


Noise reduction in absorption imaging for quantum gases experiments

Student Paper

Author(s):

Montalti, Nicolò 

Publication date:

2024

Permanent link:

<https://doi.org/10.3929/ethz-b-000673736>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)



Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Noise reduction in absorption imaging for quantum gases experiments

Semester Project

Nicolò Montalti
nmontalti@student.ethz.ch

Quantum Optics Group
Departement of Physics, D-PHYS
ETH Zürich

Supervisor:
Prof. Tilman Esslinger

Advisor:
Dr. Panagiotis Christodoulou

16 April 2024

Abstract

In many experiments with quantum gases, the information about the final state of the system is extracted by the density profile of the cloud. One of the most common techniques used to image the cloud is absorption imaging, which is limited by various kinds of noise. One of the main sources of noise is the presence of fringes that originates from vibrations in the experiment. In our case, the problem is increased by the presence of an optical cavity around the atoms which diffracts the imaging light. These fringes can be eliminated by appropriate fringe-removal algorithms. We show that these algorithms reduce the noise down to being close to the limit imposed by the CCD camera. We also analyse some patterns that persist in the optical density after removing the fringes. We relate these patterns to shifts of the atomic cloud or of the imaging system between different runs of the experiment and to diffraction by the atoms and the cavity.

Contents

| | |
|--|-----------|
| Abstract | i |
| Introduction | 1 |
| 1 Fringe removal in absorption imaging | 3 |
| 1.1 Absorption imaging | 3 |
| 1.1.1 Absorption-imaging fundamentals | 4 |
| 1.1.2 Absorption imaging in our experiment | 5 |
| 1.2 Fringe-removal algorithms | 5 |
| 1.2.1 The optimized fringe-removal algorithm (OFRA) | 8 |
| 1.2.2 The effective statistical fringe-removal algorithm (ESFRA) | 11 |
| 2 Optimization and characterization of fringe removal algorithms | 12 |
| 2.1 Optimization and comparison | 12 |
| 2.1.1 Benchmarking strategy | 12 |
| 2.1.2 Optimization of the OFRA | 14 |
| 2.1.3 Comparison between the OFRA and the ESFRA | 16 |
| 2.2 Origin of the residual noise | 17 |
| 2.2.1 Introduction to CCD noise | 18 |
| 2.2.2 Results | 19 |
| 2.3 Effectiveness of the algorithms on real data | 20 |
| 2.3.1 Benchmarking strategy | 20 |
| 2.3.2 Results | 21 |
| 3 Identification of noise patterns in the optical density | 23 |
| 3.1 PCA for identification of noise patterns | 23 |
| 3.2 Cavity diffraction artefacts | 26 |
| 4 Conclusions | 28 |
| Acknowledgements | 29 |
| Bibliography | 30 |

Introduction

“Nature isn’t classical, dammit, and if you want to make a simulation of nature, you’d better make it quantum mechanical, and by golly it’s a wonderful problem, because it doesn’t look so easy.”

— Richard Feynman

Predicting the behaviour of complex systems has always been an important goal of scientific research, and many-body quantum mechanics offers us some of the most complex problems known in physics. The possibility of solving the dynamics of many body systems computationally or analytically is very often limited by the size of the Hilbert space involved. Richard Feynman was the first to propose to use a quantum system to handle this complexity, making use of the laws of quantum mechanics itself to solve the problem of interest [1]. The problem is then shifted to finding a system that can be easily tuned and controlled, such that it can be used to simulate a wide range of problems. In the last years this has become possible, both through digital quantum computation [2–4] and analogical quantum simulation, for which the use of cold atoms has gained increasing popularity [5]. Many experiments were recently realized to simulate the behaviour of different quantum systems. Among them, we cite the use of optical lattices to investigate insulating and superfluid quantum phases [6], the observation of quantized conductance in neutral matter [7] and the use of cavities to mediate long range interactions [8].

In all of these experiments a fundamental role is played by the measuring of the final state of the gas. The typical quantity that is measured at the end of the experiment is the density profile of the cloud. If done after time of flight, this can give us information about the momentum distribution of the gas, and therefore also its temperature, but in general there are many observables that can be derived from the density distribution. For example, for the experiments performed in our group where the atoms are placed inside a cavity, the density profile can be used to detect phase transitions [8, 9] or the spin state of the atoms [10]. This report will focus on this specific aspect of quantum gases experiments, with a special emphasis on a particular technique used to detect the cloud density profile: absorption imaging.

Outline

The goal of this report is to provide a brief introduction to absorption imaging and its limitations, with a specific focus on how fringes appear in the density profile and how they can

be eliminated. In the first chapter we will introduce absorption imaging and fringe removal algorithms, focusing on two algorithms that have been proved to be particularly useful for our experiment: the optimized fringe removal algorithm [11] and the effective statistical fringe removal algorithm [12]. These two algorithm will then be implemented, characterized and optimized in the second chapter. In the third chapter we will turn our attention to some residual noise patterns in the optical density that cannot be removed by fringe removal algorithms. Some of these patterns will be analysed using principal component analysis, while another one will be investigated through optical simulations. The discussion on both fringe removal algorithms and noise patterns identification will be based on experimental data obtained from a cavity experiment. The results will be discussed with the goal of optimizing our specific setup, but most of the conclusions drawn can be extended to more general experimental configurations.

Chapter 1

Fringe removal in absorption imaging

“In every way, then, such prisoners would recognize as reality nothing but the shadows of those artificial objects.”

— Plato, The Allegory of the Cave

In quantum-gases experiments, the information encoded in the system is often extracted by looking at the density distribution of the atomic cloud. For example, the density can be measured after releasing the atoms from the trap and letting the gas expand during time-of-flight (TOF). After expansion, the density distribution contains the information on the original momentum distribution of the trapped gas, from which quantities like the temperature can be extracted. In order to measure the density, several optical imaging schemes have been proposed, which usually rely either on the absorptive or dispersive properties of the atomic cloud. The two most common techniques are absorption and phase-contrast imaging. In absorption imaging a laser beam is sent to the atomic cloud, and the density of the cloud is recovered from the shadow that the atoms produce absorbing part of the light. In phase-contrast imaging, the light sent to the atoms is far detuned, such that the cloud only alters the phase of the beam. Making the part of the beam altered by the atoms interfere with the unaltered beam, this phase modulation can be translated into an intensity modulation and detected [13]. In the following, we will focus on absorption imaging, which is the technique used in our experiment to image the atomic cloud at the end of the experimental cycle.

1.1 Absorption imaging

In this section we want to offer a brief introduction to absorption imaging. We will start by introducing some basic concepts, and then we will focus on the specific setup used in our experiment. Finally, we will discuss the limitations of absorption imaging and how some of them can be overcome. The content of this section is mostly adapted from the work of Tammuz [14] and Kosucu [15].

1.1.1 Absorption-imaging fundamentals

If a laser beam resonant with an atomic transition is sent to an atomic cloud, part of it is absorbed. Placing a charge-coupled device (CCD) behind the cloud allows us to detect the shadow generated by the atoms. The higher is density of the cloud, the darker the shadow will be. The goal of this section is to translate this general idea into a procedure to quantitatively recover the density of the cloud.

We start by calculating the amount of light absorbed by a single atom, modelled as a simple two-level system with a single excited state. We can quantify it with the scattering rate $R_{\text{scatt}} = \Gamma \rho_{ee}$, where $\Gamma = \tau^{-1}$ is the inverse of the transition lifetime τ and ρ_{ee} is the density matrix element for the excited state. ρ_{ee} can be obtained solving the optical Bloch equations (OBE) for the steady state. The resulting scattering rate is

$$R_{\text{scatt}} = \frac{\Gamma}{2} \frac{I/I_s}{1 + I/I_s + (2\delta/\Gamma)^2} \quad (1.1)$$

where I is the intensity of the beam, δ is the detuning and $I_s = \pi \hbar c / 3 \lambda_0 \tau$ is the saturation intensity, with λ_0 resonance wavelength. If the laser beam has intensity $I(x, z)$, after travelling through a gas of density n for a distance dy , it will be absorbed for an amount proportional to the column density ndy and to the scattering rate R_{scatt}

$$\frac{dI}{dy} = -\hbar\omega R_{\text{scatt}} n = -\sigma n I \quad (1.2)$$

where ω is the angular frequency of the light, and we define σ to be the absorption cross-section. From Eqs. (1.1) and (1.2) we find

$$\sigma = \frac{3\lambda^2}{2\pi} \frac{1}{1 + I/I_s + (2\delta/\Gamma)^2} \quad (1.3)$$

Equation (1.2) can be integrated, resulting in an exponential decay of the form

$$I(x, z) = I_0(x, z) \exp \left[-\sigma \int n(x, y, z) dy \right] = I_0(x, z) e^{-OD} \quad (1.4)$$

where I_0 is the intensity distribution before the absorption and OD is the optical density, defined as

$$OD = \sigma \int n(x, y, z) dy = \sigma n_{2D} = -\ln \left(\frac{I(x, z)}{I_0(x, z)} \right) \quad (1.5)$$

From Equation (1.5) we see how the column density $n_{2D}(x, z)$ is related to the optical density through $n_{2D}(x, z) = OD/\sigma$. Knowing the optical density is therefore sufficient to recover the momentum distribution of the expanding cloud, and from that quantities like the temperature or the total number of atoms N . For example, the total number of atoms can be found integrating the column density over the whole area

$$N = \int n dx dy dz = \int n_{2D} dx dz \quad (1.6)$$

To find the OD , it is sufficient to compare the intensity distribution of the beam before and after passing the cloud (respectively I_0 and I in Eq. (1.5)). Since taking an image of the beam before the cloud is not experimentally feasible, in practice the two images are taken at different times. After having acquired the image with the atoms (atoms image A), a second image is acquired with the laser propagating freely in space (bright image B). To avoid the presence of background light in the lab to affect the result, a third image is also taken with the laser off. This image is usually called dark image D . The three images are finally combined to find the optical density

$$OD = -\ln\left(\frac{A - D}{B - D}\right) \quad (1.7)$$

1.1.2 Absorption imaging in our experiment

In our experiment, when imaged, the atoms are contained inside a cavity. The cavity is formed by two mirrors with a diameter of 3 mm placed at a central distance of 178 μm . Because of the curvature of the mirror, their edges are separated by just 150 μm . The cavity itself is placed inside a vacuum chamber, and light can be sent inside and detected through two viewports. After exiting the vacuum chamber from the second viewport, the laser beam is magnified with a telescope composed of three lenses of 300, 400 and -100 mm. The result is then imaged with a CCD camera. A sketch of this setup is given in Fig. 1.1.

