



UPPSALA
UNIVERSITET

*Digital Comprehensive Summaries of Uppsala Dissertations
from the Faculty of Science and Technology 960*

Flash Diffractive Imaging in Three Dimensions

TOMAS EKEBERG



ACTA
UNIVERSITATIS
UPSALIENSIS
UPPSALA
2012

ISSN 1651-6214
ISBN 978-91-554-8439-2
urn:nbn:se:uu:diva-179643

Dissertation presented at Uppsala University to be publicly examined in B22, BMC, Husargatan 3, Uppsala, Friday, October 5, 2012 at 09:15 for the degree of Doctor of Philosophy. The examination will be conducted in English.

Abstract

Ekeberg, T. 2012. Flash Diffractive Imaging in Three Dimensions. Acta Universitatis Upsaliensis. *Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Science and Technology* 960. 68 pp. Uppsala. ISBN 978-91-554-8439-2.

During the last years we have seen the birth of free-electron lasers, a new type of light source ten billion times brighter than synchrotrons and able to produce pulses only a few femtoseconds long. One of the main motivations for building these multi-million dollar machines was the prospect of imaging biological samples such as proteins and viruses in 3D without the need for crystallization or staining. This thesis contains some of the first biological results from free-electron lasers.

These results include 2D images, both of whole cells and of the giant mimivirus and also contains a 3D density map of the mimivirus assembled from diffraction patterns from many virus particles. These are important proof-of-concept experiments but they also mark the point where free-electron lasers start to produce biologically relevant results. The most noteworthy of these results is the unexpectedly non-uniform density distribution of the internals of the mimivirus.

We also present Hawk, the only open-source software toolkit for analysing single particle diffraction data. The Uppsala-developed program suite supports a wide range of algorithms and takes advantage of Graphics Processing Units which makes it very computationally efficient.

Last, the problem introduced by structural variability in samples is discussed. This includes a description of the problem and how it can be overcome, and also how it could be turned into an advantage that allows us to image samples in all of their conformational states.

Keywords: X-ray, diffraction, mimivirus, three dimensional, phase retrieval, EMC, manifold embedding, CXI, FEL, free-electron laser, single particle

Tomas Ekeberg, Uppsala University, Department of Cell and Molecular Biology, Molecular biophysics, Box 596, SE-751 24 Uppsala, Sweden.

© Tomas Ekeberg 2012

ISSN 1651-6214

ISBN 978-91-554-8439-2

urn:nbn:se:uu:diva-179643 (<http://urn.kb.se/resolve?urn=urn:nbn:se:uu:diva-179643>)

To Maria

List of papers

This thesis is based on the following papers, which are referred to in the text by their Roman numerals.

- I **Ekeberg, T.** et al. (2012) Three-dimensional structure determination with an X-ray laser, *In preparation*
- II Seibert, M.¹, **Ekeberg, T.**¹, Maia, F. R. N. C.¹ et al. (2011) Single mimivirus particles intercepted and imaged with an X-ray laser, *Nature*, 470:78-U86
- III **Ekeberg, T.**, Maia, F. (2012) Data requirements for single-particle diffractive imaging, *In preparation*
- IV Maia, F., **Ekeberg, T.**, Timneanu, N., van der Spoel, D., Hajdu, J. (2009) Structural variability and the incoherent addition of scattered intensities in single-particle diffraction, *Physical review E*, 80(3):031905
- V Maia, F., **Ekeberg, T.**, van der Spoel, D., Hajdu, J. (2010) Hawk: the image reconstruction package for coherent X-ray diffractive imaging, *Journal of applied crystallography*, 43:1535-1539
- VI Seibert, M., Boutet, S., Svenda, M., **Ekeberg, T.** et al. (2010) Femtosecond diffractive imaging of biological cells, *Journal of physics b-atomic molecular and optical physics*, 43(19):194015

Reprints were made with permission from the publishers.

¹These authors contributed equally to the publication.

List of additional papers

- VII Koopmann, R. et al. (2012) In vivo protein crystallization opens new routes in structural biology, *Nature Methods* 9(3):259-U54
- VIII Johansson, L. C. et al. (2012) Lipidic phase membrane protein serial femtosecond crystallography, *Nature Methods*, 9(3):263-U59
- IX Barty, A. et al. (2012) Self-terminating diffraction gates femtosecond X-ray nanocrystallography measurements, *Nature Photonics*, 6(1):35-40
- X Lomb, L. et al. (2011) Radiation damage in protein serial femtosecond crystallography using an x-ray free-electron laser, *Physical Review B*, 84(21)
- XI Yoon, C. H. et al. (2011) Unsupervised classification of single-particle X ray diffraction snapshots by spectral clustering, *Optics Express*, 19(17):16542-9
- XII Martin, A. V. et al. (2011) Single particle imaging with soft X-rays at the Linac Coherent Light Source, *Advances in X-ray Free-Electron Lasers: Radiation Schemes, X-ray Optics, and Instrumentation*, 8078
- XIII Loh, N. D. et al. (2010) Cryptotomography: Reconstructing 3D Fourier Intensities from Randomly Oriented Single-Shot Diffraction Patterns, *Physical Review Letters*, 104:225501
- XIV Loh, N. D. et al. (2012) Fractal morphology, imaging and mass spectrometry of single aerosol particles in flight, *Nature*, 486(7404):513-517
- XV Martin, A. V., Morgan, A. M., **Ekeberg, T.** et al. (2012) The extraction of single-particle diffraction patterns from exposures of multiple particles, *Submitted*

List of abbreviations

AMO	Atomic, Molecular and Optical Sciences
CXI	Coherent X-ray Imaging
DESY	Deutsches Elektron Synchrotron
EM	Electron Microscopy
EMC	Expansion Maximization Compression
ER	Error Reduction
FEL	Free-electron laser
FIB	Focused ion beam
FLASH	Free-Electron Laser in Hamburg
fs	femtosecond
FWHM	Full Width at Half Maximum
GPU	Graphics Processing Unit
GTM	Generative Topographic Mapping
HHG	High Harmonic Generation
HIO	Hybrid Input Output
LCLS	Linac Coherent Light Source
MD	Molecular Dynamics
MTF	Modulation Transfer Function
OTF	Optical Transfer Function
PRTF	Phase Retrieval Transfer Function
RAAR	Relaxed Averaged Alternating Reflections
SASE	Self Amplified Stimulated Emission
SNR	Signal to Noise Ratio
SOM	Self Organizing Map
SVD	Singular Value Decomposition
XFEL	X-ray Free-Electron Laser

