangle = \hat{E}^{(-)}(x, z) |\text{vac}\rangle \quad , \tag{5.5}$$

where $|\text{vac}\rangle$ is the vacuum state of the field, and the creation operator $\hat{E}^{(-)}$ carries information about the optical setup. When we carry an intensity measurement at the point (x, z), we are projecting the field onto the state $|E(x, z)\rangle$ described above.

The next step is to introduce the slit states into this formalism. We can express the j-th slit field (5.1) at the plane $z = 0$ in terms of plane-wave fields:

$$E_j(x, 0) = E_0 \sqrt{\frac{X}{\pi}} \int_{-\infty}^{\infty} e^{-i q x_j} \text{sinc}\,(Xq)\, dq \quad , \tag{5.6}$$

where q denotes the transverse momentum of the plane waves. Now, we can make the analogy

$$E_0 dq \to |N_0, q\rangle\, dq = \frac{1}{\sqrt{N_0!}} \left[\hat{a}^\dagger(q)\right]^{N_0} |\text{vac}\rangle\, dq \tag{5.7}$$

between classical and quantum plane-wave fields, where N_0 is the number of photons that would correspond to the electric field amplitude E_0, and $\hat{a}^\dagger(q)$ is the creation operator for

the plane wave mode. We then arrive at the following expression for the slit fields:

$$|j\rangle = \sqrt{\frac{X}{\pi}} \int_{-\infty}^{\infty} e^{-iqx_j} \mathrm{sinc}\,(Xq)\,|N_0, q\rangle\, dq \quad . \tag{5.8}$$

We are now able to express the restriction $|m(x,z)\rangle$ of an electric-field state $|E(x,z)\rangle$ to the slit-qudit space:

$$|m(x,z)\rangle = \sum_{l=0}^{d-1} |l\rangle\,\langle l|E(x,z)\rangle \tag{5.9}$$

$$= \sum_{l=0}^{d-1} |l\rangle\,\langle l|\,\hat{E}^{(-)}(x,z)\,|\mathrm{vac}\rangle \tag{5.10}$$

$$= \sum_{l=0}^{d-1} \phi_l(x,z)\,|l\rangle \quad , \tag{5.11}$$

where $\phi_l(x,z) = \langle l|\,\hat{E}^{(-)}(x,z)\,|\mathrm{vac}\rangle$. In the first step, we used the fact that the identity operator in the slit-state space is $\hat{\mathbf{1}}_d = \sum_{l=0}^{d-1} |l\rangle\,\langle l|$.

Finally, let us consider an optical system consisting of a converging spherical lens of focal length f, and which uses a laser source of spatial frequency k. Also, let s be the horizontal distance between two consecutive slits. It is possible to show[6] [54] that taking measurements in the focal plane at the horizontal positions

$$x_p = -2\pi \frac{f\mu_p}{ksd} \quad , \tag{5.12}$$

with p taking the values $p = 0, 1, \ldots, d-1$ and

$$\mu_p = \begin{cases} p, \text{if } p \le d/2 \\ p - d, \text{if } p > d/2 \end{cases} \quad , \tag{5.13}$$

corresponds to taking projective measurements at the states

$$|m(x_p, f)\rangle = \frac{1}{\sqrt{d}} \sum_{l=0}^{d-1} e^{i2\pi\mu_p l/d}\,|l\rangle \quad , \tag{5.14}$$

which are the Fourier-basis states. These are indeed the projective measurements that we

---

[6]One needs to find an expression for $\hat{E}^{(-)}(x,f)$ in this system. This can be done by starting from the classical field, and then using the same quantization analogy we used earlier.

want to take in order to carry out the ptychographic protocol.

5.4   Experimental setup

We now understand how to accomplish each task in the ptychographic protocol in the context of our experiment. Evoking once more the "three constituents" of ptychography, we have discussed slit qudits and measurements in certain positions of the focal plane, which will play the roles of object and Fourier-basis measurements respectively. The post-processing algorithm, our third and final "part", will remain unchanged. We are ready to see how the three fit together in our experiment. Let us begin by discussing the concrete setup, which is shown in figure 5.5.

We used a single-mode diode laser at $\lambda = 687\,\mathrm{nm}$, whose beam profile was spatially filtered, expanded and collimated. The treated beam would then illuminate our SLM (Holoeye PLUTO) approximately as a uniform-intensity, normally-incident, plane wave with vertical polarization. The phase of the reflected field was modulated by a computer-generated mask addressed at the SLM given by an array of blazed diffraction gratings. Typical masks are shown in the insets of figure 5.6. For a display with pixels $8\,\mu\mathrm{m}$ wide, the gratings have a period of 12 pixels, width and separation of 11 pixels for $d < 20$, or 9 and 5 pixels, respectively, for $d \geq 20$. The modulated field was transmitted through a spherical lens (L_1) and, at its focal plane, the first diffracted order was filtered by a slit diaphragm at both output arms of a beam splitter (BS). The filtered field emerged as a coherent superposition of d slit fields (equation (5.3)). The magnitude and phase of each mode were controlled by the phase depth of the grating and its relative lateral displacement, respectively [53,56]. The high level of purity of such states have been characterized in many other experiments, e.g., [51,53,57–59]. CMOS cameras (Thorlabs DCC1545M) at the transmitted and reflected arms record the far-field and near-field intensity distributions, respectively.

The insets of figure 5.5 show the role of each of the three modules (separated by boxes in the figure) in our setup. The first step emulated the preparation of the projected states ($\hat{P}_\ell|\psi\rangle$). The projectors $\hat{P}_\ell$ act as binary filters in $\mathcal{H}_d$, as each level either is or is not selected, as seen in figure 5.7. This corresponded to applying binary filters to the mask that generated $|\psi\rangle$. In the other two modules the spatial intensity distributions, associated with $|\psi_\ell\rangle$ and recorded by the cameras, were used to obtain the outcomes of the projective measurements in the computational (near-field) and Fourier basis (far-field). Note that
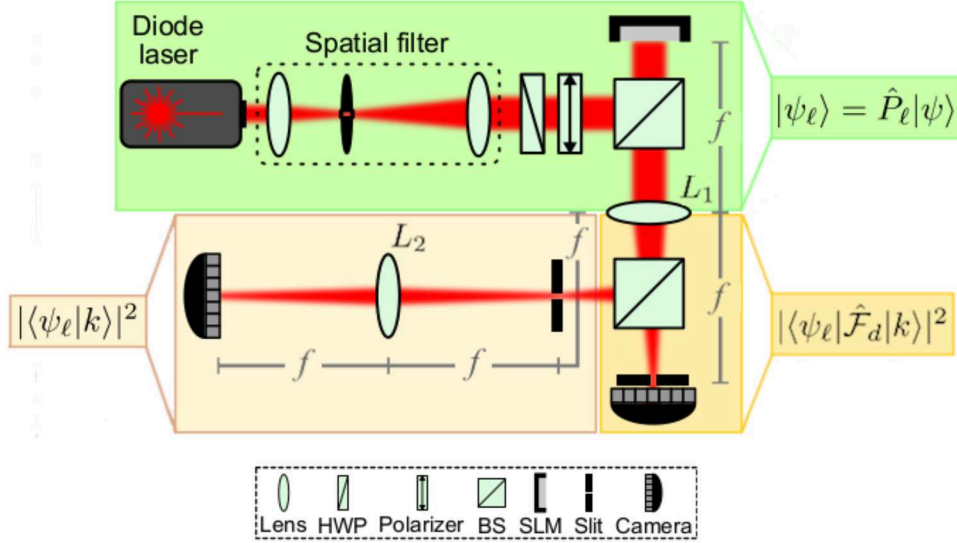
Figure 5.5: **Stages of the experimental setup**. This figure shows the role of each stage (shaded regions) in our setup regarding the ptychographic protocol. Lenses L_1 and L_2 have focal length $f = 30\,\mathrm{cm}$. HWP: half-wave plate; BS: beam splitter; SLM: spatial light modulator.

the lens L_1, shared between two modules, assisted both preparation and measurement stages, simplifying our setup [51].