The configuration described above presents an important limitation. Since it is beneficial to illuminate the atoms with a uniform intensity, the waist of the incoming laser beam must be large. In her precedent work [15], Kosucu reports a predicted beam waist of 700 μm , obtained through ray-tracing calculations. Since the beam is bigger than the distance between the mirrors, the cavity acts like a thick slit, producing diffractive patterns. An example is provided in Fig. 1.2. In principle these patterns should not be visible in the optical density, since they are present both in the atoms and bright pictures. In practice, if we calculate the optical density as suggested by Eq. (1.7), we observe some residual fringes (cfr. Fig. 1.2). This can be explained by assuming that some components of the setup, which determine the direction of the imaging beam, move in the time between when the two pictures are taken. As a result, the two pictures are not perfectly aligned and the diffraction patterns do not cancel out completely.

Reducing these fringes was the first goal of this project. In the rest of the chapter, we will present two algorithms that try to mitigate this effect.

1.2 Fringe-removal algorithms

Reducing the fringes generated by the vibrations in the imaging system is an essential task to improve the ability to count atoms in our experiment. Researchers have proposed different post-processing algorithms to achieve this result [17]. Many of them share the same basic idea: generate an ideal bright reference image that plugged into Eq. (1.7) produces the correct OD without additional fringes. The methods to generate this ideal reference image vary and include the use of singular value decomposition (SVD) [18], principal component analysis (PCA) [11] and deep learning methods [19]. In this report we will focus on PCA-based methods. In




Figure 1.1: Cavity inside the vacuum chamber and imaging system. A laser beam is sent to the atoms inside the cavity through the first viewport, it is partially absorbed by the atoms, and, after being magnified, it is imaged by a CCD camera. Figure taken from Ref. [16]




Figure 1.2: In order: atoms (A), bright (B) and dark (D) image and optical density (OD). The intensity of the first four images is given in an arbitrary scale, whereas OD has the right dimensionless units. The scale of OD_{sat} is adjusted to increase the visibility of the fringes. The vertical fringes generated by the cavity are clearly visible in A and B . In OD most of these fringes cancel out, but some vertical and diagonal fringes can still be observed in OD_{sat} .

the next section we will present the optimized fringe removal algorithm (OFRA) proposed by Niu *et al.* [11], and in the following we will discuss an extension of it, namely the effective statistical fringe removal algorithm (ESFRA) proposed by Song *et al.* [12]. During the project, the enhanced principal component analysis (EPCA) proposed by Xiong *et al.* [20] was also tested. Since we could not achieve results of comparable quality with this algorithm, we decided not to include further details into this report.

1.2.1 The optimized fringe-removal algorithm (OFRA)

In the optimized fringe-removal algorithm (OFRA), the basic idea is to construct a set of bright images that can serve as a basis to generate the best ideal bright image B'_R for each atoms image A_R . Once the ideal bright image has been generated, the OD can be found as

$$OD = -\ln\left(\frac{A_R}{B'_R}\right) \quad (1.8)$$

The subscript R denotes *real* images, in contrast to *centred* images A and B . In fact, in order to generate the basis, it is useful to have the data centred at zero intensity. Therefore, after subtracting the corresponding dark images D_R , we calculate the pixel-wise mean image \bar{B} from the bright pictures and centre each picture as $A = A_R - D_R - \bar{B}$ and $B = B_R - D_R - \bar{B}$.

As anticipated, the first step is to decompose the bright images B_n into a set of orthogonal basis images P_m , with n being the index of the original bright images and m the index of the basis images. To do so, we first define an area where we are sure not to find any atom in any image, which we call *edge area*. In this way, we can find a basis set that only depends on the fringes pattern and not on the density distribution of the atoms. For example, looking at Fig. 1.2, we can use the bottom part of the picture. We then crop the bright images to only include this area, and we reshape each cropped 2D vector B_n in a 1D vector \mathbf{b}_n . Therefore, in total we have N vectors \mathbf{b}_n of length L , where N is the number of bright pictures we are using to construct the basis and L is the number of pixels in the edge area. Typically, $L \sim 10^4$ is much larger than $N \sim 10^2$.

Keeping in mind that the expectation value of each pixel in the vectors \mathbf{b} , $\langle \mathbf{b}_n \rangle_n$, is zero, it is possible to define the covariance matrix between two pixels (i, j) as $S_{ij} = \sum_n (b_i)_n (b_j)_n$, where $(b_i)_n$ is the i^{th} pixel of the n^{th} image. Since S_{ij} is a symmetric matrix, it can be diagonalized using singular value decomposition (SVD). The resulting eigenvectors \mathbf{v}_n form a set of $N - 1$ independent orthogonal basis vectors. The missing independent vector is the result of fixing the mean of the data set.

The decomposition of vectors in a set of orthogonal vectors that diagonalize the covariance matrix is usually known as principal value analysis (PCA) [21]. The key advantage of PCA is that the resulting eigenvectors are the directions in the feature space (i.e. the space spanned by the vectors \mathbf{b}_n) of maximal variance. In other words, they can explain the variance in the original data, with the most important components being the ones with higher eigenvalues. In our situation this corresponds to the ability of explaining the fringe patterns, with the components with higher eigenvalues showing the main fringes and the ones with lower eigenvalues being mostly noise. An example of some components with their eigenvalues is shown

in Fig. 2.1a. Since most of the components do not contribute to the fringe pattern, it can be convenient to reduce the basis and keep only the first $M < N$ components.

After having found a set of M basis vectors for the edge area $\{\mathbf{v}_m\}_{m < M}$, we need to use it to find a suitable basis of images that also include the area with the atoms. Using the orthogonality of the basis vectors, we can find a coefficient matrix $C_{nm} = \mathbf{b}_n \cdot \mathbf{v}_m$ such that

$$\mathbf{b}_n = C_{nm} \mathbf{v}_m \quad (1.9)$$

where we have used Einstein's notation of repeated indices. We will keep using this notation in the rest of the report unless differently specified. Equation (1.9) can be seen as L systems (one for each pixel) of $N \times M$ linear equations with M unknowns \mathbf{v}_m . Since we have more equations than unknowns, the system is over-determined, and it must be solved using a least-squares method. This means finding the matrix \tilde{C} such that

$$\tilde{\mathbf{v}}_m = \tilde{C}_{mn} \mathbf{b}_n \quad (1.10)$$

minimizes the error $\sum_m (\tilde{\mathbf{v}}_m - \mathbf{v}_m)^2$ for each element (pixel) of \mathbf{v} . The matrix \tilde{C} is called the pseudo-inverse of C , and it can be calculated numerically. The matrix \tilde{C} can be finally used to construct a basis of bright images V_m . For each edge vector \mathbf{b}_n there is a corresponding full image B_n . So we can use Eq. (1.10) to find a basis of images $V_m = \tilde{C}_{mn} B_n$. Note that here V and B are again 2D vectors.

The final goal is to generate an ideal bright image B'_n for each atoms image A_n . To do so, we crop the atoms image A_n in the same edge area defined above and we convert the result in a 1D vector \mathbf{a}_n . Now we want to express \mathbf{a}_n as a linear combination of the basis vectors \mathbf{v}_m .

$$\mathbf{a}_n = K_{nm} \mathbf{v}_m \quad (1.11)$$

Since the basis vectors are orthogonal, the coefficients K_{nm} can be found as $K_{nm} = \mathbf{a}_n \cdot \mathbf{v}_m$, and the ideal bright image is given by a linear combination of the basis bright images V_m with the same coefficients K_{nm}

$$B'_n = K_{nm} V_m \quad (1.12)$$

In practice, in the implementation algorithm, we use a more general method to find the coefficients K_{nm} in Eq. (1.11) that also works for non-orthonormal basis, which will be useful when discussing the ESFRA (see section below). We start by multiplying Eq. (1.11) by \mathbf{v}_p on both sides, and we define the autocorrelation matrix $G_{mp} = \mathbf{v}_m \cdot \mathbf{v}_p$ and the projection matrix $P_{np} = \mathbf{a}_n \cdot \mathbf{v}_p$. Equation (1.11) becomes

$$P_{np} = K_{nm} G_{mp} \quad (1.13)$$

Note that when the vectors are orthonormal, $G_{mp} = \delta_{mp}$, and $P_{np} = K_{np}$. In general, Eq. (1.13) can be solved for K_{nm} minimizing the least-square error as explained above, and used to reconstruct the ideal bright image

$$B'_n = K_{nm} V_m \quad (1.14)$$

We can then recentre the result adding \bar{B} and use Eq. (1.8) to find the optical density.

For better clarity, we also show a pictorial representation of the algorithm in Fig. 1.3a and a summary in pseudocode in Alg. (1.1). The real code can be found in the Quantum Optics group archive.