Mathematical conventions

\vec{x}	\vec{x} is a positional vector
\mathbf{x}	\mathbf{x} is a data vector
x^\dagger	Conjugate transpose of x
x^*	Complex conjugate of x
\hat{x}	Fourier transform of x
E	The identity matrix

Contents

List of abbreviations	ix
1 Introduction	13
2 X-ray lasers	15
2.1 Undulator radiation	15
2.2 Microbunching and SASE	15
2.3 Seeding	16
2.4 Optically driven X-ray lasers	17
3 Diffraction	18
3.1 Scattering from an inhomogeneous body	18
3.2 The Born approximation	18
3.3 The Fraunhofer approximation	19
3.4 Scattering factors	20
3.5 The Ewald sphere	20
4 Phase retrieval	22
4.1 The phase problem	22
4.2 Oversampling	22
4.3 Convex optimization	23
4.3.1 Error metrics	25
4.3.2 The Hybrid Input Output algorithm (HIO)	26
4.3.3 The Relaxed Averaged Alternating Reflections algorithm (RAAR)	27
4.3.4 Other algorithms	27
4.3.5 Additional constraints	27
4.3.6 The Phase Retrieval Transfer Function	28
4.4 Support recovery	29
4.4.1 The shrinkwrap algorithm	30
4.4.2 Variation on the shrinkwrap	31
4.5 The Hawk software package	31
4.6 Missing data	32
4.6.1 Mathematical description of the problem	33
4.6.2 Efficient calculation	33
4.6.3 Relation to noise	34
4.7 Reconstructing biological samples from experimental data: Imaging of cells on solid supports	35
5 The giant mimivirus	38
5.1 Background	38

5.2	Two-dimensional imaging of injected mimivirus particles	38
6	Aligning diffraction patterns	42
6.1	Common lines	42
6.2	Manifold embedding	44
6.3	Expansion Maximization Compression (EMC)	45
6.4	Estimating the required number of diffraction patterns	48
7	Experimental three-dimensional imaging of the mimivirus	51
8	Recovering conformations	55
8.1	Effects of structural variability on phase recovery	55
8.2	Extensions of alignment algorithms	56
9	Outlook	58
10	Sammanfattning på svenska	60
	References	64

1. Introduction

Seeing and imaging very small things have not only fascinated humans ever since we first opened our eyes, but is also the source of some of the greatest scientific breakthroughs. Discoveries such as the existence of cells, bacteria and DNA as well as our fundamental understanding of the particles that make up the universe all originate in this need.

How small things that are visible to us is fundamentally limited to half of the wavelength used to probe the sample and with visible light we can therefore not see things smaller than 200 nm. For this reason, X-rays with wavelength down to 1 Å are very attractive for imaging since this wavelength is short enough to see individual atoms. This has been utilized for more than 50 years in X-ray crystallography to determine the atomic structures of proteins and other macromolecules. Since X-rays interact very weakly with matter, many billion identical proteins are assembled to a crystal which enhances the scattered signal enough to permit structure determination.

X-ray crystallography has been an incredibly successful method with over 70000 structures solved. Its greatest weakness is however the requirement for the protein to be crystallized. The crystallization process is often hard and in many cases even impossible, which means that we are blind to the structure of any protein or macromolecule that can not be crystallized.

Free-electron lasers, a new type of light source might be the key to solve this problem. They produce X-ray pulses less than 100 femtoseconds long and with a peak brilliance 10 billion times higher than any previously existing X-ray source. This kind of intensity can produce strong enough scattering to allow imaging even from a single molecule. The extreme amounts of energy deposited in the molecule will rapidly turn it into a plasma and the sample will be destroyed, but the pulses delivered by the free-electron laser are however short enough to outrun this damage process. The protein will thus be damaged after the entire pulse has passed and the diffracted light will correspond to the undamaged protein. This principle is called *diffract and destroy* and was first proposed in [37].

In 2005, FLASH, the Free-electron LASer in Hamburg, started operation as the worlds first free-electron laser. FLASH only operates in the soft X-ray regime and it was mainly built as a test facility to develop the technique. Even so, it did open for users and has greatly advanced the field of coherent X-ray imaging[8].

Four years later, in 2009, the Linac Coherent Light Source (LCLS) at the Stanford Linear Accelerator Center (SLAC) opened to users. LCLS was the first hard X-ray source and has performed above expectations since, pushing many fields of science forward as a result[50][9].

In this thesis, I aim to present recent applications of both FLASH and LCLS that have been advancing our understanding of biology as well as the methods for coherent imaging. It includes a description of the difficulties of imaging in three dimensions and the methods to overcome them. Further, fundamental differences between single-particle and crystallographic techniques are discussed and how the single-particle techniques can make it possible to image a particle in all of its conformations.

2. X-ray lasers

Powerful coherent X-ray sources are very important in many areas of science. Synchrotrons have changed the world of structural biology and today we have access to free-electron lasers capable of producing pulsed radiation with a peak brilliance 10 billion times higher than that of synchrotrons.

2.1 Undulator radiation

An undulator is a periodic array of opposing magnets that will cause a stream of relativistic electrons to wiggle due to the Lorentz force imposed on the moving charges. The wiggling causes the electrons to emit radiation and each turn will cause more radiation to be emitted. Since the electrons travel at relativistic speeds they will co-propagate with the radiation and the intensity increases over the length of the undulator. Undulators are used as the means of creating radiation in both third-generation synchrotrons and free-electron lasers.