To illustrate the results of the setup, figure 5.6 shows the (normalized) one-dimensional intensity distributions measured for a $d = 6$ dimensional state. The near- and far-field measurements (red points) are plotted in the first and second rows, respectively; the solid black curves are the theoretical predictions and the insets show the mask addressed to the SLM in each case. In figure 5.6(a), we plot the results for the target state $|\psi\rangle$. Figures 5.6(b) to 5.6(e) show the results for $|\psi_0\rangle$ to $|\psi_5\rangle$, respectively. The red circles in the far-field graphs indicate the transverse positions given by equation (5.12) where we took the ptychographic data.

Resorting to the n probe projectors defined in equation (4.23) and fully described after that, we used two families of that form with the following specifications: (i) $n = 5$ and $s_\ell = \ell\lfloor d/5 \rfloor$; (ii) $n = d$ and $s_\ell = \ell$. For both, $r = \lceil d/2 \rceil$. Figure 5.7 illustrates their action on 6-dimensional states. Family (i) uses $5d$ measurement outcomes, which is comparable to other reconstruction techniques [21, 37, 57–59]; family (ii) requires d^2 outcomes, which is an overcomplete dataset to reconstruct pure states[7].

---

[7]We have also investigated a family that uses 4 projectors, similar to that proposed in chapter 4. The results were generally good, but we were not satisfied with their consistency as bad reconstructions were happening with some frequency. We still need to uncover what could have caused this performance
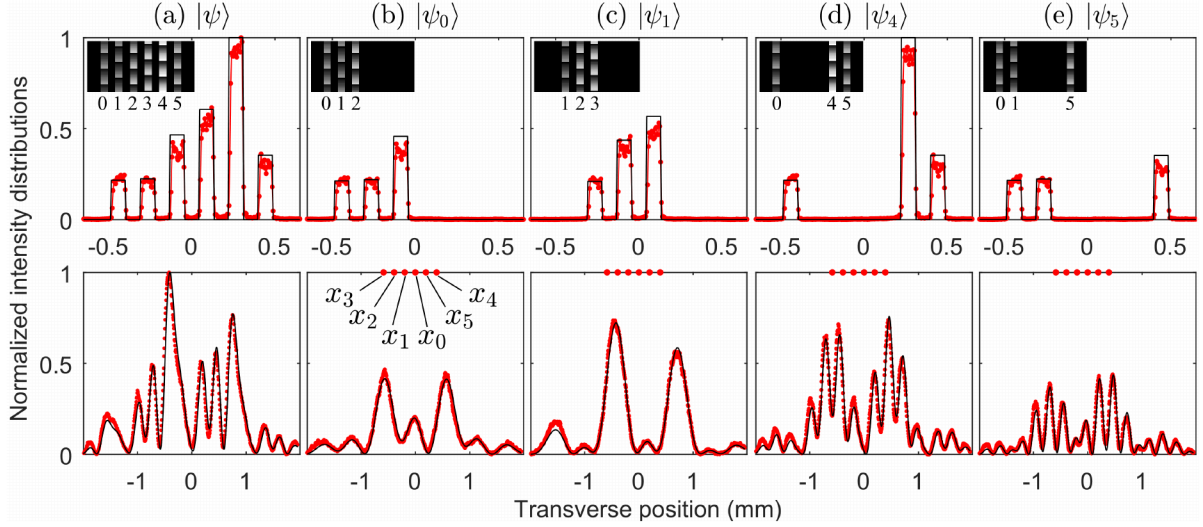
Figure 5.6: **Near- and far-field intensities of a typical state.** Normalized near- (first row) and far-field (second row) intensity distributions for a 6-dimensional target state $|\psi\rangle$ (a) and the corresponding projected states $|\psi_\ell\rangle$ (b)–(e). Experiment (red points); theory (solid lines). The insets show the mask for $|\psi\rangle$ and the "filtered" masks for each $|\psi_\ell\rangle$. In the far-field plots, the red circles indicate the positions given by eq. (5.12) where the ptychographic data are taken. Normalizations use the maximum intensities generated by (a) $|\psi\rangle$ and (b)–(e) $|\psi_\ell\rangle$ (in this case, $|\psi_2\rangle$).
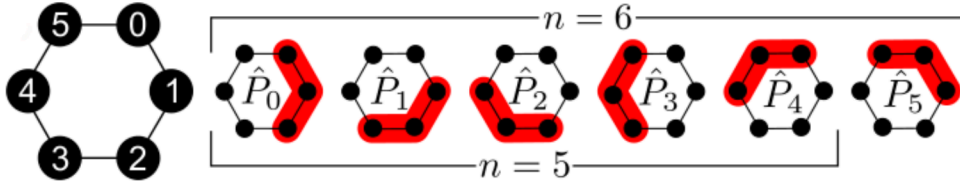

Figure 5.7: **Illustration of the projector families.** Schematic representation of the projectors $\hat{P}_\ell$ for $d = 6$ and $r = 3$. The red shade highlights the selected levels in $\mathcal{H}_d$. Bottom: family (i); top: family (ii).

We performed the ptychographic protocol for several dimensions from $d = 3$ to $d = 32$. For each d we reconstructed a number of pure states that were randomly generated according to the Haar measure. Given a target state $|\psi\rangle$, we built its greyscale mask M_ψ, and the corresponding $n = d$ "filtered" masks $\{M_{\psi_\ell} \to |\psi_\ell\rangle\}_{\ell=0}^{d-1}$, following the prescriptions for the family (ii) of projectors. Each mask was addressed to the SLM one at a time, and we recorded both near- and far-field images. In both cases, we took three images per mask, averaged over them, integrated over the transverse direction y, and subtracted the background noise, obtaining one-dimensional patterns as seen in figure 5.6. All far-field

---

shortcoming, and decided to use one more projector in our experiment.

images were taken within the region between the first two single-slit diffraction minima.

To characterize and check our state preparation, we used least-squares fitting on both the near- and far-field data. This procedure worked as a trustworthy reconstruction, which would take into account our experimental imperfections. The retrieved source states, $|\psi_{\mathrm{src}}\rangle$, served as ground truths to which we later compared the ptychographic reconstructions.

In the ptychographic reconstructions, however, we *only* used the postprocessed far-field images for each M_{ψ_ℓ}. In the ℓ-th image we took *only* the intensities $I_{\ell j}$ at the d pixels in the positions x_j given by equation (5.12) (see also figure 5.6). Gathering these intensities, the ptychographic data are settled as $\{\Pi_\ell = \{I_{\ell j}^{1/2}\}_{j=0}^{d-1}\}_{\ell=0}^{d-1}$. To perform the reconstruction with the family (ii) of projectors, we use all Π_ℓ. To do the same with the family (i), we select the subset of five Π_ℓ associated with the M_{ψ_ℓ} which accomplish the prescriptions for that family. In either case, the data is fed into the PIE algorithm (section 4.1) and its normalized estimate, $|\phi_{\mathrm{PIE}}\rangle$, is used to calculate the fidelity of the ptychographic reconstruction as $F = |\langle\phi_{\mathrm{PIE}}|\psi_{\mathrm{src}}\rangle|^2$.