Algorithm 1.1: OFRA

```

1   Input: "array of real atoms images"  $A_{Rn}$ 
2   Input: "array of real bright images"  $B_{Rn}$ 
3   Input: "array of real dark images"  $D_{Rn}$ 
4   Input: numberComponents
5
6   Output: "optical density"  $OD_n$ 
7
8    $A_{Rn} -= D_{Rn}$ 
9    $B_{Rn} -= D_{Rn}$ 
10
11   $\bar{B} = \text{mean}(B_{Rn})$ 
12   $A_n = A_{Rn} - \bar{B}$ 
13   $B_n = B_{Rn} - \bar{B}$ 
14
15  edgeArea = "area without atoms"
16
17   $\mathbf{b}_n = \text{toVector}(\text{crop}(B_n, \text{edgeArea}))$ 
18
19  # construction of the basis
20   $S_{ij} = \sum_n (b_i)_n (b_j)_n$ 
21   $\mathbf{v}_n, \lambda_n = \text{SVD}(S)$  # eigenvectors, eigenvalues
22  keep only  $\mathbf{v}_m$  with  $m < \text{numberComponents}$ 
23   $C_{nm} = \mathbf{b}_n \cdot \mathbf{v}_m$ 
24   $\tilde{C} = \text{pinv}(C)$  # pseudo-inverse
25   $V_m = \tilde{C}_{mn} B_n$ 
26
27  # construction of the ideal bright image
28   $P_{np} = \mathbf{a}_n \cdot \mathbf{v}_p$ 
29   $G_{mp} = \mathbf{v}_m \cdot \mathbf{v}_p$ 
30   $K_{nm} = [\text{linearSolve}(G^T, P^T)]^T$  # solve  $P=KG$  for  $K$ 
31   $B'_n = K_{nm} V_m$ 
32
33   $B'_{Rn} = B'_n + \bar{B}$ 
34   $OD_n = -\log(A_{Rn} / B'_{Rn})$ 
35
36  return  $OD_n$ 

```

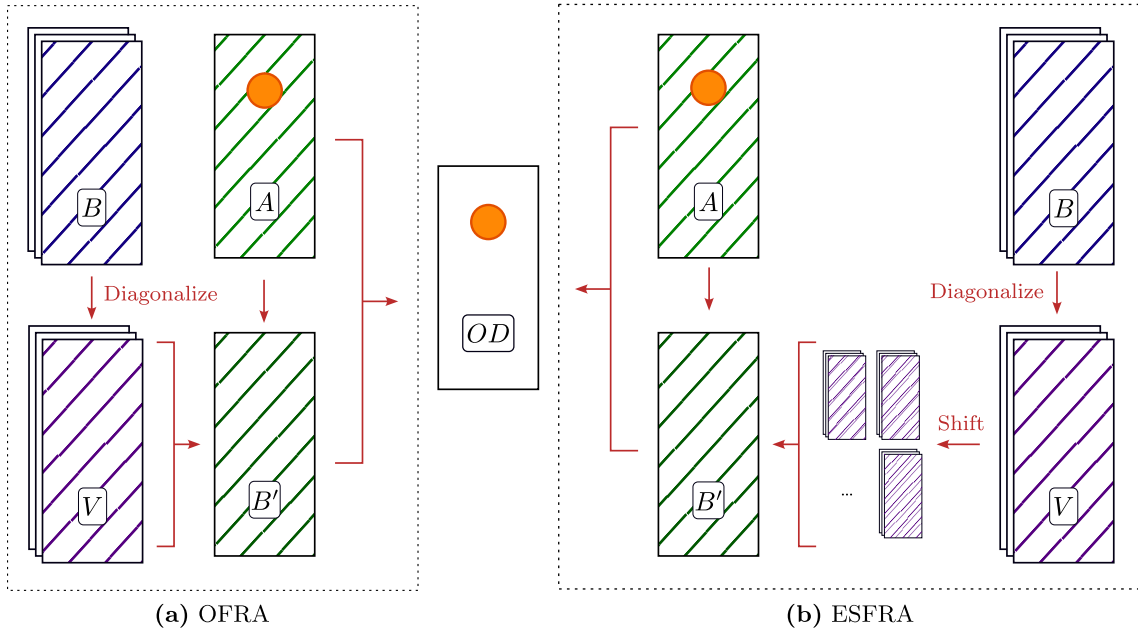


Figure 1.3: Graphical description of the (a) OFRA and (b) ESFRA. The OFRA uses a set of bright images to find a basis that can be used to generate the ideal bright image for each atoms image. The ESFRA extends the OFRA basis by shifting the images by a couple of pixels in the two directions.

1.2.2 The effective statistical fringe-removal algorithm (ESFRA)

The main origin of the fringes are believed to be vibrations of the acquisition system, for example of the optics that shapes the imaging beam. The OFRA is an effective way to reduce the noise introduced by these effects, but it may require a big dataset to find a good basis. This means longer acquisition times, since each run of the experiment takes around 45 s. To tackle this inconvenience, the effective statistical fringe-removal algorithm (ESFRA) developed by Song *et al.* [12] proposes a method to increase the number of elements of the basis without taking more data. Since the fringes originate from displacements of the camera, they propose to increase the size of the basis by shifting the images by a couple of pixels in the two directions. We will indicate the maximum displacement, measured in pixels by d . After performing the translation, the basis is increased by a factor $(2d + 1)^2$. For example for $d = 1$ we have 4 additional images, two for each direction of translation. After enhancing the dataset, the basis is not orthonormal any more, and to reconstruct the bright image we need to use the general method described above. The rest of the algorithm is completely analogous to the OFRA. A pictorial representation of the ESFRA is shown in Fig. 1.3b.

Chapter 2

Optimization and characterization of fringe removal algorithms

“You know that children are growing up when they start asking questions that have answers.”

— John J. Plomp

Now that we have introduced the principles of fringe removal algorithms, we want to characterize and optimize them. The algorithms presented in the previous chapter, namely the OFRA and ESFRA, were implemented in Python and tested on different datasets to try to characterize their behaviour and efficiency. In this chapter we will present the tests that we made and their results. The chapter is divided in three parts. In the first, we optimize the parameters of the OFRA and ESFRA, and we compare their efficiencies on pictures without atoms. In the second, we try to understand what is the limit to the residual noise that is left after running the algorithms, focusing on the noise generated by the camera. In the last section we test the algorithms on pictures with atoms, and evaluate their performance with a particular focus on counting the number of atoms.

2.1 Optimization and comparison

In order to use the algorithms we introduced before, it is important to understand how to implement them to achieve the best results. This means finding the best set of parameters, like the size of the training set or the number of principal components kept after the diagonalization, and evaluating the effectiveness of the different algorithms in removing the fringes. In the next sections, we will explain the benchmarking method we used to evaluate the performance of the algorithms, and we will apply it to the study of the OFRA and the ESFRA.

2.1.1 Benchmarking strategy

In order to characterize, optimize and confront different algorithms, we need to develop a benchmarking strategy to evaluate quantitatively how “good” an algorithm is performing. We

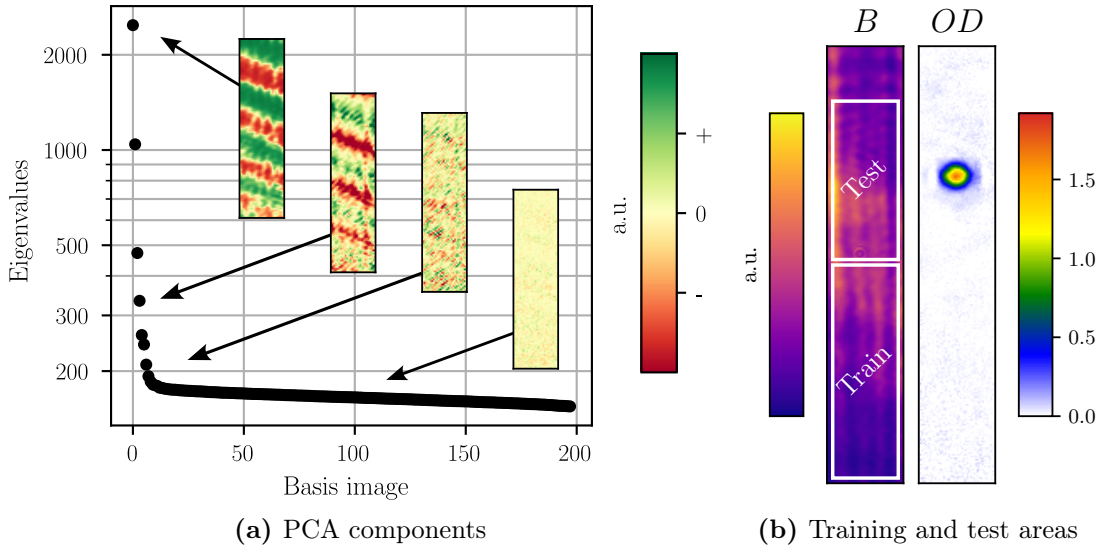


Figure 2.1: (a) PCA components of the OFRA algorithm. The basis is generated from 200 bright images, resulting in a basis of 199 independent vectors. The eigenvalues suggest that only less than 10% of the vectors (20) are relevant. This is confirmed qualitatively looking at the images corresponding to the respective eigenvectors. Note that the eigenvectors are centred at zero. (b) The training and test areas are shown on the bright picture. Next to it, we show the OD in the corresponding region. The training area is used to find the basis of bright images, whereas the test area is used to evaluate the performance of the algorithm.

can then change the input parameters, the training data and the experimental conditions to judge which configuration works better.