The wavelength, λ , of undulator radiation is a function of the undulator period λ_u and the magnetic field strength of the undulator, B_o . The relation is described by the *undulator equation*:

$$\lambda = \frac{\lambda_u}{2\gamma^2} \left(1 + \frac{1}{2}K^2 + (\gamma\theta)^2 \right) \quad (2.1)$$

where θ is the divergence angle, K is called the *undulator strength* and is defined as

$$K = \frac{\lambda_u e B_0}{2\pi m_e c} \quad (2.2)$$

and γ is the Lorentz factor from special relativity:

$$\gamma = \frac{1}{\sqrt{1 - v^2/c^2}} \quad (2.3)$$

2.2 Microbunching and SASE

In a continuous stream of relativistic electrons traveling through an undulator, the total intensity of the radiated field scales linearly with the number of electrons since there is no phase correlation between the fields emitted from the individual electrons. If

the electrons are distributed into *microbunches* spaced one wavelength apart they will however all radiate in phase and the positive interference will cause the intensity of the radiated field to scale as the square of the number of electrons.

The free-electron laser achieves microbunching through a process called SASE while the synchrotron does not use microbunching at all which is the reason for the 10^{10} fold difference in peak brilliance.

Microbunching can be achieved by introducing an electric field of the same wavelength as the wiggling of the electrons $\lambda_r = \frac{\lambda_u}{2\gamma^2} \left(1 + \frac{K^2}{2}\right)$. Electrons traveling together with the nodes of the field will be unaffected by it while electrons traveling out of phase will feel a ponderomotive force that drives them towards the nodes. The effect will create the microbunching required for the electrons to emit in phase.

The field emitted by the microbunches will add coherently to the field that is causing the bunching and amplify it. This feedback mechanism causes an exponential increase in bunching and emitted power. Because of this, even a very weak initial field can be enough to cause a strong bunching. In the free-electron lasers active today no external field is actually used, instead the self generated field from the first part of the undulator is amplified and causes the microbunching. This process is called *Self Amplified Stimulated Emission*, or SASE. The modes of the random radiation from the beginning of the undulator that is close to the wavelength λ_r will be amplified and cause microbunching. Several such modes can exist and the final spectrum of the pulse might thus contain several energy peaks.

Since the initial field is the product of a stochastic process, SASE radiation is characterized by a shot-to-shot intensity fluctuation and a small wavelength fluctuation. Since the overwhelming majority of the radiation is emitted when the electrons are distributed in microbunches the transverse coherence of the beam is very high which is attractive for imaging applications.

2.3 Seeding

To avoid the problems caused by the stochastic start of the SASE process, several techniques of seeding free-electron lasers are under development. Seeding here refers to some process which imposes a single mode that will then be amplified.

One promising technique is to apply an external field created by a high harmonic generation laser source, see section 2.4. This attempt was tested successfully at the Linac Coherent Light Source (LCLS) at SLAC.

Another proposed method is to seed the laser by bunching the electrons in advance. This method is called Enhanced SASE or ESASE.

Both current and future free-electron lasers will likely make use of seeding and thereby deliver more stable and more powerful X-ray pulses.

2.4 Optically driven X-ray lasers

Free-electron lasers are huge and expensive facilities and this is a strong motivation to develop cheaper and smaller alternatives. High harmonic generation (HHG) is a technique for generating pulsed x-rays that has been around for several decades. From a powerful laser pulse in the visible spectrum, higher harmonics can be generated by leading the pulse through for example a gas. Electrons in the gas will tunnel out from the atom and when the driving field reverses they will recombine. The energy of the released electron will be a multiple of the photon energy of the driving laser and thus the photon emitted when the electron recombines will have higher energy and be a multiple of the energy of the driving field.

HHG sources can produce equally short pulses as an FEL and of intensities about three orders of magnitude lower than from FELs. They can not currently reach sub nanometer wavelengths and are therefore not a viable alternative for many biological applications. If this changes it could however be possible to have machines with FEL like specifications in the basement of a lab, the development of HHG sources is therefore of great importance for structural biology.

3. Diffraction

3.1 Scattering from an inhomogeneous body

Diffraction from an inhomogeneous medium, assuming monochromatic light, is described by the following equation.

$$U^{(s)}(\vec{r}) = \int \rho(\vec{r}') U(\vec{r}') \frac{e^{i2\pi s|\vec{r}-\vec{r}'|}}{|\vec{r}-\vec{r}'|} d\vec{r}' \quad (3.1)$$

Where $U^{(s)}(\vec{r})$ is the scattered field at point \vec{r} , $U(\vec{r})$ describes the field in the medium and $s = \frac{1}{\lambda}$. ρ is the scattering potential defined as

$$\rho(\vec{r}, \omega) = \pi s^2 (n^2(\vec{r}, \omega) - 1) \quad (3.2)$$

where n is the refractive index of the material.

A full derivation of this formula can be found in [5].

3.2 The Born approximation

In equation 3.1 the variable $U(\vec{r})$ in the integrand is the total field inside the object, i.e. $U(\vec{r}) = U^{(i)}(\vec{r}) + U^{(s)}(\vec{r})$, the total field is the sum of the incoming and the scattered fields. Since the scattered field is present both inside and outside of the integral, this equation is very hard to use for predicting diffraction. When the scattered wave is much weaker than the incoming wave we can however make the simplification $U(\vec{r}) \approx U^{(i)}(\vec{r})$ and thus remove $U^{(s)}$ from the integrand. Equation 3.1 then takes the form

$$U^{(s)}(\vec{r}) = \int \rho(\vec{r}') U^{(i)}(\vec{r}') \frac{e^{2\pi i s|\vec{r}-\vec{r}'|}}{|\vec{r}-\vec{r}'|} d\vec{r}' \quad (3.3)$$

This simplification is called the first-order Born approximation[5] and greatly simplifies the equation. The approximation usually holds for viruses and smaller objects at X-ray wavelengths and will be assumed for all the examples of diffractive imaging in this work.