5.5 Results and discussion

For $d \leq 10$, $11 \leq d \leq 17$, and $d > 17$, we have reconstructed 100, 50, and 13 random states per dimension, respectively.

The fidelities we obtained were collected and their distribution is shown in figure 5.8. The distributions in grey correspond to the $n = 5$ projector families, while those in red to the $n = d$ families. For a better visualization, we scaled the width of each violin by area, i.e., they all have the same area, regardless of the number of reconstructions. The horizontal solid line cutting each violin represents the average fidelity.

Since $F \in [0, 1]$, with $F = 1$ characterizing a perfect reconstruction, we can see that the fidelities achieved for all dimensions and both families of projectors were consistently high. This means all states prepared by the source were faithfully reconstructed by the pytchographic method.

Considering only the state-space dimension, figure 5.8 shows that, in general, the fidelities decrease as d increases. This trend is related with the growing experimental imperfections in the preparation and measurement stages. An increasing d decreases the purity of the source states since the coherence between the transverse spatial modes is
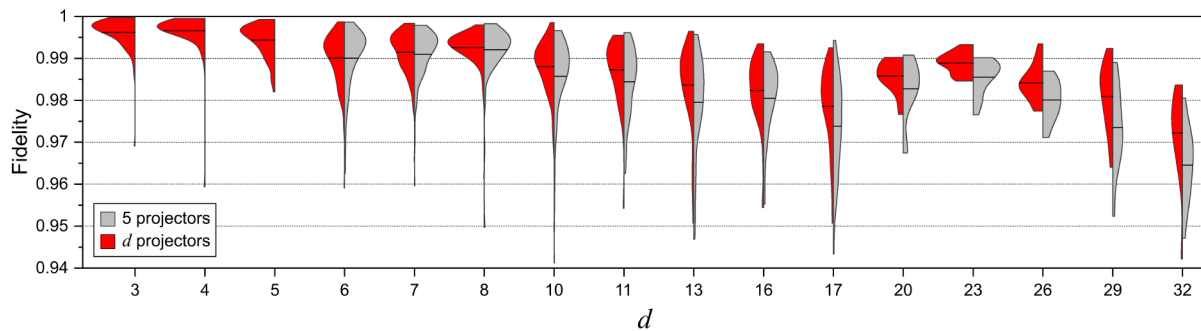
Figure 5.8: **Reconstruction fidelities for several dimensions.** Split violin plots with the distribution of fidelities as a function of d for the ptychographic reconstructions using $n = 5$ and $n = d$ projectors. The horizontal solid lines represent the average fidelities.

affected. Note that for $d = 20$ the trend is broken and restarted at a higher level because, as described earlier, we reduced the width and separation of the spatial modes (for $d \geq 20$), thus reducing decoherence effects. At the measurement stage, the errors in the detection process grow with the number of Fourier basis states, d. This remark on decoherence concerns, of course, our particular implementation and not the ptychographic method in general.

Let us now compare the results of the two projector families. As seen in figure 5.8 for $d \geq 6$, the dispersion of both fidelity distributions for each d are similar, while the averages for $n = d$ are greater than for $n = 5$, which becomes more evident as d increases. This shows that, in general, using more projectors leads to better reconstructions since the ptychographic datasets will have more diversity and redundancy, which are key elements for the method, as discussed in chapter 4. On the other hand, using less projectors implies in less measurement settings and, as we observed here, provides reconstructions with comparable quality.

As for the postprocessing times, t_{PIE}, for $n = 5$ the algorithm took few milliseconds per state for all d; for $n = d$, the times scaled roughly as $t_{\mathrm{PIE}} \approx d^3$, but did not last more than 230 milliseconds for $d = 32$ on a modest laptop.

We have shown that the present form of quantum ptychography is resilient to noise and works quite well for state purities up to $\sim 90\%$. Below that, in general, the reconstruction algorithm would probably fail to converge. This is expected, since one of its premises is the state being pure. Thus, our results confirm the high level of purity of the source states.

5.6   Chapter summary

We performed a proof of concept demonstration of quantum state ptychography. The method faithfully reconstructed pure states in dimensions up to 32, and it was shown to be simple and flexible in regard to the measurement settings, and robust to noise. Although we used a simplified classical optical setup, similar results could be achieved in a truly quantum regime (single photons with high transverse coherence and a single-photon detector array) and a full arrangement (preparation of $|\psi\rangle$ and generation of $|\psi_\ell\rangle$'s via mode blockers), by keeping the errors at the same level of our experiment. Just as ptychography, quantum ptychography may be generalized in many ways; here we just demonstrated its first application.

Chapter 6

Future perspectives

The modern form of classical ptychography, when it was first proposed by Faulkner and Rodenburg [31], had some important restrictions. For example, it demanded that the illumination probe had a high degree of coherence [33], and that the beam profile was very well characterized and known *a priori* [34]. Such limitations are but natural when a new method is in its infancy.

Over the years, the community has managed to improve and generalize the classical method so as to overcome some of those limitations. For example, probe-retrieving ptychography [29] has been able to recover the probe profile as well as the object's, and extensions to partially coherent beams have also been achieved [33].

Many of those extensions do have an interesting analogue in the quantum version. We would like to briefly discuss some of them in this chapter.

6.1   Extension to mixed states

When dealing with quantum systems, it is often that they deviate from pure states. In fact, in many experiments we are able to prepare the systems in nearly-pure states, but experimental imperfections often reduce their purity.

Making an analogy with classical ptychography, if a pure state corresponds to an object with a fixed complex-valued transmission function, then a mixed state would be a statistical mixture of such objects (and therefore an statistical ensemble of complex-valued transmission functions) [33]. Physically speaking, this could represent an object with some

kind of internal dynamics, which happens much faster than the measurement window of our setup, for example. And while these will not be the typical objects which we subject to classical ptychography (though that is certainly possible), these are the natural objects subject to quantum ptychography.

We did obtain some interesting results in this regard, but they are mostly preliminary by the time of writing. We would like to briefly discuss what changes we have made to the method in order to make this generalization possible.

6.1.1 Projector families

The space of mixed states is embedded in the space of observables, which is d^2-dimensional. Although operators that represent physical states have further restrictions, we have to characterize them using observables themselves. In other words, mixed states are embedded in the space of physical observables, and it is the latter space that we have access to during our measurements. Therefore, we have to make our measurement set complete in the whole set of observables, and thus need at least d^2 members in our basis. It is expected, though, that we would need more measurements for mixed states than for pure states. The d-projector family that we saw in chapter 4 standed as our natural choice.

6.1.2 Complete set of measurements

After choosing a projector family and a unitary operation, it is possible to cast the set of measurement operators as a matrix (by expressing them in any complete basis in the operator space), and check its rank in order to verify if the measurement set is indeed complete.

We verified that the measurement set that results from the d-projector family and the Fourier transform is not informationally. We verified that fractional Fourier transform (which we will present shortly), instead of the usual Fourier transform, can overcome this issue, and chose to use them instead.