For this first part of the work, until differently specified, we used the data collected on 21.10.2022 in the main run of the experiment. On that day, the experiment was run 590 times, producing 590 atoms, bright and dark pictures. In most of the analysis we do not use the full 590 pictures, but only a subset. When doing so, we use the most recent pictures. So when we say we train the algorithm on 200 pictures, we mean we train it on the last 200 pictures taken on that day. As customary, we usually split this dataset in two parts with different pictures: a training set and a test set. The training set is used to train the algorithm and produce the basis set of bright images. The test set is used to evaluate how good the algorithm performs on new data. If not differently specified, to extract the training and test sets, we shuffle the bright images, and we extract two random subsets with the desired sizes. In addition to a training and a test set, we also define a training and a test area. The training area correspond to the edge area discussed in the previous section, and it is used to find the basis of images. The test area is where the result of the algorithms is evaluated, and it usually corresponds to the area in the picture where the atoms are located. In Fig. 2.1b we show the training and test areas used in the following analysis. The training area is at the bottom part of the picture, whereas the test area is at the top.

To simplify the analysis of the quality of the algorithm, we initially neglect the atoms images. The idea is to use bright images in place of atoms images because we know the

expected optical density: zero everywhere, since there are no atoms. Any deviation from zero is purely noise. So we use the root-mean-square (RMS) optical density $\text{RMS}_{OD} = \sqrt{\langle OD^2 \rangle}$ as a benchmarking parameter. The lower the RMS, the lower the noise. Since we are dealing with multiple pictures with multiples pixels, we first calculate the RMS_{OD} for each picture, and then we average the result over the different pictures. To be precise we should then talk about *average* RMS, but in the following we will simply say RMS.

In a typical benchmarking session we therefore proceed in the following way.

1. We define a training and a test set of bright pictures B_{train} and B_{test} , and a training and test area.
2. We use the training set and the training area to find a basis of bright pictures V .
3. We construct the ideal bright pictures B'_{test} corresponding to the test set B_{test} using the basis V .
4. We calculate the OD between B'_{test} and B_{test}
5. We compute RMS_{OD} and use it as a benchmarking parameter

2.1.2 Optimization of the OFRA

We begin by characterizing the optimized fringe removal algorithm. In particular, we focus on how the RMS_{OD} is affected by the number of components that is kept from the diagonalization of the covariance matrix, the size of the training set and the size of the edge area.

Number of components

While introducing the OFRA in Section 1.2.1 we have anticipated that not all the components resulting from the diagonalization of the covariance matrix are necessary. We have already noticed in Fig. 2.1a that most of the components have small eigenvalues. We can identify a clear turning point in Fig. 2.1a, where the slope of the plot changes drastically, at around 5% of the images (10 components). It is therefore interesting to analyse if including more components benefits or worsens the final result.

The dataset was divided as explained above in 200 training images and 100 test images. The training set was used to create a basis of which only a fraction f was kept. The $200 \cdot f$ images of the basis were then used to reconstruct the images of the test set. Finally, the corresponding OD and RMS_{OD} were calculated. The result is shown in Fig. 2.2a. We can see that at around $f = 5\%$ the RMS_{OD} suddenly drops from around 0.055 to 0.051. Despite this drop being relatively small, it happens at the same point of the drop in the eigenvalues shown in Fig. 2.1a. It is therefore reasonable to limit the components to about 10% of the size of the original training set. This allows us to only consider the most important principal components, reducing the execution time, while keeping some margin from the turning point.

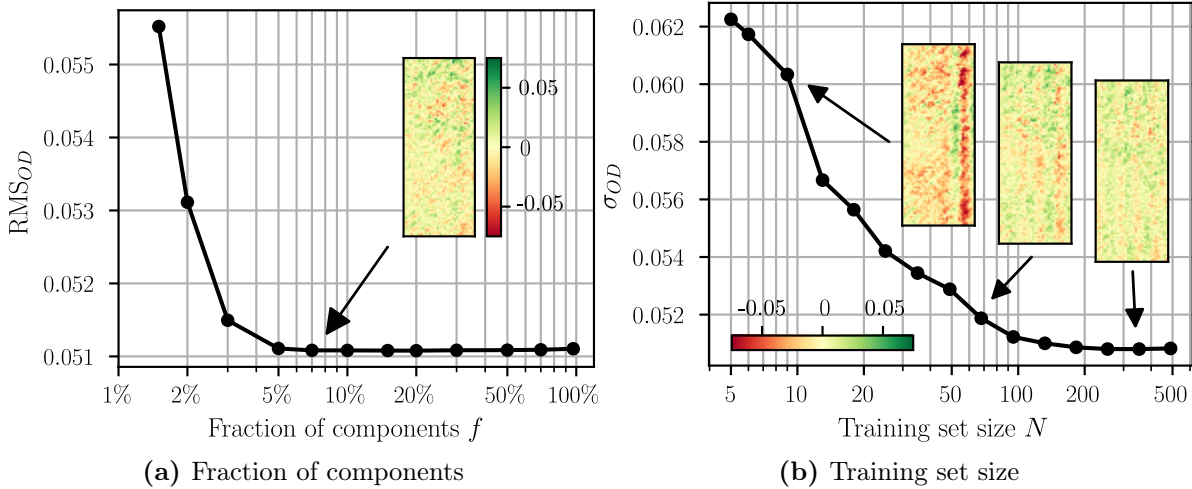


Figure 2.2: (a) RMS_{OD} for different number of components kept from the diagonalization of the covariance matrix. We observe a sudden drop in RMS_{OD} at around 5%, in agreement with what observed from the eigenvalues in Fig. 2.1a. (b) RMS_{OD} for different training set sizes. We observe a decrease in RMS_{OD} until a training set size of 100 images. After that, it stabilizes between 0.0051 and 0.0052. The insets show the OD calculated from the test set for specific points in the plots.

Training set size

We now focus on the size of the training set. It is reasonable to expect that the more images we use to construct the basis, the better the result will be. However, we also expect that there is a point where adding more images does not improve the result any more. In Fig. 2.2b we show the RMS_{OD} for different training set sizes N . We observe a decrease in RMS_{OD} until $N = 100$. After that, it stabilizes between 0.051 and 0.052. This suggests that the optimal training set size is around 100 images. After that, adding more images would require more acquisition time without improving the result.

Edge area

Another element that could affect the performance of the algorithm is the edge area. We therefore want to test if changing its size affects the result. In Section 2.1.2 we show the RMS_{OD} for different training areas. The training area is taken as a rectangle with the same width of the whole picture and a variable height, and the size is expressed as a fraction of area of the whole image. We observe that the RMS_{OD} decreases until a training area of 10% of the total area. After that, it stabilizes below 0.051. However, since taking bigger areas does not require more acquisition time, and the computational time to run the algorithm on 100-500 bright images is already in the order of seconds or minutes on a standard computer, it is reasonable to use an area as big as possible, verifying that there are no atoms.

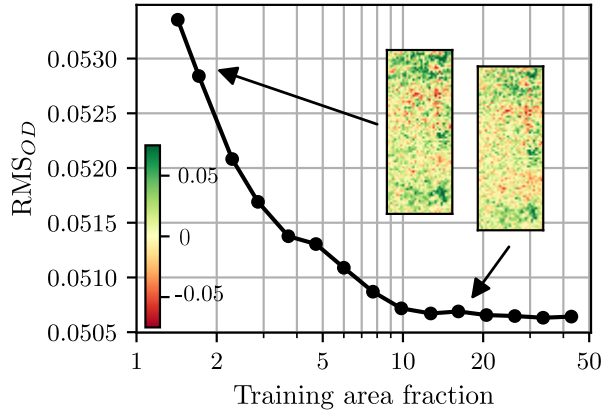


Figure 2.3: RMS_{OD} for different training area sizes. The size of the training area is expressed as a fraction of the total (test+training) area. The RMS seems to reach a stable point at a training area fraction of 20%.

Summary

In conclusion, we have proved that the OFRA is an effective way to reduce noise in absorption imaging data. We have seen how only around the first 10% of the components seem to be relevant for the reconstruction of the ideal bright image. However, using more does not worsen the result, so the decision of how many to consider should take into account the performance of the available hardware. The size of the training area has also an influence on the quality of the result, but except for a minimum threshold to obtain good results, there does not seem to be an upper bound, if not the one given by computational requirements. The parameter that seem to influence the result the most is the size of the training set. In fact, if for different training areas and principal components fraction the RMS_{OD} showed variations between 0.055 and 0.050, using small training set sizes ($N < 10$) results in RMS_{OD} higher than 0.060. Reducing this number to the limit of 0.051 requires more than 100 images. This has drawbacks not only on the computational time, which remains contained, but especially on the acquisition time, which is the most important limitation. In the next section we will try to see if the use of the ESFRA can overcome these limitations.

2.1.3 Comparison between the OFRA and the ESFRA

In the previous section we have seen how the OFRA is limited to an RMS_{OD} of around 0.051. We have also seen how this level of noise is reached using sets of at least 100 images. In this section we want to compare the OFRA to the ESFRA, and see if the latter can improve this result, either by reducing the noise or by requiring smaller training sets.