3.3 The Fraunhofer approximation

We will now assume that the extent of the diffracting object is much smaller than the distance to the point \vec{r} . This will allow us to further simplify equation 3.3 into a form which is very easy to work with. The region where this assumption is true is called the *far field*. We start with defining the *Fresnel number* F as

$$F = \frac{a^2}{r\lambda} \quad (3.4)$$

where a is the extent of the scatterer and r is the length of \vec{r} . When $F \ll 1$ we are in the far field[5]. That means that we can rewrite the last term in the integrand of equation 3.3.

$$\frac{e^{2\pi i \vec{s} \cdot (\vec{r} - \vec{r}')}}{|\vec{r} - \vec{r}'|} \approx \frac{e^{2\pi i s r}}{r} e^{-2\pi i \vec{s}_{\text{out}} \cdot \vec{r}'} \quad (3.5)$$

where \vec{s}_{out} is a vector pointing in the direction of the outgoing light and of length $\frac{1}{\lambda}$. We have used the fact that the denominator is dominated by \vec{r} and we have also rewritten the exponent:

$$|\vec{r} - \vec{r}'| \approx r - \vec{s}_{\text{out}} \cdot \vec{r}' \quad (3.6)$$

Substituting this into equation 3.3 then gives

$$U^{(s)}(\vec{r}) = \frac{e^{2\pi i s r}}{r} \int \rho(\vec{r}') U^{(i)}(\vec{r}') e^{-2\pi i \vec{s}_{\text{out}} \cdot \vec{r}'} d\vec{r}' \quad (3.7)$$

We now assume $U^{(i)}$ to be a plane wave of amplitude U_0 and directionality given by \vec{s}_{in} :

$$U^{(i)}(\vec{r}') = U_0^{(i)} e^{2\pi i \vec{s}_{\text{in}} \cdot \vec{r}'} \quad (3.8)$$

Equation 3.7 then becomes

$$U^{(s)}(\vec{r}) = \frac{e^{2\pi i s r}}{r} \int \rho(\vec{r}') U_0^{(i)} e^{-2\pi i (\vec{s}_{\text{out}} - \vec{s}_{\text{in}}) \cdot \vec{r}'} d\vec{r}' = U_0^{(i)} \frac{e^{2\pi i s r}}{r} \int \rho(\vec{r}') e^{-2\pi i \vec{s} \cdot \vec{r}'} d\vec{r}' \quad (3.9)$$

where $\vec{s} = \vec{s}_{\text{out}} - \vec{s}_{\text{in}}$.

This is a really interesting result since equation 3.9 is similar to the well known Fourier transform[48].

$$\hat{f}(\vec{s}) = \mathfrak{F}\{f(\vec{r})\}(\vec{s}) = \int f(\vec{r}) e^{-2\pi i \vec{s} \cdot \vec{r}} d\vec{r} \quad (3.10)$$

This leads to an important conclusion.

In the far field, the scattering of a plane wave is proportional to the Fourier transform of the scattering potential evaluated at the vector \vec{s} .

This is expressed in the following equation.

$$U^{(s)}(\vec{s}) = U_0^{(i)} \frac{e^{i k r}}{r} \mathfrak{F}\{\rho(\vec{r})\}(\vec{s}) \quad (3.11)$$

In the rest of this thesis we will be assuming that equation 3.11 accurately describes diffraction. This implies that we are assuming the Born approximation, that we measure the diffraction data in the far field and that incoming waves are plane waves.

3.4 Scattering factors

So far we have described the diffracting material in terms of its refractive index n . Another common way to describe the same properties is by the *scattering factors* of the material[13]. They relate to the diffractive index of the material as

$$n = 1 - \frac{1}{2\pi} N r_0 \lambda^2 (f_1 + i f_2) \quad (3.12)$$

where r_0 is the classical electron-radius and $f = f_1 + i f_2$ is the scattering factor. The scattering factor describes the scattering from a single atom in relation to the scattering from a free electron. This way of quantifying diffraction is therefore more useful when working with atomic positions, such as protein structures, or in general when the atomic composition is known. Tables of scattering factors exist for most elements at many X-ray wavelengths[24].

The real part of the scattering factor, f_1 , describes the strength of scattering while the complex part of the scattering factor, f_2 , describes the absorption of the material. f_2 can be related to the attenuation length of the material in the following way.

$$\mu = 2n r_0 \lambda f_2 \quad (3.13)$$

where n is the number of atoms in a unit volume. The attenuation length is the distance into a material where the intensity of an incoming wave has dropped to half.

3.5 The Ewald sphere

The Fourier transform of a three-dimensional object is also three dimensional, but the diffracted signal is only two dimensional. The part of three-dimensional Fourier space that is sampled by a diffraction experiments is given by the vector $\vec{s} = \vec{s}_{\text{out}} - \vec{s}_{\text{in}}$ introduced in equation 3.9.

Since \vec{s}_{in} is constant and \vec{s}_{out} is of fixed length, the vector \vec{s} will cover a sphere in diffraction space (see figure 3.1) and the sphere will intersect the origin at forward scattering. This sphere is called the *Ewald sphere*[19]. In the experiments in this thesis, only a small part of the diffraction angles around forward scattering is sampled by the detector. In that case, this section of the sphere can be approximated as a plane that cuts the Fourier space through the origin.

A back Fourier-transform of the diffracted wave will then, according to the Fourier-slice theorem[7], give a projection of the samples scattering potential.




Figure 3.1. The momentum transfer $\vec{s} = \vec{s}_{\text{out}} - \vec{s}_{\text{in}}$ will always reside on a sphere of radius $\frac{1}{\lambda}$ that intersects the origin. This sphere is called the Ewald sphere.

4. Phase retrieval

4.1 The phase problem

As described in section 3.3, the diffracted wavefront can often be described as the Fourier transform of the scattering potential of the object, sampled at the Ewald sphere. The value of the Fourier transform is a complex number and the complex amplitude corresponds to the amplitude of the electromagnetic wave and the complex argument (or phase) corresponds to the phase shift of the wave.

In a flat Ewald geometry, i.e. when the scattering angle is small, we can determine the projection image of the sample from the scattered wave. This is done by performing a back Fourier transform of the scattered wave. It is however impossible to directly measure the phase of an X-ray wavefront, the detector only records the intensity of the wave which is given by the square of the amplitude.

$$I(\vec{s}) = A(\vec{s})^2 = \hat{\rho}(\vec{s}) \hat{\rho}(\vec{s})^* \quad (4.1)$$

where $\hat{\rho}(\vec{s})$ is the Fourier transform of the scattering potential $\rho(\vec{r})$ and $A(\vec{s})$ is the amplitude of the scattered wave:

$$A(\vec{s}) = |\hat{\rho}(\vec{s})| \quad (4.2)$$

Since we don't know the phases, there is in general no direct way of recovering an image of the object. This is called the phase problem and is a well known issue in diffractive imaging. To solve it, some additional information is always required to compensate for the information lost with the phases.