The fractional Fourier transform is closely related to Fresnel diffraction [3], and both free propagation and lens systems can be fomulated with it [32]. It has has interesting mathematical properties, such as having Hermite-Gauss polynomials as its eigenfunctions [60], and being interpretable as a rotation in the time-frequency domain [61]. It has also found many applications in signal processing [60]. The fractional Fourier transform of

order α, $\mathcal{F}_\alpha$, of a function f is defined as

$$\mathcal{F}_\alpha[f](u) = \sqrt{1 - i \cot \alpha}\, e^{i\pi \cot \alpha u^2} \int_{-\infty}^{\infty} e^{-i\, 2\pi\left(\csc \alpha u x - \cot \alpha x^2/2\right)} f(x) dx \quad , \qquad (6.1)$$

and the order α can be interpreted as the rotation angle in the time-frequency domain. As it would be expected, by taking $\alpha = \pi/2$, we can see that $\mathcal{F}_\alpha$ reduces to the usual Fourier transform $\mathcal{F}$. It can be shown that [62], for $a = 2\alpha/\pi$, $\mathcal{F}_a$ corresponds to the usual transform raised to the power a, that is,

$$\mathcal{F}_a = (\mathcal{F})^a \quad . \qquad (6.2)$$

In fact, it is possible to tackle the problem in the inverse direction [62], and start with equation (6.2) as the definition of the fractional transform and then arrive at its integral formulation (6.1).

Just as the usual Fourier transforn, the fractional transform also has a discrete version. In fact, the discrete fractional Fourier transform can be defined by a relation to the usual discrete transform by invoking a relation similar to equation (6.2). Let us start by remembering the matrix form of the discrete Fourier transform of dimension d, whose jk-th element is given by

$$F_{jk} = \frac{1}{\sqrt{d}} \left(e^{i\, 2\pi/d}\right)^{jk} \quad . \qquad (6.3)$$

It is possible to diagonalize the matrix F, and write it in the form

$$F = V\Lambda V^T \quad , \qquad (6.4)$$

with the matrix V containing the eigenvectors of F in its columns (or, equivalently, corresponding to the transformation to the basis of eigenvectors), and Λ having the eigenvalues in the diagonal. Now, it happens that F only has four distinct eigenvalues (though they have some multiplicity), namely the fourth roots of unity: $\{1, i, -1, -i\}$. Because of their multiplicity, the eigenvectors are not uniquely determined, as a linear combination of vectors in the same eigenspace yields another eigenvector. Irrespectively of this non-uniqueness, after settling for a given matrix V, the fractional discrete Fourier transform is defined [63] as

$$F_a = V\Lambda^a V^T \quad . \qquad (6.5)$$

In our mixed state simulations, we chose to use the discrete fractional Fourier

transform as the unitary transformation in place of the usual transform. This choice was motivated by the relative simplicity of implementing fractional transforms in optical setups [3, 32]. In our optical setup, which we presented in chapter 5, we would simply have to make measurements in intermediary planes, instead of the focal plane. We believe, though, that many other unitary transformations can be used, and encourage the reader to seek those which bring greater simplicity to their setup.

6.1.3   Choosing good measurement sets

After finding a complete set of measurements, we can also inspect the measurement set matrix's eigenvalues. If there are any whose modulus is very small (compared to the others), the part of the state that falls into the corresponding subspace is likely to be poorly reconstucted. It is recommended, therefore, to change either the projectors or the unitary operation so the smallest eigenvalue is not too small.

In our case, we saw that choosing the fractional order a of our Fourier transform, was an effective strategy to overcome this problem. This could be a practical shortcoming in an experiment though, as a very fine tuning of the fractional order might not be easy to achieve. Figure 6.1 shows how the least eigenvalue behaves as a function of fractional order, and indicates that there may be regions where it is somewhat stable, which would make this experimental implementation much more feasible.

6.1.4   Post-processing via SDP

The iterative post-processing discussed in chapter 4 did have some performance shortcomings when reconstructing mixed states. Even after some adaptations, we were still having significant stagnation problems, so we decided to use semi-definite programming[1] [22, 64]. There are many open-source libraries [23–26], for many programming languages, which make it fairly simple and straightforward to implement SDPs, and in fact they have been used in many fields for quite some time now, including quantum state reconstruction [21, 65].

---

[1]We have already encountered semidefinite programming earlier in this thesis. See footnote 7 of chapter 2.
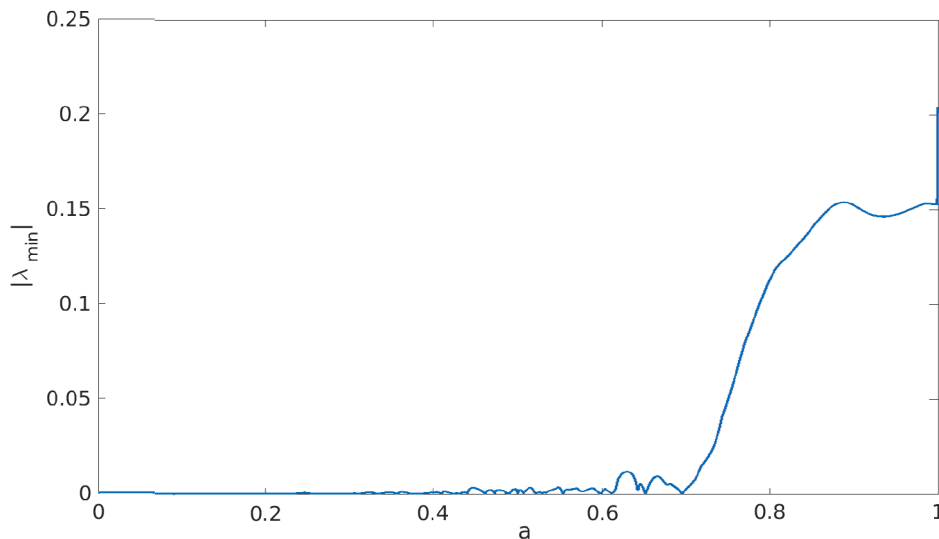
(a) $d = 11$, $r = 6$
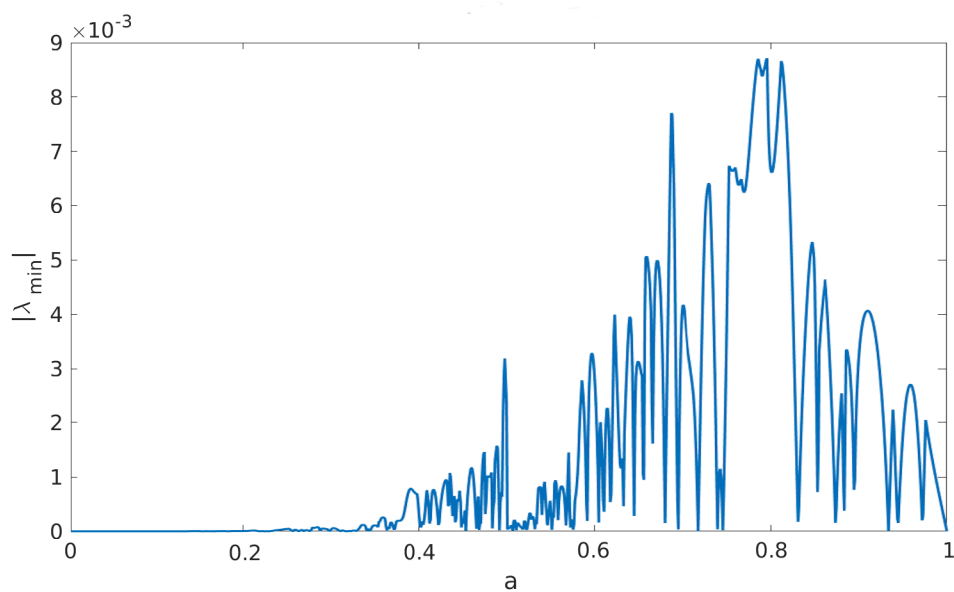
(b) $d = 16$, $r = 9$

Figure 6.1: **Minimum eigenvalues of measurement sets.** These graphs show the smallest modulus of the eigenvalues, $|\lambda_{\min}|$, of the measurement matrix vs. fractional Fourier order a. If $|\lambda_{\min}|$ is too small, the corresponding fractional order is likely to give bad reconstructions. It is advisable, therefore, to use values for a that fall in regions with higher values.