To characterize the ESFRA, we use the same procedure described in Sections 2.1.1 and 2.1.2. We train the algorithm on different dataset sizes, and, keeping only 10% of the principal components, we observe the behaviour of RMS_{OD} . The result is shown in Fig. 2.4 for shifts $d = 0, 1, 2$ pixels. The ESFRA with a translation of zero pixels is equivalent to the OFRA. The first thing we notice is that the performance with a translation of one or two pixels is




Figure 2.4: RMS_{OD} for different training set sizes for the OFRA and the ESFRA with different shifts d . The ESFRA with $d = 1$ is able to reduce the noise to 0.050 using 100 images. Shifts $d > 1$ do not seem to improve the result.

equivalent. This suggests that the vibrations of the imaging system cause translations smaller than one pixel. Therefore, we can limit ourselves to $d = 1$ when using the algorithm to reduce the computational resources. We also observe that the ESFRA is able to reduce the RMS_{OD} to 0.050, using 100 images. At the same time, to have the same noise level of the OFRA (0.051), we only need 50 images. It is also interesting to observe that for larger training sets, the noise increase again to 0.051. This behaviour is unexpected and can probably be attributed to some overfitting or minimization failure of the numerical algorithm. If we take into account the statistical significance of the result, we see from the error bars that we cannot claim that the difference in performance of the three algorithm is significant. Nevertheless, since the ESFRA consistently outperforms the OFRA on average, it is still reasonable to choose the former algorithm to post process the data.

In conclusion, the ESFRA seems to be able to improve the results of the OFRA. Using it, it is possible to achieve smaller RMS_{OD} using smaller training set sizes. Increasing the maximum displacement d over one does not seem to further improve the result. This is in agreement with what described by Song *et. al.* in their paper [12].

2.2 Origin of the residual noise

In the previous sections we have seen how we could reduce the noise in the optical density down to $RMS_{OD} \approx 0.050$. It is natural to ask ourselves if there is an intrinsic limitation to this value, and what it is limited by. We know that the pictures captured by the camera have some noise given both by the functioning of a CCD camera, like the readout noise, the clock-induced-charge (CIC) noise and the dark noise, and by the quantum nature of the detection process, i.e. the shot noise. For more details about the functioning and noise of CCD cameras one can refer to Mohan's thesis [22]. Recently, we substituted the old camera

Point Grey CMLN-13S2M-CS with a new iKon-M 934. As previously discussed by Kosucu in her thesis, the new camera has a very low CIC and dark noise, so that we can only consider the contribution of the shot and readout noise [15].

2.2.1 Introduction to CCD noise

Shot noise

Shot noise is an intrinsic phenomenon arising from the quantum nature of light and of the detection process. Since the electromagnetic field is quantized in photons, when detecting the intensity we are effectively counting the number of photons. A process of this kind follows a poissonian statistics, where the variance of the photon count σ_P^2 is equal to the average number of photons $\langle P \rangle$. In a CCD camera, the photons are converted into electrons with an efficiency, called quantum efficiency, η_{QE} . The quantum efficiency of the iKon camera at 780 nm is 93.8%. For comparison, the Point Grey camera has an efficiency of $\sim 25\%$. The poissonian statistics of the photon distribution is therefore inherited by the electrons distribution. The variance of the electrons produced in the process is then

$$\sigma_E^2 = \langle E \rangle = \eta_{QE} \langle P \rangle \quad (2.1)$$

Finally, the electrons are counted, and the signal is digitized with an analog-to-digital converted (ADC). The number of electrons needed to generate a count in the ADC (one analog-to-digital unit, or ADU) is called sensitivity s , so the number of counts is

$$\langle C \rangle = \langle E \rangle / s \quad (2.2)$$

with a variance

$$\sigma_{SN}^2 = \sigma_C^2 = \sigma_E^2 / s^2 = \langle E \rangle / s^2 = \langle C \rangle / s \quad (2.3)$$

Equation (2.3) is often used to measure the sensitivity of a camera. This is possible by looking at the variance of counts at different intensities and fitting $\langle C \rangle$ over σ_C^2 . For the iKon camera used with a pre-amplification gain of 4, the sensitivity reported by the manufacturer is 1.4 electrons/ADU. This was also confirmed experimentally by Kosucu in her work [15].

Readout noise

Another important source of noise is the readout noise. It is introduced by the ADC, and it has the characteristic of being Gaussian and independently distributed (white noise). From the manual of the iKon camera, we find that the readout noise for the camera operated under our conditions is $\sigma_{RD} = 11.0$ electrons / pixel, or 7.86 ADU / pixel.

Noise propagation in the OD

When measuring the RMS_{OD} for pictures without atoms as we did in the previous sections, the CCD noise provides a lower bound for the result, which can be found propagating the uncertainties in the calculation of the OD. When calculating the OD as

$$OD = -\log\left(\frac{A}{B'}\right) \quad (2.4)$$

we have an uncertainty

$$\sigma_{OD}^2 = \frac{B'^2}{A^2} \left[\frac{\sigma_A^2}{B'^2} + \frac{A^2 \sigma_{B'}^2}{B'^2 B'^2} \right] \simeq \frac{\sigma_A^2}{A^2} + \frac{\sigma_{B'}^2}{B'^2} \geq \frac{\sigma_A^2}{A^2} \quad (2.5)$$

where in the second step we have used the fact that the atoms and the reconstructed bright images should be on average the same, since there are no atoms. In the last step we also limited our attention to the uncertainty in the atoms image, since for an ideal fringe-removal algorithm the reconstructed bright image would be noise-free. If we only consider shot noise and readout noise, the variance in the OD is therefore limited by

$$\sigma_{OD}^2 \geq \frac{\sigma_A^2}{A^2} \geq \frac{\sigma_{SN}^2 + \sigma_{RD}^2}{A^2} = \frac{1}{A_s} + \frac{\sigma_{RD}^2}{A^2} \quad (2.6)$$

The previous relation holds for each pixel in the OD . We can find the limit of the RMS_{OD} by summing over all the pixels of the image

$$RMS_{OD} = \sqrt{\langle OD^2 \rangle_{\text{pixels}}} = \sqrt{\langle \sigma_{OD}^2 \rangle_{\text{pixels}}} \geq \sqrt{\left\langle \frac{1}{A_s} + \frac{\sigma_{RD}^2}{A^2} \right\rangle_{\text{pixels}}} \quad (2.7)$$

Note that if we did not apply a fringe-removal algorithm, but we calculated the optical density with the simple method, we should also include the noise of the bright image, which would result in an extra factor of $\sqrt{2}$.

2.2.2 Results

In order to verify how our algorithms are performing in comparison to the shot and CCD noise limit, we collected data at different powers for the imaging beam. The data was collected with the new iKon camera, so it should not be compared directly to the previous results. The power was modulated with an acousto-optic modulator (AOM). The data was collected on 18.12.2023, running the experiment 42 times for each imaging power. 35 images were used for training and 7 for testing. The results are shown in Fig. 2.5, where we displayed the RMS_{OD} obtained after applying the ESFRA ($d = 1$) with the CCD noise limit calculated as explained above. We include both the CCD noise limit neglecting the noise in the bright picture and including it. The two simply differ by a factor $\sqrt{2}$. For comparison, we also show the RMS_{OD} for the same test images with the corresponding bright images taken randomly from the dataset.

The data seems to agree well with the assumption that the RMS_{OD} is limited by the CCD noise of the atoms picture, at least for high imaging powers. For the lowest imaging power, the power is too low to give a reasonable result, so we will exclude it from the rest of the discussion. The noise in the RMS_{OD} is higher than the CCD noise limit (readout + shot noise) of the single atoms picture, but lower than the noise of the atoms and bright pictures together. This confirms that the fringe-removal algorithms we are employing not only remove the fringes from the image, but create an ideal bright image with lower noise. This could be attributed to neglecting the higher components in the PCA decomposition, which contain most of the shot noise. On the other hand, when calculating the optical density with random bright images, the resulting RMS_{OD} is higher than the atoms + bright picture CCD noise limit, as

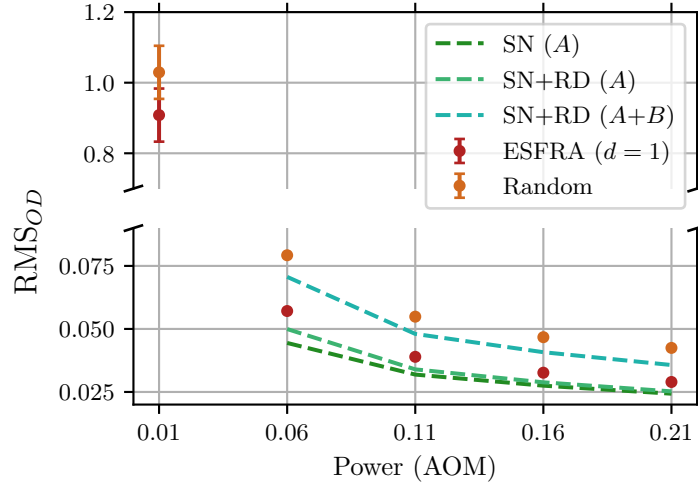


Figure 2.5: Root-mean-square optical density for images without atoms after applying the ESFRA ($d = 1$) compared to the CCD noise limit. For the atoms image (A), we show the shot noise (SN) and the shot noise + readout noise (SN + RD) separately. We also include the estimation of SN+RD for the case where the CCD noise of the bright image cannot be neglected ($A+B$). The error bars of the last four data points are smaller than the marker. The orange dots show the RMS_{OD} for the same test images with the corresponding bright images taken randomly from the dataset.

expected. It is also worth noticing that the scale on which the RMS_{OD} varies on the plot is much larger than what we observed when characterizing and optimizing the fringe removal algorithms in the previous sections (e.g. cfr. Fig. 2.4). On this scale, the performances of the OFRA and ESFRA look identical.

2.3 Effectiveness of the algorithms on real data

In the previous sections we have focused on the description and implementation of algorithms to reduce fringes in the optical density. To simplify the benchmarking, we decided to use pictures with a known optical density, namely zero. Despite this method having proved to be useful to optimize and compare the algorithms, the final goal is to use them on atoms. Similarly to what was done before, we will start by quickly explaining a new benchmarking method valid for non-zero optical densities, and then we will use it to characterize and compare the OFRA and the ESFRA with new data.