4.2 Oversampling

One of the strengths of single-particle diffractive-imaging is that there is an appealing solution to the phase problem. The additional information in this case, is that the extent of the particle is limited. In this section we will outline how this extra information solves the phase problem and quantify how much information is needed.

The Shannon sampling theorem (sometimes called the Nyquist-Shannon sampling-theorem) states that *A function that contains no frequencies higher than B, is completely determined by its values at a series of points spaced $\frac{1}{2B}$ apart.*[45]

In single particle imaging, the diffracted signal is band-limited since the particle has a finite size. Let us label the particle diameter d_p , the band limit will then be $d_p/2$. This implies, according to Shannon, that the diffracted signal will contain no frequency higher than $d_p/2$ and can be fully described by its value at points spaced $1/d_p$ apart. This distance defines a critical pixel density required to recover the object if we were detecting the phases as well. Since our samplings doesn't contain the phases, a $\frac{1}{d_p}$ sampling is however not enough to recover the sample. If we can sample at twice as many points we will however in general have collected enough additional information to compensate for the lack of phases.

This was realized by David Sayre[43] in 1952 and although at his time the applications were rather artificial, it has gained immense importance in single particle imaging today where oversampling is easily achieved simply by making sure that the pixels of the detector cover a small enough angle.

This method, called the *oversampling method* is used in almost all applications of single particle X-ray imaging today. While the method describes the requirements for phase recovery and thus the necessary experimental conditions, it only states that recovery is theoretically possible and doesn't outline the way to do it. The most common way to solve it in practice is by a family of optimization techniques called *convex optimization*. They are the subject of the next section.

4.3 Convex optimization

The phase problem in single particle imaging can be described as a search for an object that fulfills two constraints. The first constraint is given by the experimental data, the amplitudes of the Fourier transform of the object should match the square root of the measured intensities. This constraint is referred to the *Fourier-space constraint*. The second constraint is given by the oversampling discussed above. In real space, this constraint is simply a limit to the size of the particle. The real space can therefore be split in two regions, one region where the object is allowed to have density and one region that must be empty. The part that allows density is usually called the object's *support* and the constraint imposed by it is called the *real-space constraint*.

Figure 4.1 shows a schematic drawing of the set of all objects fulfilling the Fourier-space constraint together with the set of all objects fulfilling the real-space constraint. The solution will fulfill both of the constraints and is found where the two sets intersect. The figure also hints at a simple technique to find the solution. We note that projecting onto one of the sets in the figure will always bring us closer to the solution and interchangeably projecting on the two sets will eventually bring us all the way to the solution. This optimization technique is called *error reduction* (ER) and comes from a family of techniques called *convex optimization*.




Figure 4.1. The solution to the phase problem has to fulfill both the real-space constraint and the Fourier-space constraint. The right figure shows how repeated projections onto the two sets will bring us to the solution.

The projections on the Fourier-space constraint (P_f) and on the real-space constraint (P_r) used above can be expressed mathematically as

$$P_f \rho(\vec{r}) = \mathfrak{F}^{-1} \sqrt{I(\vec{s})} e^{-i \arg(\mathfrak{F} \rho(\vec{r}))} \quad (4.3)$$

$$P_r \rho(\vec{r}) = \begin{cases} \rho(\vec{r}) & \text{if } \vec{r} \in S \\ 0 & \text{if } \vec{r} \notin S \end{cases} \quad (4.4)$$

where S is the support and $I(\vec{s})$ are the experimentally measured intensities.

In terms of these projectors, one iteration of the ER algorithm is given by

$$\rho_{n+1}(\vec{r}) = P_r P_f \rho_n(\vec{r}) \quad (4.5)$$

i.e. first projecting on the Fourier-space constraint followed by a projection on the real-space constraint.

The seed for the algorithm is normally an object created by applying random phases to the experimental amplitudes.

The following list explains the implementation of the algorithm in plain text.

1. Assign a random phase to every pixel of the diffraction image.
2. Inverse Fourier transform the pattern.
3. Set all pixels outside the support to zero.
4. Fourier transform.
5. Replace the amplitudes with the experimentally measured amplitudes, but keep the phases.
6. Repeat from step 2

A third way to describe the algorithm is by describing the effect of one iteration on real space.

$$\rho_{n+1}(\vec{r}) = \begin{cases} P_f \rho_n(\vec{r}) & \text{if } \vec{r} \in S \\ 0 & \text{if } \vec{r} \notin S \end{cases} \quad (4.6)$$

The pixels in the support are replaced with the last model after applying the Fourier-space constraint while the pixels outside of the support are set to zero. The above list and equations 4.5 and 4.6 are all different ways to think of the same algorithm.

We note that the algorithm will only work if the two sets are both convex since that is what guarantees that the projections will always bring us closer to the solution. If the sets are non-convex, local minima might exist that will be stable points for the algorithm.

Figure 4.1 shows the two sets of the constraints as blobs but what do they actually look like, and are they convex? Before looking into this, we need the definition of a convex set, which is given by: *a set is convex if for any two points in the set, all points along a straight line between them (called the convex combinations of the points) are also in the set.*[42]

Objects fulfilling the real-space constraint can all have different values along the pixels inside the support while the pixels outside the support are all zero. Any convex combination of two such objects will obviously still have only zeros outside the support and will thus still fulfill the real-space constraint. The set is therefore convex and has the shape of a hyperplane.

Now consider two different objects fulfilling the Fourier-space constraint, they both have the same amplitudes but different phases. Any convex combination of these objects will have a lower amplitude, and since the amplitude is different, the convex combination will not fulfill the Fourier-space constraint. The set of objects fulfilling the Fourier-space constraint is therefore not convex.