6.1.5 Preliminary results

The results we obtained after making the aforementioned adaptations were very encouraging. We have simulated blind reconstructions, using realistic noise models as in the pure case (section 4.4), and drawing random mixed states according to the Hilbert-Schmidt measure[2] [39]. The reconstruction fidelities[3] were very satisfactory. Figure 6.2 shows fidelities histograms for a few dimensions.

Overall, these results show that the ptychographic strategy of keeping the unitary transformation fixed, while attaining data diversity with different projectors, is still suitable for mixed states. This might be an interesting manner to simplify experimental setups in the future.

6.2 Process ptychography

As it was initially proposed by Faulkner and Rodenburg [31], classical ptychography assumed perfect knowledge of the probe beam profile. It was later shown that, in case the profile was incorrectly characterized, the reconstructions would be consequently jeopardized [29]. Later on, new extensions of the method showed that the probe could be reconstructed as well as the object [29,34].

In brief, this is possible because probe profile and object play analogous roles in the post-processing steps. It was still an important assumption, though, that the probe profile was constant throughout the translations along the object.

The analogous of probe retrieval, in the quantum protocol, would be to recover the probe projectors. This could be and interesting form of process tomography, in case the

---

[2]The Hilbert-Schmidt measure is naturally induced on the space of mixed states from the Haar measure of the unitary matrices, and can therefore be seen as the natural measure over the space of mixed states.

To draw a random mixed state ρ according to this measure, we start with a random unitary matrix $\hat{U}$, and then we generate a normalized hermitian matrix with the following correspondence:

$$\hat{\rho} = \frac{\hat{U}\hat{U}^\dagger}{\text{tr}\left(\hat{U}\hat{U}^\dagger\right)} \quad .$$

[3]Given two density matrices $\hat{\rho}$ and $\hat{\sigma}$, their fidelity is defined as

$$F(\hat{\rho}, \hat{\sigma}) = \left[\text{tr}\left(\sqrt{\sqrt{\hat{\sigma}}\hat{\rho}\sqrt{\hat{\sigma}}}\right)\right]^2 \quad .$$

It is easy to see that this definition reduces to the usual squared-fidelity for pure states.
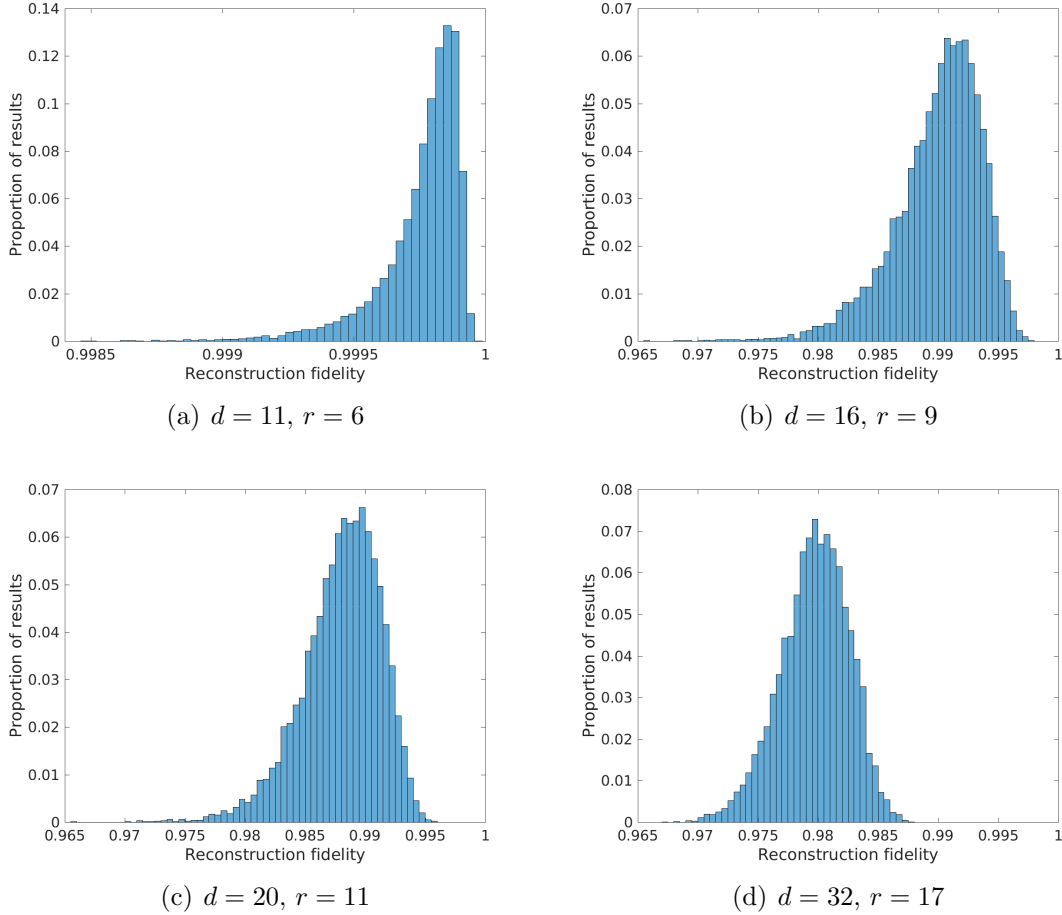
Figure 6.2: **Preliminary reconstruction results.** These figures show histograms of reconstructed fidelities for several dimensions. Our simulations included the same noise sources as in the pure state case. The reconstructions were very satisfying overall.

projectors were shifts of an unknown operator (i. e., the projectors are cyclic permutations of the operator relative to the computational basis), in a situation analogous to the shifting illumination in classical ptychography.

6.3 Post-processing improvements

It is somewhat often that improvements on optimization algorithms, and even entirely new algorithms, are proposed. These advances can bring benefits to the ptychographic method, by making the post processing faster and more accurate.

In [66], the authors use the momentum technique to avoid stagnation, and show

significant improvements on the performance of classical ptychography. They also make changes to the Wiener filter, and arrive at similar results. The two works also demonstrate that these performance gains also translate to the method being capable of reconstructing datasets which were previously not possible.

Therefore, we could expect that the use of higher-order gradient descent methods (e.g. the BFGS algorithm and other methods that use approximations to the Hessian matrix) could improve quantum ptychography speed, performance, and perhaps even make it work in increasingly difficult datasets.