2.3.1 Benchmarking strategy

In the next sections we will assume to have collected a set of atoms, bright and dark images of sufficient size. The data will be processed with the fringe-removal algorithms described above, and the resulting images used to calculate the optical density. Now we need a way to measure the quality of the result.

The first benchmarking parameter is similar to what we defined before, namely the stan-

standard deviation of the optical density in a region of interest. The region of interest is defined as a square centred on the cloud of atoms. The key difference is that before we were calculating the RMS optical density $\text{RMS}_{OD}^{(i)}$ for each image i , and then we were averaging over all the images in the test set

$$\text{RMS}_{OD} = \langle \text{RMS}_{OD}^{(i)} \rangle_i \quad (2.8)$$

Now we are not interested in the variation inside a single image, but in the variation in the optical density among different images. The most convenient quantity to use is therefore the total variance

$$[\sigma_{OD}^{\text{tot}}]^2 = \sum_{x,y} [\sigma_{OD}^{(x,y)}]^2 \quad (2.9)$$

where $\sigma_{OD}^{(x,y)}$ is the standard deviation of each pixel (x,y) calculated over the different test images. We will usually express the results in terms of the square root of the above expression σ_{OD}^{tot} . This quantity gives us an idea of how different the optical densities are between each other. Since each of them comes from different runs of the experiment, we do not expect to have zero standard deviation. However, the more the fringe removal algorithm is effective, the more noise generated by fringes will be suppressed, reducing the total standard deviation.

Another interesting quantity to consider is the total number of atoms N , which is directly related to the integrated optical density

$$\mathcal{N} = \int OD(x, z) dx dz \quad (2.10)$$

by $N = \mathcal{N}/\sigma$, where σ is the absorption cross-section. When running the experiment multiple times, we are therefore interested in the mean $\langle \mathcal{N} \rangle$ and standard deviation $\sigma_{\mathcal{N}}$, which are proportional to the average number of atoms and its variation during different runs. We will also use the standard deviation of the mean

$$\sigma_{\langle \mathcal{N} \rangle} = \frac{\sigma_{\mathcal{N}}}{\sqrt{\# \text{ test images}}} \quad (2.11)$$

We expect the presence of fringes in the optical density to be able to affect both the average number of atoms and its standard deviation.

2.3.2 Results

To evaluate the performance of the algorithms, we used the data collected on 02.12.2022 and 04.12.2022, for a total of 320 atoms, bright and dark images. The test area was chosen as a square of side 120 pixels centred on the atoms, and the training area was a rectangle of 120×520 pixels just below the test area. The pictures were acquired with the Point Grey camera. The OFRA and ESFRA were run on this data with the parameters obtained in Section 2.1.2 and Section 2.1.3 for different training set sizes. For a given training set size k , the 320 bright images were divided in $320/k$ groups of k images. For each group, a basis was constructed and used to reconstruct the ideal bright images of the corresponding atoms images. After this process, for each value of k , we had 320 optical densities. For each of

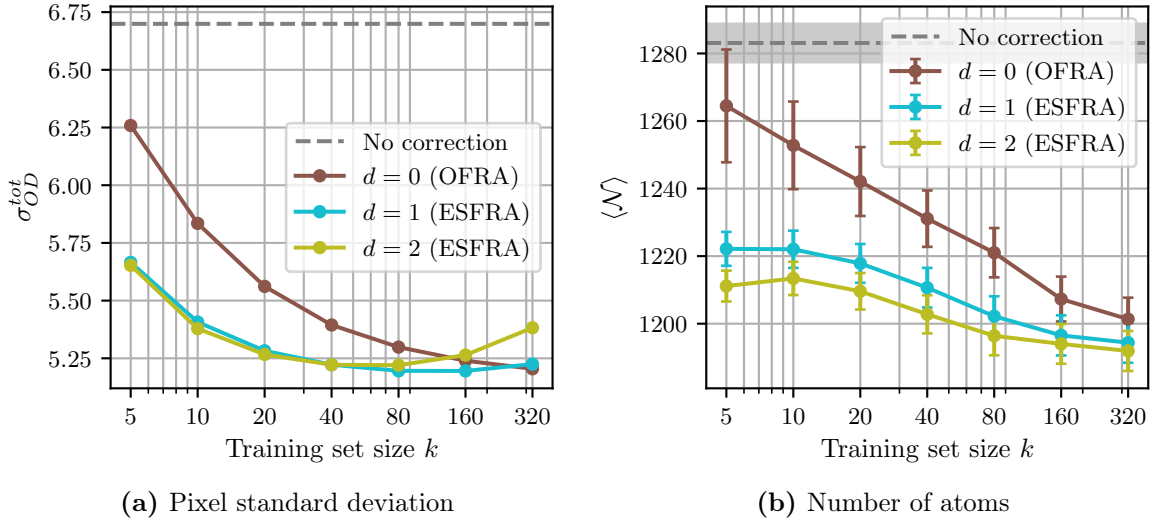


Figure 2.6: Benchmarking of the OFRA and the ESFRA on images with non-zero optical densities. The performance are compared to what would be obtained without running any fringe removal correction. In (a) we show the standard deviation of the optical density between different runs σ_{OD}^{tot} , in (b) the average integrated optical density $\langle \mathcal{N} \rangle$ and its standard deviation $\sigma_{\langle \mathcal{N} \rangle}$. The standard deviation is indicated with error bars for the OFRA and ESFRA and with the grey shaded region for the non-corrected data. Both plots suggest better performances for the ESFRA with bigger shifts d and bigger training set sizes k .

this set of 320 optical densities, we calculated σ_{OD}^{tot} , $\langle \mathcal{N} \rangle$ and $\sigma_{\langle \mathcal{N} \rangle}$. The results are shown in Fig. 2.6.

The plot in Fig. 2.6a is very similar to the one in Fig. 2.4, where we were showing the RMS_{OD} for pictures without atoms. Here we also observe an improvement of the standard deviation σ_{OD}^{tot} for the ESFRA compared to the OFRA for smaller training set sizes. In any case, both algorithms perform better than the non-corrected data. The results for the integrated OD shown in Fig. 2.6b are also interesting. From the plot, it seems that running the fringe-removal algorithm affects the average number of atoms $\langle \mathcal{N} \rangle$. In particular, the bigger the training set size and the more efficient the algorithm (bigger d), the less the number of atoms. This suggests that the fringes artificially increase the density. For big training sets, all three algorithms converge to the same average $\langle \mathcal{N} \rangle$, corroborating this hypothesis. The standard deviation $\sigma_{\langle \mathcal{N} \rangle}$ is also smaller for higher k and d , hinting again at a better performance of the ESFRA algorithm for higher d and bigger training set sizes.

Chapter 3

Identification of noise patterns in the optical density

“In the symphony of life, even noise has its own unique pattern, weaving unexpected melodies into the tapestry of existence.”

— ChatGPT

In this chapter we try to address the question of how to identify and reduce the noise in the optical density. The goal is to understand how some patterns in the optical density emerge even after the data has been processed with the fringe-removal algorithms. In the first section we will present a technique based on PCA proposed by Cao *et al.* [23] to isolate different noise components. In the second section we will try to explain the origin of a particular noise pattern, generated by the interference of the light diffracted by the cavity and by the atoms.

3.1 PCA for identification of noise patterns

In the previous chapter we have seen how to use PCA to generate a basis useful to implement fringe-removal algorithms. Now we want to use PCA to better understand the noise in the resulting optical densities. The basic idea is to perform PCA on a set of optical density images that have been calculated and corrected with the methods explained in the previous chapter. As already mentioned, PCA allows to decompose a set of data into components that capture its principal variations. The idea of Cao *et al.* [23] was to calculate these components for the optical density images and try to give a physical interpretation of them. This can help understand their origin, and therefore also suggest a way to limit their influence.

Analysis without corrections

To perform the analysis, we used the data of Section 2.3.2. We generated the ideal bright images with the ESFRA ($d = 1$) and calculated the optical density. This left us with a set of 320 optical densities. The optical densities were then cropped to squares of 100×100 px centred approximately around the atoms. The question was now how the different pictures

differed from each other. To address this question, we performed PCA on this set of images. In Fig. 3.1a we show the first four components with the corresponding explained variance. The explained variances of all the components add up to the total variance $[\sigma_{OD}^{\text{tot}}]^2$ as defined in Eq. (2.9). In the same plot we also show the explained variance fraction, defined as the explained variance of a single component divided by the total.

The first thing we notice is that only the first few components look relevant, with the first component explaining alone more than 15% of the total variance and the fourth less than 1%. While it is easy to give a physical interpretation of the first three components, the fourth one is a bit less intuitive. The first noise pattern is identical to the original density. This means that this component represents a deviation from the mean image where the density is rescaled by a constant factor. In other words, it represents a variation of the total number of atoms, which makes the cloud more or less dense. The second and the third components both represent a displacement of the centre of the cloud on two orthogonal directions. Any displacement in the 2D plane from the mean position can therefore be described in terms of these two components. It is important to emphasize that this displacement can be caused either by a movement of the camera or of the atoms themselves, and that it cannot be eliminated by the fringe-removal algorithms that we have seen in the previous chapter. Another possibility, investigated in Ref [24], is that these components indicate dipole excitations. In this case this would mean that the cloud oscillates back and forth excited by fluctuations of the trapping potential. We consider this interpretation less likely, since we are not deliberately exciting the BEC. The fourth component shows a variation of the density in the horizontal direction, where an increment (decrement) in the central region is compensated by a decrement (increment) on the left and on the right of the cloud. Similarly to what discussed for the second and third components, this component could also indicate a quadratic excitation of the system. Another possible physical origin of this pattern will be better explained in Section 3.2.