The fact that the Fourier-space constraint is not convex implies that the ER algorithm might not work very well since it is unable to escape local minima. Later in this section we will describe modifications of the ER algorithm that partially solves this issue.

4.3.1 Error metrics

To monitor how well the current model complies with both sets we can measure the distance from it to the respective set. This introduces two error metrics, the *real-space error*, E_r , and the *Fourier-space error*, E_f [30].

The real-space error is defined as

$$E_r = |P_r \rho - \rho| = \left(\sum_i \rho_i^2 \right)^{\frac{1}{2}} \quad (4.7)$$

The Fourier-space error is defined as

$$E_f = |P_f \rho - \rho| = \left(\sum_i (A_i - \hat{\rho}_i)^2 \right)^{\frac{1}{2}} \quad (4.8)$$

The error associated with a certain set will obviously be zero after projecting on that set. Therefore the real-space error is calculated after applying the Fourier-constraint and the Fourier-space error is calculated after applying the real-space constraint. We also note that at a perfect solution both of these errors will be zero. Since noise is usually present in diffraction data, perfect solutions are however rare which means that the two sets don't really intersect. We will then simply be searching for the point that minimizes the two errors metrics.

The name of the error-reduction algorithm actually refers to the fact that every iteration will reduce both of these errors with every iteration. This feature is also what makes it impossible for it to escape from a local minimum. We will now take a look at some algorithms that overcome this drawback and are therefore more useful to real-world problems.

4.3.2 The Hybrid Input Output algorithm (HIO)

The *Hybrid Input Output* (HIO) algorithm[21] has become one of the most used algorithms for single particle phase recovery. The workflow is similar to that of the error-reduction algorithm with the only difference being a different real-space projection. The update is instead

$$\rho_{n+1}(\vec{r}) = \begin{cases} P_f \rho_n(\vec{r}) & \text{if } \vec{r} \in S \\ \rho_n(\vec{r}) - \beta P_f \rho_n(\vec{r}) & \text{if } \vec{r} \notin S \end{cases} \quad (4.9)$$

The reason for the change was to speed up the usually very slow error-reduction algorithm by means similar to a negative feedback. In addition to making it faster than the ER algorithm this also gives it the ability to escape from local minima.

The parameter β in the algorithm determines how much negative feedback is added and should be in the range $[0, 1]$. A lower feedback gives a slow algorithm that takes small steps but is less likely to miss a minimum but also more likely to spend much time in local minima. A higher β gives a fast algorithm that can quickly cover large parts of the search space and is more likely to not even notice local minima. The downside of a high β is the risk of escaping or totally missing even the global minimum.

HIO has the slightly counter-intuitive property that if the current minimum is not perfect, i.e. the two sets don't intersect, it will eventually escape[35]. This means that local minima are less of a problem, however, in the presence of noise, even the true solution will not be perfect and the algorithm will, given enough time, escape from it. A common solution to this problem is to introduce a threshold on the error metrics and stop the algorithm when it is reached, then finish with a few iterations of error reduction to refine the solution to the best point of the current minimum. Another solution is to run the algorithm for a long time and afterwards select the iterate that had the lowest errors as the starting point for the ER refinement.

4.3.3 The Relaxed Averaged Alternating Reflections algorithm (RAAR)

Another common phasing algorithm is the Relaxed Averaged Alternating Reflectors (RAAR) algorithm [29]. Together with HIO and ER it makes up more than 90% of the author's usage.

Its behavior can be described as an intermediate between error reduction and HIO. It doesn't escape all minima, but can escape shallower ones. When the data is of high quality it usually finds the solution much faster than HIO and stays in the true minimum but if the true minimum is only slightly deeper than some local minimum it will struggle much more than HIO.

The update is described by the following equation, where β plays a similar role as for HIO.

$$\rho_{n+1}(\vec{r}) = \begin{cases} P_f \rho_n(\vec{r}) & \text{if } \vec{r} \in S \text{ and } \rho(\vec{r}) \geq -(1 + \beta) P_f \rho(\vec{r}) \\ \beta \rho_n(\vec{r}) - (1 - 2\beta) P_f \rho_n(\vec{r}) & \text{otherwise} \end{cases} \quad (4.10)$$

4.3.4 Other algorithms

Many other phasing algorithms exist and are in use. Most of them are also based on the idea of convex optimization and, just like HIO and RAAR, also differ from ER only in how the real-space constraint is applied. The following are three of the more common ones.

- Difference map[17]
- Saddlepoint optimization[33]
- Hybrid projection reflection[2]

A few algorithms are based on slightly different ideas. One example is the charge flipping algorithm that doesn't use a support but instead treats low-density regions in real-space different from high-density regions[38].

Another algorithm that doesn't use a support is called Espresso[31]. It assumes that the object being recovered is sparse and the real-space constraint is therefore replaced with a sparsifying operation. A generalization of the algorithm also works for objects that can be transformed to become sparse, such as an object having several patches of flat densities.

4.3.5 Additional constraints

The only information about the sample that we have used so far is the size limit. Sometimes we do however know more and this extra information can really aid the

algorithm. The two most common such extra constraints are the *reality constraint* and *positivity constraint*.

The real part of the scattering factor is often much larger than the imaginary part and the object can, in such cases, often be assumed to be purely real. This can then be enforced, usually by simply setting the complex component to zero after applying the real-space constraint.

One caveat here is that the algorithm becomes sensitive to badly centered diffraction patterns since a translated pattern will result in the object being multiplied with a phase ramp. A remedy for this when the misalignment is small is to replace the reality constraint with a *phase ramp constraint*. We then project the phases onto the best-fitting ramp instead of projecting them to the real axis. This method can even handle cases where the misalignment is less than one pixel. If the phase shift across the object is larger than a full turn the problem of fitting the ramp becomes computationally much harder, and the implementation of the phase ramp constraint in Hawk (see section 4.5) therefore assumes the ramp to be linear, thus requiring the misalignment to be small.

The positivity constraint works in the same way as the reality constraint and since negative scattering factors don't exist it can usually be enforced. Only when the model of diffraction doesn't hold, for example when the Born approximation is not satisfied, is it sensible not to use it. Then, the faults of the model might manifest themselves as negative densities and a positivity constraint might make the reconstruction harder and the result worse.