6.4   Chapter summary

We have seen some interesting perspectives for future developments in quantum ptychography. In special, we have briefly described its generalization to mixed states and showed encouraging preliminary results. Classical ptychography, as initially proposed in 2004 by Rodenburg and Faulkner, has seen a number of interesting applications and extensions. We hope that quantum ptychography can follow a similar path.

Chapter 7

Conclusion

We have proposed quantum ptychography, a new method for reconstructing pure quantum states. It is based on an analogy with classical ptychography, a method for reconstructing wavefronts in classical optics. In this thesis we have formulated quatum ptychography mathematically and understood its workings by showing its equivalence to an alternated gradient search. Furthermore, we showed that it yields excellent reconstructions, both in numerical simulations and in an experiment.

The method does compare favorably with other methods in the literature. While having similar limitations, it requires a simpler experimental implementation since measurements are performed in a single basis, as opposed to the several bases that are generally required. We can say that using several intermediate projectors and a single measurement basis is indeed an effective strategy to attain experimental simplicity.

Another striking property of quantum ptychography is that it brings plenty of freedom in the choice of the intermediate projectors. This makes the method flexible with respect to the number of measurements. The experimenter can add more measurements for better reconstructions, or refrain from doing so to favor experimental simplicity (and therefore cost).

In order to give proper weight to this simplification, let us look at a concrete example. In the context of slit qudits, this means we can use a single lens and path blockers, instead of a setup to implement arbitrary unitary transformations, to reconstruct states. And this is without jeopardizing the quality of the reconstructions. This might be especially valuable in times of low investments in science, or any context of material scarcity.

The method also has many interesting extensions. We have discussed briefly our initial results regarding its generalization to mixed states, and discussed other interesting possibilities, such as using it for process tomography, how it could benefit with more sophisticated numerical methods.

It is worth noting that we have only studied the use of the Fourier transform and a few specific projector families. Though these specific choices were motivated by experimental simplicity, they should in no manner be taken as fundamental restrictions. In fact, we encourage the reader to pursue other projectors and unitary transformations that may better suit their needs.

As was the case of classical ptychography, we hope that these further developments can unfold in the near future, and that quantum ptychography can become a useful tool in the quantum information toolbox.

References

[1] Frank Arute et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019.

[2] Leslie E Ballentine. *Quantum mechanics: a modern development.* World Scientific Publishing Company, 2014.

[3] Joseph W Goodman. *Introduction to Fourier optics.* Roberts and Company Publishers, 2005.

[4] Godel escher bach letter cube. https://peterbeshai.com/experiments/godel-escher-bach/. Accessed: 2021-01-23.

[5] U. Fano. A stokes-parameter technique for the treatment of polarization in quantum mechanics. *Physical Review*, 93(1), 1954.

[6] Daniel F. V. James, Paul G. Kwiat, William J. Munro, and Andrew G. White. Measurement of qubits. *Physical Review A*, 64, 2001.

[7] H. Häffner, W. Hänsel, C. F. Roos, J. Benhelm, D. Chek al kar, M. Chwalla, T. Körber, U. D. Rapol, M. Riebe, P. O. Schmidt, C. Becher, O. Gühne, W. Dür, and R. Blatt. Scalable multiparticle entanglement of trapped ions. *Nature*, 438, 2005.

[8] David, Yi-Kai Liu, Steven T. Flammia, Stephen Becker, and Jens Eisert. Quantum state tomography via compressed sensing. *Physical Review Letters*, 105, 2010.

[9] Max S Kaznady and Daniel FV James. Numerical strategies for quantum tomography: Alternatives to full optimization. *Physical Review A*, 79(2):022109, 2009.

[10] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing,*

STOC '96, page 212–219, New York, NY, USA, 1996. Association for Computing Machinery.

[11] Peter W Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM review*, 41(2):303–332, 1999.

[12] David Deutsch and Richard Jozsa. Rapid solution of problems by quantum computation. *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences*, 439(1907):553–558, 1992.

[13] Charles H Bennett and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. *Proc. of IEEE Int. Conf. on Comp., Syst. and Signal Proc.*, 1984.

[14] Michael A Nielsen and Isaac Chuang. *Quantum computation and quantum information.* American Association of Physics Teachers, 2002.

[15] Hans J. Briegel and Robert Raussendorf. Persistent entanglement in arrays of interacting particles. *Phys. Rev. Lett.*, 86:910–913, Jan 2001.

[16] Steven T. Flammia, Andrew Silberfarb, and Carlton M. Caves. Minimal informationally complete measurements for pure states. *Foudations of Physics*, 35, 2005.

[17] Boris Zelikovich Moroz and Askol'd Mikhailovich Perelomov. On a problem posed by pauli. *Theoretical and Mathematical Physics*, 101(1):1200–1204, 1994.

[18] D. Goyeneche, G. Cañas, Etcheverry, E. S. Gómez, G. B. Xavier, G. Lima, and A. Delgado. Five measurement bases determine pure quantum states on any dimension. *Physical Review Letters*, 115, 2015.

[19] Howard Georgi. *Lie algebras in particle physics: from isospin to unified theories*, volume 54. Westview press, 1999.

[20] Quimey Pears Stefano, Lorena Rebón, Silvia Ledesma, and Claudio Iemmi. Determination of any pure spatial qudits from a minimum number of measurements by phase-stepping interferometry. *Physical Review A*, 96, 2017.

[21] Claudio Carmeli, Teiko Heinosaari, Michael Kech, Jussi Schultz, and Alessandro Toigo. Stable pure state quantum tomography from five orthonormal bases. *Europhysics Letters*, 115, 2016.

[22] Stephen Boyd, Stephen P Boyd, and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

[23] Irene: a python toolkit for general optimization. https://irene.readthedocs.io/en/latest/introduction.html. Accessed: 2021-02-27.

[24] Pylmi-sdp: Symbolic linear matrix inequalities (lmi) and semi-definite programming (sdp) tools for python. https://github.com/cdsousa/PyLMI-SDP. Accessed: 2021-02-27.

[25] Cvxopt: Python software for convex optimization. https://cvxopt.org/index.html. Accessed: 2021-02-27.

[26] Cvxpy: Cvxpy is a python-embedded modeling language for convex optimization problems. https://www.cvxpy.org/. Accessed: 2021-02-27.

[27] R. W. Gerchberg and W. O. Saxton. A practical algorithm for the determination of phase from image and diffraction plane pictures. *Optik*, 35:237–246, 1972.

[28] J R Fienup. Phase retrieval algorithms: a comparison. *Appl. Opt.*, 21(15):2758–2769, 1982.

[29] Manuel Guizar-Sicairos and James R Fienup. Phase retrieval with transverse translation diversity: a nonlinear optimization approach. *Optics express*, 16(10):7264–7278, 2008.

[30] James R Fienup. Phase retrieval algorithms: a personal tour. *Applied optics*, 52(1):45–56, 2013.

[31] H. M. L. Faulkner and J. M. Rodenburg. Movable aperture lensless transmission microscopy: A novel phase retrieval algorithm. *Phys. Rev. Lett.*, 93:023903, Jul 2004.

[32] Haldun M Ozaktas and David Mendlovic. Fractional fourier optics. *JOSA A*, 12(4):743–751, 1995.

[33] Pierre Thibault and Andreas Menzel. Reconstructing state mixtures from diffraction measurements. *Nature*, 494(7435):68–71, 2013.