Analysis after corrections

Now that we know the main components that contribute to the variation of the optical density between different runs of the experiment, we can use this information to correct our data. The shift of the atoms can be easily corrected by translating the images. Following the work presented in Ref [23], the centre of each optical density is calculated as a weighted mean of the pixel intensities, with weights corresponding to the pixel positions. The image is then cropped to a 100×100 px square centred around the centre. In order to deal with non-integer centres, the intensity is linearly interpolated to estimate the correct value for each pixel. After centring the image a first time, the procedure is repeated three more times to let the process reach a stable point.

The result is shown in Fig. 3.1b, where we show the explained variance of the first four components and the total variance $[\sigma_{OD}^{\text{tot}}]^2$ before and after applying the correction. We notice that after two iterations, the shift correction decreases the total variance from 20 to almost 5. The main reduction is seen in the second component, the one representing a vertical displacement. More iterations of the correcting algorithm do not seem to further improve the result. We also notice that the other components are only slightly affected by this operation.

After performing the first correction, the remaining variance is completely dominated by

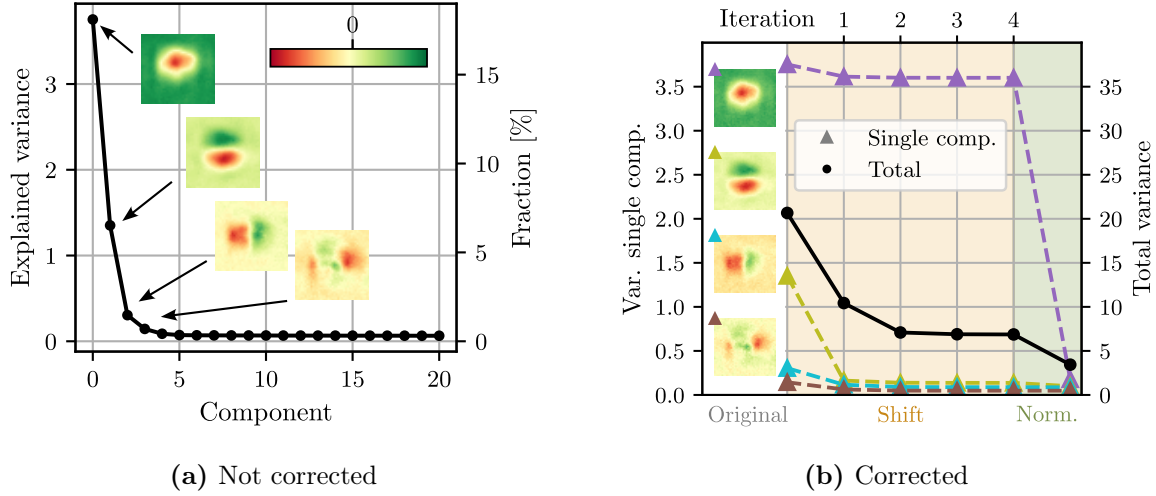


Figure 3.1: Principal component analysis of the optical density obtained from different runs of the experiment. In (a) we show the principal components with their explained variance for the uncorrected images. In (b) we analyse how the variance is decreased by applying successive corrections. In particular, the yellow area shows the result of correcting for the displacement of the atoms for different number of iterations of the centring algorithm. The green area shows the variance after renormalizing the density profile.

the first component, which can now explain around half of the total variance. In order to confirm the interpretation of this component that we offered before, i.e. a rescaling of the total density, we normalize each image by dividing for the integrated optical density, and we repeat the analysis. The result is also shown in Fig. 3.1b. After normalizing, the first component disappears completely, confirming the interpretation we gave. The total variance is also reduced by an amount equal to the contribution of this component (around 3.5). However, even if normalizing the data decreases the variance, a variation of the density is likely due to a real variation of the number of atoms in the experiment. This effect cannot be simply regarded as noise, and it should not be eliminated when processing the data.

In conclusion, simply centring the images, we were able to reduce the variance $[\sigma_{OD}^{tot}]^2$ by 75%. To achieve such a result we suggest iterating the centring algorithm twice, since it seems to be enough to reach a stable point. This correction has no influence on the total number \mathcal{N} . This is in agreement with the expectations, since a shift of the atoms region of the order of one pixel should not change the sum of the optical density contained into it by any considerable amount. This means that the correction we have presented cannot improve the resolution when counting the number of atoms in the cloud, but it can still be useful when we have to extract data from the shape of the cloud, for example when measuring the temperature or the condensate fraction, or if we are interested in investigating correlations in the density with a single-pixel resolution.




Figure 3.2: (a) Optical density averaged over 320 runs of the experiment. (b) Result of a simulation based on Fresnel propagation of a light beam passing through the cavity with a Gaussian absorber inside. The simulation is able to capture the rise of fringes at the edges of the atomic cloud. The scale of the simulated OD is arbitrary.

3.2 Cavity diffraction artefacts

Even after removing the fringes with the fringe removal algorithms and the aberrations in the optical density generated by a shift of the cloud, an interesting artefact remains visible in the optical density. This pattern appears when averaging over many images, such that the noise fluctuations cancel out. An example is shown in Fig. 3.2a, where we see the result of averaging the 320 optical densities centred with the algorithm presented above. We notice the presence of a periodic modulation of the optical density at the edges of the cloud in the horizontal direction. This pattern reminds us of the fourth component given by PCA in the previous section. We mentioned before how that component corresponds to an increase of the density at the centre of the cloud associated to a drop at the edges. This would be compatible with some diffraction mechanism originated by the atomic cloud, whose effect is also expected to scale with the density.

To further investigate this possibility, we performed a numerical simulation of our system. The simulation was performed applying Fourier propagation in the Fresnel approximation. More details about this computational method can be found in Ref [25]. Since the diffraction pattern was expected to be generated by the interference of the light scattered by the cavity and by the atoms, we tried to reproduce this experimental configuration numerically. The cavity acts like a thick slit that cuts the light both before and after the atoms (cfr. Fig. 1.1). To simulate this effect, we modelled it with two thin slits of width equal to the distance between the mirrors of the cavity. The atomic cloud was modelled as a partial Gaussian filter at the centre of the cavity. After the cavity, we reproduced the imaging system as a set of three lenses of focal length 300 mm, 400 mm and -100 mm as previously shown in Fig. 1.1. The distance between the lenses was measured directly. We measured a distance of 230 mm between the first two lenses, 320 mm between the second the third and 270 mm between the

third and the camera. We also measured a distance of 70 mm between the viewport on the vacuum chamber and the first lens. Knowing that the atoms are approximately at a distance of 250 mm from the viewport, we expected the distance between the atomic cloud and the first lens to be around 320 mm. Mathematical calculations with ray tracing methods suggest a distance of 304 mm to have a focused image.

The light of Gaussian beam of waist $300\ \mu\text{m}$ was propagated through this setup and the intensity at the end was recorded. Then, the same operation was done without the atoms Gaussian filter. The two resulting images are the equivalent of the atoms and bright images described in Chapter 1. The two images were then used to calculate the optical density shown in Fig. 3.2b. We can see how this simulation is able to capture the basic pattern of drops in the density that we observe in Fig. 3.2a. In some preliminary calculations, from the simulated results, we noticed that the diffraction pattern can affect the observed total number of atoms when integrating over a region around the atomic cloud. Moreover, given the non-uniformity of the variation in the OD , the choice of the exact dimension and shape of the region influences the result, since it might include or exclude more or less fringes. We were not able to understand neither theoretically nor numerically the exact influence of diffraction on the observed number of atoms, but this might be an interesting challenge for future work. It could also be possible to find a way to deconvolve the diffraction contribution to the density from the rest of the signal, either physically or with post-processing techniques, to clear the final image from these fringes. Possible solutions could be to image a slit on the cavity at the position of the atoms or to use spatially incoherent light.

Chapter 4

Conclusions

“I may not have gone where I intended to go, but I think I have ended up where I needed to be.”

— Douglas Adams, *The Long Dark Tea-Time of the Soul*

The OFRA and ESFRA have proved to be good fringe-removal algorithms candidates to be used in our experiment. We have shown how the ESFRA with a $d = 1$ shift seems to give the best compromise between acquisition time, execution time and residual noise. We also highlight the importance of using training set of sufficient sizes, at least greater than 50 and possibly above 100 images. At the same time, we observed how the residual noise in the optical density seems to be limited by the CCD camera noise, even after the camera has been upgraded with a new model to reduce it. Another interesting finding was that, when counting the number of atoms in the cloud, the presence of fringes does not only increase the variation of the count but also its average value. Since in our experiment counting how many atoms are in each momentum state is essential, a careful implementation of fringe-removal algorithms is important not only to decrease the error in the result, but also to get a result closer to the real value.

Some time was also spent in investigating the residual noise patterns in the optical density. We have identified the principal noise components and interpreted them as a variation of the number of atoms and a shift of the position of the cloud in the images. Although the former cannot be eliminated with post-processing method, the latter can be fixed by choosing an area centred around the atoms when analysing the data. We have also identified a noise pattern produced by the diffraction of the cavity and of the atoms, which results in a drop of the optical density at the horizontal edges of the cloud. Finding a way to remove or reduce this effect could be the subject of future work.

Acknowledgements

I would like to thank Prof. Tilman Esslinger and Dr. Tobias Donner for giving me the opportunity to continue working in the Quantum Optics Group. Another thanks goes to Panos for having thought about this project, that I have really enjoyed doing, and for the guidance that he offered me during its development. I would also like to thank the other members of the Cavity team: Rodrigo, Fabian, Nicola and Jakob for welcoming me in their group. I am looking forward to continuing my experience in the team.