4.3.6 The Phase Retrieval Transfer Function

In optical systems a *transfer function* (OTF) is the Fourier transform of the point-spread function which is the shape a delta function obtains after passing through the system. In other words, an image will be convoluted by the point-spread function when passing through the system. Of particular usefulness is the modulation transfer function (MTF) which is the absolute value of the transfer function. The radial average of this function is often used to indicate the resolution of the system, where values close to one means that the image is preserved at that resolution and low values means that the image is distorted at that resolution.

The reconstruction process can also be thought of as such an optical system and its effect on the reconstructed image can be described with a transfer function called the *Phase Retrieval Transfer Function* (PRTF)¹.

To calculate the PRTF, multiple reconstructions are performed using the same parameters but with different randomly chosen starting points. The PRTF is then calculated

¹More precisely, what is usually referred to as the PRTF has an equivalent meaning as the MTF. The corresponding OTF can also be calculated but is very rarely used

by averaging the Fourier images and then normalizing with the amplitude:

$$\hat{f}_{\text{average}} = \frac{\sum_n \hat{f}_{ni}}{N} \quad (4.11)$$

$$\text{PRTF}_i = \left| \frac{\hat{f}_{\text{average}}}{|\hat{f}_i|} \right| \quad (4.12)$$

Where N is the number of repeats, \hat{f}_{ni} is the recovered value of the pixel i in Fourier space at repeat n and the absolute value $|\hat{f}_i|$ is the same for all repeats since it is given by the experimental intensity. Just like the MTF, the PRTF describes how the Fourier amplitudes have decreased, in this case from the averaging of phased diffraction-patterns.

The back Fourier-transform of \hat{f}_{average} gives the average real-space reconstruction, which is much more reliable than any individual reconstruction since the effects of the starting phases have been washed out and only the features that are supported by the data are left.

The phase retrieval is normally reliable for where we have strong signal and less reliable where the signal is weak. Since the diffracted signal tends to be stronger in the central, low resolution, region we expect a PRTF with high values at low q that drops towards higher q . One common way to quantify the resolution is to threshold the radial average of the PRTF (also often referred to as the PRTF). The point where it first drops below the threshold defines the resolution. In CXI, a threshold of e^{-1} is commonly used but values ranging from 0.1 to 0.5 are also found. In this thesis we have exclusively used the e^{-1} convention.

The resolution is usually expressed as the *full-period resolution* which is the inverse of the amplitude of the Fourier-space coordinate \vec{s} where the PRTF drops below the threshold. Full-period resolution is common in crystallography and is sometimes referred to as crystallographic resolution. Another common definition is the *half-period resolution* which is half of the full-period resolution. This resolution corresponds to the size of a pixel when Fourier space is cropped at the resolution limit and is therefore often more intuitive.

The step of averaging and calculating a PRTF is an essential part of any reconstruction since it provides a way to judge the quality and resolution of the reconstruction. It also provides the only standardized method of quantifying the resolution of a reconstruction, and was done for all the reconstructions in this thesis. For an example of a radially averaged PRTF, see figure 7.3.

4.4 Support recovery

Most phasing algorithms assume that the support is known, or in other words, that the shape of the sample is known. Although sometimes our prior knowledge of the particle is enough to estimate a decent support, this is not generally true.

A common way to estimate the shape of the support of an unknown object is through calculating the *autocorrelation* of the object, which is the objects cross-correlation with itself:

$$a(\vec{r}) = \int \rho(\vec{r} + \vec{t}) \rho(\vec{t})^* d\vec{t} \quad (4.13)$$

The autocorrelation has a size that is twice as large as the size of the sample along any direction. What makes it very useful to us is that it can be calculated from the diffracted intensity alone, which can be shown through the convolution theorem.

$$a = \mathfrak{F}^{-1} \{ \mathfrak{F} \{ \rho \} \mathfrak{F} \{ \rho \}^* \} = \mathfrak{F}^{-1} \{ A^2 \} = \mathfrak{F}^{-1} \{ I \} \quad (4.14)$$

We can thus calculate the autocorrelation by simply inverse Fourier-transforming the diffracted intensities and thereby get an envelope that is twice as large as the object.

4.4.1 The shrinkwrap algorithm

It is often important to have a tight support, i.e. one that closely resembles the shape of the particle without giving it room to move around. This is especially true for experimental data, and the autocorrelation method is then usually not good enough. The most common way of handling this problem is to recover the support during the reconstruction through an algorithm called *shrinkwrap*[32].

The name is borrowed from the type of plastic wrapping that shrinks when exposed to heat and can thus be made to fit very tightly around objects. In the same way, this algorithm can start with a very loose support and successively shrink it until it fits very tightly around the actual shape of the object.

The idea is that when we start the reconstruction we don't know the shape of the object very well but as the reconstruction progresses we gain more knowledge about the sample and can make a better estimate of the object. The shrinkwrap uses any phasing algorithm such as the HIO or RAAR and regularly (usually every 20 iterations) updates the support. The update is done in two steps:

1. A Gaussian blur is applied to the current real-space image.
2. A threshold is applied and all parts of the blurred image with densities higher than the threshold is included in the new support. The threshold is usually defined as a percentage of the maximum value.

The strength of the blurring, σ_{blur} , is usually decreased throughout the reconstruction. The reason is that we don't want to over-interpret the result from the first iterations while, as we gain more confidence also in the higher resolution of the reconstruction, we can use it to create a better support.

The shrinkwrap is usually good at excluding pixels from the support but not at including new pixels, meaning that with time it will shrink and not grow. This is the reason for its shrinkwrap like behavior and is also the reason for why the blurring is important

since badly recovered parts of the image could otherwise cause parts belonging to the object to be discarded from the support and probably never included again.

As stated earlier, the support was introduced as the required additional constraint to compensate for the information lost with the phases. It seems risky to give up control of the support constraint and it begs the question of how the problem is actually constrained when using the shrinkwrap algorithm. The answer is sadly not very clear since there is no analytical analysis of the algorithm but it is empirically well tested and has been shown to work in many cases. The lack of a solid understanding of the limits of the shrinkwrap is still a weakness of the convex-optimization phase-retrieval methods.