[34] Andrew M Maiden and John M Rodenburg. An improved ptychographical phase retrieval algorithm for diffractive imaging. *Ultramicroscopy*, 109(10):1256–1262, 2009.

[35] Mário Foganholi Fernandes and Leonardo Neves. Ptychography of pure quantum states. *Scientific reports*, 9(1):1–10, 2019.

[36] Kenneth Kreutz–Delgado. The complex gradient operatorand the $\mathcal{CR}$-calculus, 2005.

[37] Dardo M. Goyeneche and Alberto C. de la Torre. State determination: An iterative algorithm. *Phys. Rev. A*, 77:042116, Apr 2008.

[38] AB Klimov, C Muñoz, A Fernández, and C Saavedra. Optimal quantum-state reconstruction for cold trapped ions. *Physical Review A*, 77(6):060303, 2008.

[39] Karol Zyczkowski and Hans-Jürgen Sommers. Induced measures in the space of mixed quantum states. *Journal of Physics A: Mathematical and General*, 34(35):7111, 2001.

[40] Karol Życzkowski, Karol A. Penson, Ion Nechita, and Benoît Collins. Generating random density matrices. *Journal of Mathematical Physics*, 52(6):062201, 2011.

[41] Michael Reck, Anton Zeilinger, Herbert J Bernstein, and Philip Bertani. Experimental realization of any discrete unitary operator. *Physical review letters*, 73(1):58, 1994.

[42] Je a Chiaverini, Joseph Britton, Dietrich Leibfried, Emanuel Knill, Murray D Barrett, RB Blakestad, Wayne M Itano, John D Jost, C Langer, R Ozeri, et al. Implementation of the semiclassical quantum fourier transform in a scalable system. *science*, 308(5724):997–1000, 2005.

[43] Philipp Schindler, Daniel Nigg, Thomas Monz, Julio T Barreiro, Esteban Martinez, Shannon X Wang, Stephan Quint, Matthias F Brandl, Volckmar Nebendahl, Christian F Roos, et al. A quantum information processor with trapped ions. *New Journal of Physics*, 15(12):123012, 2013.

[44] Matteo Mariantoni, Haiyan Wang, Takashi Yamamoto, Matthew Neeley, Radoslaw C Bialczak, Yu Chen, Mike Lenander, Erik Lucero, Aaron D O'Connell, Daniel Sank, et al. Implementing the quantum von neumann architecture with superconducting circuits. *Science*, 334(6052):61–65, 2011.

[45] Y. S. Weinstein, M. A. Pravia, E. M. Fortunato, S. Lloyd, and D. G. Cory. Implementation of the quantum fourier transform. *Phys. Rev. Lett.*, 86:1889–1891, Feb 2001.

[46] Kouichi Hosaka, Hiroyuki Shimada, Hisashi Chiba, Hiroyuki Katsuki, Yoshiaki Teranishi, Yukiyoshi Ohtsuki, and Kenji Ohmori. Ultrafast fourier transform with a femtosecond-laser-driven molecule. *Phys. Rev. Lett.*, 104:180501, May 2010.

[47] Chao-Yang Lu, Daniel E. Browne, Tao Yang, and Jian-Wei Pan. Demonstration of a compiled version of shor's quantum factoring algorithm using photonic qubits. *Phys. Rev. Lett.*, 99:250504, Dec 2007.

[48] B. P. Lanyon, T. J. Weinhold, N. K. Langford, M. Barbieri, D. F. V. James, A. Gilchrist, and A. G. White. Experimental demonstration of a compiled version of shor's algorithm with quantum entanglement. *Phys. Rev. Lett.*, 99:250505, Dec 2007.

[49] Andrea Crespi, Roberto Osellame, Roberta Ramponi, Marco Bentivegna, Fulvio Flamini, Nicolò Spagnolo, Niko Viggianiello, Luca Innocenti, Paolo Mataloni, and Fabio Sciarrino. Suppression law of quantum states in a 3d photonic fast fourier transform chip. *Nature communications*, 7(1):1–8, 2016.

[50] Mehul Malik, Mohammad Mirhosseini, Martin PJ Lavery, Jonathan Leach, Miles J Padgett, and Robert W Boyd. Direct measurement of a 27-dimensional orbital-angular-momentum state vector. *Nature communications*, 5(1):1–7, 2014.

[51] M. A. Solís-Prosser, M. F. Fernandes, O. Jiménez, A. Delgado, and L. Neves. Experimental minimum-error quantum-state discrimination in high dimensions. *Phys. Rev. Lett.*, 118:100501, Mar 2017.

[52] MF Fernandes, MA Solís-Prosser, and Leonardo Neves. Ptychographic reconstruction of pure quantum states. *Optics Letters*, 45(21):6002–6005, 2020.

[53] M. A. Solís-Prosser, A. Arias, J. J. M. Varga, L. Rebón, S. Ledesma, C. Iemmi, and L. Neves. Preparing arbitrary pure states of spatial qudits with a single phase-only spatial light modulator. *Opt. Lett.*, 38(22):4762–4765, Nov 2013.

[54] Miguel Ángel Solís Prosser. Estudio de la medición de qubits y qudits espaciales codificados en fotones individuales y aplicaciones, 2011. Master's dissertation, Departamento de Física, Universidad de Concepción.

[55] Roy J. Glauber. The quantum theory of optical coherence. *Physical Review*, 130, 1963.

[56] Juan José Miguel Varga, Lorena Rebón, MA Solís-Prosser, L Neves, S Ledesma, and C Iemmi. Optimized generation of spatial qudits by using a pure phase spatial light modulator. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 47(22):225504, 2014.

[57] D. Goyeneche, G. Cañas, S. Etcheverry, E. S. Gómez, G. B. Xavier, G. Lima, and A. Delgado. Five measurement bases determine pure quantum states on any dimension. *Phys. Rev. Lett.*, 115:090401, Aug 2015.

[58] Q P Stefano, L Rebón, S Ledesma, and C Iemmi. Determination of any pure spatial qudits from a minimum number of measurements by phase-stepping interferometry. *Phys. Rev. A*, 96(6):062328, 2017.

[59] Quimey Pears Stefano, Lorena Rebón, Silvia Ledesma, and Claudio Iemmi. Set of $4d - 3$ observables to determine any pure qudit state. *Opt. Lett.*, 44(10):2558–2561, May 2019.

[60] Ervin Sejdić, Igor Djurović, and LJubiša Stanković. Fractional fourier transform as a signal processing tool: An overview of recent developments. *Signal Processing*, 91(6):1351–1369, 2011.

[61] L.B. Almeida. The fractional fourier transform and time-frequency representations. *IEEE Transactions on Signal Processing*, 42(11):3084–3091, 1994.

[62] Victor Namias. The fractional order fourier transform and its application to quantum mechanics. *IMA Journal of Applied Mathematics*, 25(3):241–265, 1980.

[63] Adhemar Bultheel and Héctor E Martınez Sulbaran. Computation of the fractional fourier transform. *Applied and Computational Harmonic Analysis*, 16(3):182–202, 2004.

[64] Lieven Vandenberghe and Stephen Boyd. Semidefinite programming. *SIAM review*, 38(1):49–95, 1996.

[65] Thiago O Maciel, André T Cesário, and Reinaldo O Vianna. Variational quantum tomography with incomplete information by means of semidefinite programs. *International Journal of Modern Physics C*, 22(12):1361–1372, 2011.