Bibliography

- [1] Richard P. Feynman. “Simulating Physics with Computers”. In: *International Journal of Theoretical Physics* 21.6 (June 1, 1982), pp. 467–488. ISSN: 1572-9575. DOI: [10.1007/BF02650179](https://doi.org/10.1007/BF02650179) (cit. on p. 1).
- [2] Alexandre Blais et al. “Circuit Quantum Electrodynamics”. In: *Reviews of Modern Physics* 93.2 (May 19, 2021), p. 025005. ISSN: 0034-6861, 1539-0756. DOI: [10.1103/RevModPhys.93.025005](https://doi.org/10.1103/RevModPhys.93.025005) (cit. on p. 1).
- [3] Xiaoling Wu et al. “A Concise Review of Rydberg Atom Based Quantum Computation and Quantum Simulation*”. In: *Chinese Physics B* 30.2 (Feb. 2021), p. 020305. ISSN: 1674-1056. DOI: [10.1088/1674-1056/abd76f](https://doi.org/10.1088/1674-1056/abd76f) (cit. on p. 1).
- [4] Colin D. Bruzewicz et al. “Trapped-Ion Quantum Computing: Progress and Challenges”. In: *Applied Physics Reviews* 6.2 (May 29, 2019), p. 021314. ISSN: 1931-9401. DOI: [10.1063/1.5088164](https://doi.org/10.1063/1.5088164) (cit. on p. 1).
- [5] Immanuel Bloch, Jean Dalibard, and Sylvain Nascimbène. “Quantum Simulations with Ultracold Quantum Gases”. In: *Nature Physics* 8.4 (4 Apr. 2012), pp. 267–276. ISSN: 1745-2481. DOI: [10.1038/nphys2259](https://doi.org/10.1038/nphys2259) (cit. on p. 1).
- [6] Markus Greiner, Olaf Mandel, and Tilman Esslinger. “Quantum Phase Transition from a Superfluid to a Mott Insulator in a Gas of Ultracold Atoms”. In: 415 (2002) (cit. on p. 1).
- [7] Sebastian Krinner et al. “Observation of Quantized Conductance in Neutral Matter”. In: *Nature* 517.7532 (7532 Jan. 2015), pp. 64–67. ISSN: 1476-4687. DOI: [10.1038/nature14049](https://doi.org/10.1038/nature14049) (cit. on p. 1).
- [8] Renate Landig et al. “Quantum Phases from Competing Short- and Long-Range Interactions in an Optical Lattice”. In: *Nature* 532.7600 (7600 Apr. 2016), pp. 476–479. ISSN: 1476-4687. DOI: [10.1038/nature17409](https://doi.org/10.1038/nature17409) (cit. on p. 1).
- [9] Kristian Baumann et al. “Dicke Quantum Phase Transition with a Superfluid Gas in an Optical Cavity”. In: *Nature* 464.7293 (7293 Apr. 2010), pp. 1301–1306. ISSN: 1476-4687. DOI: [10.1038/nature09009](https://doi.org/10.1038/nature09009) (cit. on p. 1).
- [10] Fabian Finger et al. *Spin- and Momentum-Correlated Atom Pairs Mediated by Photon Exchange and Seeded by Vacuum Fluctuations*. Oct. 27, 2023. DOI: [10.48550/arXiv.2303.11326](https://doi.org/10.48550/arXiv.2303.11326). URL: <http://arxiv.org/abs/2303.11326> (visited on 01/29/2024). preprint (cit. on p. 1).

- [11] Linxiao Niu et al. “Optimized Fringe Removal Algorithm for Absorption Images”. In: *Applied Physics Letters* 113.14 (Oct. 1, 2018), p. 144103. ISSN: 0003-6951, 1077-3118. DOI: [10.1063/1.5040669](https://doi.org/10.1063/1.5040669) (cit. on pp. 2, 5, 8).
- [12] Bo Song et al. “Effective Statistical Fringe Removal Algorithm for High-Sensitivity Imaging of Ultracold Atoms”. In: *Physical Review Applied* 14.3 (Sept. 2, 2020), p. 034006. ISSN: 2331-7019. DOI: [10.1103/PhysRevApplied.14.034006](https://doi.org/10.1103/PhysRevApplied.14.034006) (cit. on pp. 2, 8, 11, 17).
- [13] M. R. Andrews et al. “Direct, Nondestructive Observation of a Bose Condensate”. In: *Science* 273.5271 (July 5, 1996), pp. 84–87. DOI: [10.1126/science.273.5271.84](https://doi.org/10.1126/science.273.5271.84) (cit. on p. 3).
- [14] Naaman Tammuz. “Thermodynamics of Ultracold 39K Atomic Bose Gases with Tuneable Interactions”. University of Cambridge, 2011 (cit. on p. 3).
- [15] Miray Kosucu. “Characterising and Improving the Detection Abilities of an Ultracold-Gas Experiment inside an Optical Cavity”. MA thesis. ETH Zurich, 2023. DOI: [10.3929/ethz-b-000637466](https://doi.org/10.3929/ethz-b-000637466) (cit. on pp. 3, 5, 18).
- [16] Lorenz Hruby. “Metastability and Quench Dynamics in a Long-Range Interacting Hubbard Model”. Doctoral Thesis. ETH Zurich, 2018. DOI: [10.3929/ethz-b-000308115](https://doi.org/10.3929/ethz-b-000308115) (cit. on p. 6).
- [17] Gaoyi Lei, Chencheng Tang, and Yueyang Zhai. “Fringe Removal Algorithms for Atomic Absorption Images: A Survey”. In: *Chinese Physics B* 31.5 (May 1, 2022), p. 050313. ISSN: 1674-1056. DOI: [10.1088/1674-1056/ac3758](https://doi.org/10.1088/1674-1056/ac3758) (cit. on p. 5).
- [18] C. F. Ockeloen et al. “Detection of Small Atom Numbers through Image Processing”. In: *Physical Review A* 82.6 (Dec. 16, 2010), p. 061606. DOI: [10.1103/PhysRevA.82.061606](https://doi.org/10.1103/PhysRevA.82.061606) (cit. on p. 5).
- [19] Gal Ness et al. “Single-Exposure Absorption Imaging of Ultracold Atoms Using Deep Learning”. In: *Physical Review Applied* 14.1 (July 6, 2020), p. 014011. DOI: [10.1103/PhysRevApplied.14.014011](https://doi.org/10.1103/PhysRevApplied.14.014011) (cit. on p. 5).
- [20] Feng Xiong, Yun Long, and Colin V. Parker. “Enhanced Principle Component Method for Fringe Removal in Cold Atom Images”. In: *JOSA B* 37.7 (July 1, 2020), pp. 2041–2044. ISSN: 1520-8540. DOI: [10.1364/JOSAB.391297](https://doi.org/10.1364/JOSAB.391297) (cit. on p. 8).
- [21] Hervé Abdi and Lynne J. Williams. “Principal Component Analysis”. In: *WIREs Computational Statistics* 2.4 (2010), pp. 433–459. ISSN: 1939-0068. DOI: [10.1002/wics.101](https://doi.org/10.1002/wics.101) (cit. on p. 8).
- [22] Jeffrey Mohan. “Reducing Noise in Thermodynamic Measurements on Interacting Fermi Gases”. ETH Zürich, July 23, 2018 (cit. on p. 17).
- [23] Shuyang Cao et al. “Extraction and Identification of Noise Patterns for Ultracold Atoms in an Optical Lattice”. In: *Optics Express* 27.9 (Apr. 29, 2019), p. 12710. ISSN: 1094-4087. DOI: [10.1364/OE.27.012710](https://doi.org/10.1364/OE.27.012710) (cit. on pp. 23, 24).
- [24] Romain Dubessy et al. “Imaging the Collective Excitations of an Ultracold Gas Using Statistical Correlations”. In: *New Journal of Physics* 16.12 (Dec. 19, 2014), p. 122001. ISSN: 1367-2630. DOI: [10.1088/1367-2630/16/12/122001](https://doi.org/10.1088/1367-2630/16/12/122001) (cit. on p. 24).

- [25] David G. Voelz. *Computational Fourier Optics: A MATLAB Tutorial*. 2011. ISBN: 978-0-8194-8205-1 (cit. on p. 26).

Declaration of originality

The signed declaration of originality is a component of every written paper or thesis authored during the course of studies. In consultation with the supervisor, one of the following three options must be selected:

- I confirm that I authored the work in question independently and in my own words, i.e. that no one helped me to author it. Suggestions from the supervisor regarding language and content are excepted. I used no generative artificial intelligence technologies¹.
- I confirm that I authored the work in question independently and in my own words, i.e. that no one helped me to author it. Suggestions from the supervisor regarding language and content are excepted. I used and cited generative artificial intelligence technologies².
- I confirm that I authored the work in question independently and in my own words, i.e. that no one helped me to author it. Suggestions from the supervisor regarding language and content are excepted. I used generative artificial intelligence technologies³. In consultation with the supervisor, I did not cite them.

Title of paper or thesis:

Noise reduction in absorption imaging for quantum gases experiments

Authored by:

If the work was compiled in a group, the names of all authors are required.

Last name(s):

Montalti

First name(s):

Nicolò

With my signature I confirm the following:

- I have adhered to the rules set out in the Citation Guide.
- I have documented all methods, data and processes truthfully and fully.
- I have mentioned all persons who were significant facilitators of the work.

I am aware that the work may be screened electronically for originality.

Place, date

Zurich, 16.04.2024

Signature(s)

Nicolò Montalti

If the work was compiled in a group, the names of all authors are required. Through their signatures they vouch jointly for the entire content of the written work.

¹ E.g. ChatGPT, DALL E 2, Google Bard

² E.g. ChatGPT, DALL E 2, Google Bard

³ E.g. ChatGPT, DALL E 2, Google Bard