[66] Andrew Maiden, Daniel Johnson, and Peng Li. Further improvements to the ptychographical iterative engine. *Optica*, 4(7):736–745, 2017.

Appendix A

Simulation parameters

We have made numerical simulations as part of our theoretical proposal of quantum ptychography, which we saw in chapter 4. The motivation was to verify that the protocol would indeed work under conditions that are usually found in a laboratory.

In this appendix, we describe how we optimized and chose some parameters of those simulations.

A.1   Update rate

The parameter β (equation (4.11)) controls the step-size of the update in the PIE algorithm and can be adjusted to improve its convergence. For $\beta = 1$, the algorithm corrects the estimate strictly in the subspace spanned by $\hat{P}_\ell$; higher values can make it progress faster and converge in less iterations; lower values can make it slower but more stable. Therefore, it is advisable to run a few reconstructions with several values of β and compare their performance.

Using the same initially estimated and target state, we obtained the optimal β by running the PIE and recording the relative distance

$$D = \frac{||\phi'\rangle - |\phi\rangle|^2}{||\phi\rangle|^2} \tag{A.1}$$

between current and last estimate, for several values of the parameter. Figure A.1 shows the evolutions for a few values of β. The best progression was achieved by $\beta = 1.5$, which we used in all simulations.
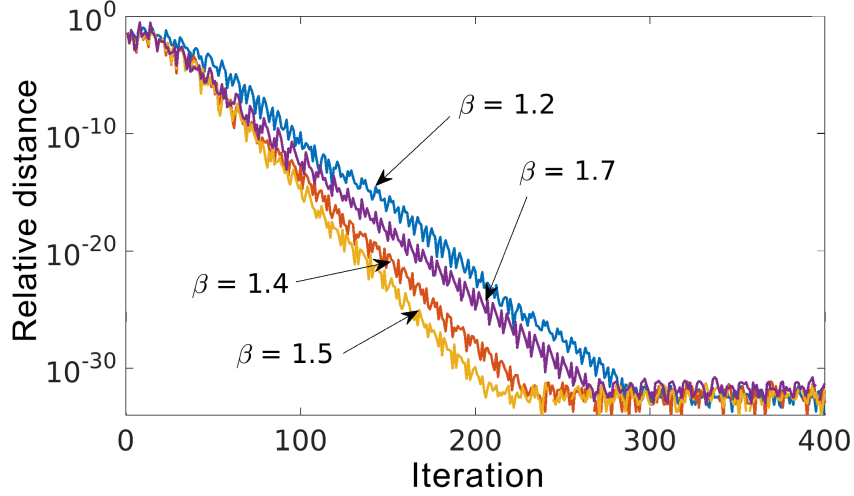
Figure A.1: **Optimization of the update rate.** The curves show the evolution of the relative distance D (eq. (A.1)) for different values of the feedback parameter, β. We found $\beta = 1.5$ to achieve the fastest decrease, meaning that the PIE algorithm converged at a higher rate.

A.2   Projector ranks

After finding the best β, we were able to determine the optimal r of the probe projectors given by equation (4.23). We studied the convergence of the PIE algorithm as a function of this rank for a few state-space dimensions ($d = 10, 15, 20$). For each combination of d and r, we reconstructed 10^4 random states and calculated the average number of iterations necessary until convergence. The results are shown in figure A.2 and indicate that a rank around $d/2$ works best. In our simulations we alternated between $\lfloor d/2 \rfloor$ and $\lceil d/2 \rceil$. In general, we verified that both, as well as any other close value, produced similar results regarding the quality of the reconstructions.

A.3   Noise level

We introduced depolarization and Poissonian noise in the ptychographic data to study the protocol in a realistic scenario. As mentioned in the main text, we based our noise levels on experiments found in the literature, but we still wanted to verify if they were indeed realistic. For this purpose, we picked 10^4 random pure states, degraded their amplitudes with the two kinds of noise described in section 4.4, and computed the fidelities with
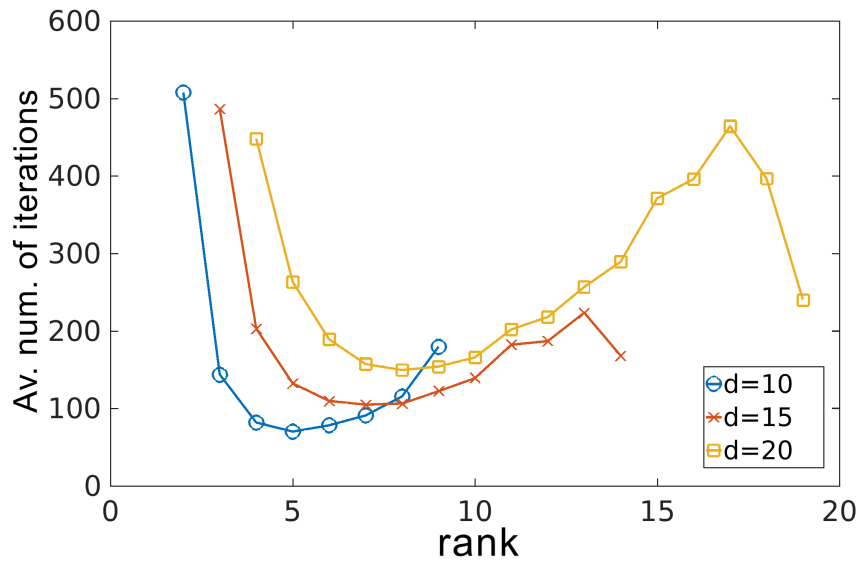
Figure A.2: **Optimization of projector ranks.** The curves show the average number of iterations until convergence of the PIE algorithm as a function of the probe projector rank. Ranks around $d/2$ made the algorithm converge faster

respect to the original states. Figure A.3 shows a histogram of the degraded fidelities, which are indeed comparable—and even lower—to actual experiments [56]. This confirms, therefore, that our simulations used realistic amounts of noise.
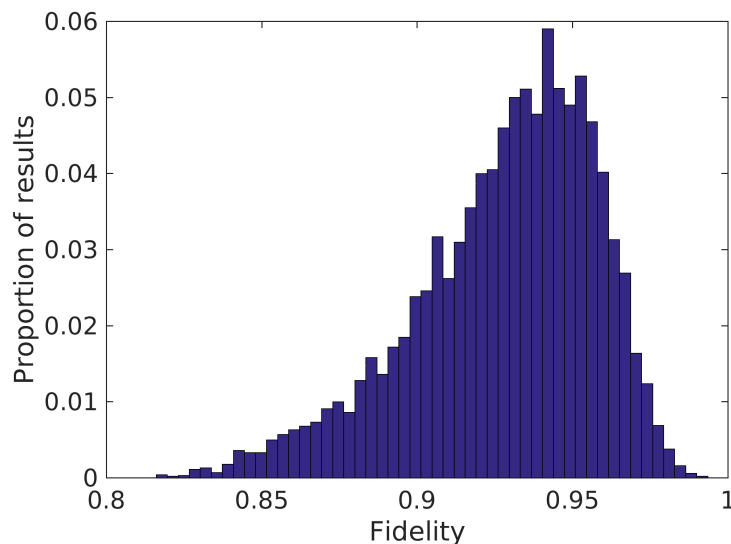

Figure A.3: **Verification of noise level.** This histogram of degraded fidelities shows how severely the noise level affected our simulated states. It is quite similar to experimental results in the literature [56], and thus confirms that the noise parameters in our simulations were indeed realistic.