The thresholding in the shrinkwrap works very well for samples with sharp boundaries since these are not very sensitive to the exact value of the threshold. This makes the algorithm very suitable for example for reconstruction of objects created with a Focused Ion Beam (FIB) which was also the type of sample used when developing the algorithm. Most biological particles of interest such as cells and viruses also have sharp boundaries and the shrinkwrap can be expected to work well for them. This is however often not true for the projection images of these samples that are accessible from single patterns, and can increase the difficulty of these reconstructions.

4.4.2 Variation on the shrinkwrap

One variation on the shrinkwrap that was used in many of the reconstructions in this thesis is the *constant-area shrinkwrap*. The only difference compared to normal shrinkwrap is that the fixed thresholding is abandoned. Instead, the strongest part of the blurred image is included up to the level where the support has a specific area.

Since the original shrinkwrap very rarely grows the support, it makes one attempt at shrinking around the correct structure and, if it fails, continues to shrink all the way to a one-point support. The constant area prevents this shrinking and the algorithm can cover a large search space in a single run. It is therefore preferable if the problem is hard.

If the area of the object is not known, it is common to slowly let the area decrease during the reconstruction. When the area becomes smaller than the actual size of the object, the error metrics will rise dramatically which gives a good target area for the next run.

4.5 The Hawk software package

The Uppsala-developed Hawk package is the only open-source software for single-particle diffraction analysis and is described in **paper V**. It includes the methods described in this section and many more, and a rich library of utility functions combined

with being open source makes it possible for any user to extend the software. Another strength of Hawk is an efficient implementation in C and the ability to use Graphics Processing Units (GPU) when available for very fast computation.

Hawk and its underlying libraries were used for all the reconstructions in this thesis.

4.6 Missing data

Most diffraction patterns from single-particle experiments lack data. Even in the best case, a region in the center will have to be missing since the direct beam would otherwise damage the detector. This is achieved by, for example, placing a beam stop in front of the detector or by constructing the detector, or system of detectors with a hole in the center that lets the direct beam through. Regardless of implementation, this does however mean that data are lacking. In this section we will investigate what consequences this has on the reconstructed image. See [23] and [40] for a description of the detectors used at LCLS.

All algorithms mentioned here handle the missing data in the same way. When applying the Fourier-space constraint, the pixels with missing experimental amplitudes retain their recovered amplitude, just like all pixels retain their phase. Thus, the missing amplitudes are recovered together with the missing phases.

We note, that the missing parts of Fourier space are completely unconstrained by the Fourier-space constraint. Similarly, the real-space constraint only constrains the area outside the support while the inside is unconstrained. A pressing question is whether you can construct an object that is unconstrained by both the real and Fourier constraint. This would be an object that fits inside the support and has a Fourier transform that is zero outside of the missing-data region. Such an object could be added to any solution without violating either the real-space constraint or the Fourier-space constraint, and thus create an equally good solutions which creates an ambiguity in the reconstruction. Such objects are said to be unconstrained.

Totally unconstrained objects don't generally exist. What does exist are however weakly constrained objects. Objects that have only a very small contribution outside the support and outside the missing-data region. These objects will only influence the real and Fourier errors slightly and if the level of noise in the pattern is large enough, the algorithm might be insensitive to these small differences and the object will behave as if unconstrained.

The silver bullet to this problem is to minimize the size of the missing-data region. To do this, it is important to be able to predict how severe the problem is for a certain experimental setup. This is done in the next section. Sometimes it is not possible to completely eliminate the problem and then it's even more important to be able to quantify the effect well and possibly use additional constraints to solve it.

4.6.1 Mathematical description of the problem

The Fourier transform is a linear transform and the discrete version of it can be described by a matrix that we denote \mathfrak{F} . Every column in this matrix correspond to a certain pixel in real space and every row correspond to a pixel in Fourier space. We can rearrange the columns and rows to split the matrix into four submatrices based on whether they lie inside or outside the support and inside or outside the missing data region respectively. Equation 4.15 visualizes this split where S symbolizes the support and M symbolizes the missing region and the bar denotes their inverses.

$$\mathfrak{F} = \left(\begin{array}{c|c} \mathfrak{F}_{SM} & \mathfrak{F}_{\bar{S}M} \\ \hline \mathfrak{F}_{S\bar{M}} & \mathfrak{F}_{\bar{S}\bar{M}} \end{array} \right) \quad (4.15)$$

We define the problem as finding the object inside the support whose Fourier transform gives the smallest contribution to the known-data region, i.e. we are searching for a vector ρ that minimizes $|\mathfrak{F}_{S\bar{M}}\rho|$.

To solve this problem we are going to use a singular-value decomposition (SVD)[16]. The theory states that the matrix $\mathfrak{F}_{S\bar{M}}$ can be decomposed in the following way.

$$\mathfrak{F}_{S\bar{M}} = U\Sigma V^\dagger \quad (4.16)$$

where Σ is a rectangular diagonal matrix and U and V are both unitary matrices. The rows and columns are usually arranged such that the values of Σ , called the singular values, are sorted in descending order.

Any column vector, \mathbf{V}_i of the matrix V , also called a right singular vector, can be seen as a potential input to the matrix $\mathfrak{F}_{S\bar{M}}$ and the output will be the corresponding column vector \mathbf{U}_i of the matrix U , called a left singular vector, scaled by the corresponding singular value Σ_i . The smallest singular values correspond to the most weakly constrained objects, or *weakly constrained modes*.

Since V is unitary, the column vectors will be orthogonal to each other. Therefore, the set of right-singular vectors corresponding to the most weakly constrained modes will form an orthonormal basis set that spans a subspace where every vector is weakly constrained.

A way to think of this is that we have several degrees of freedom in our reconstruction that can not be found from the data alone. Figure 4.2 shows an example of such a weakly constrained mode in 1D.

4.6.2 Efficient calculation

The matrix $\mathfrak{F}_{S\bar{M}}$ can be very large. The support is typically not larger than 10000 pixels for 2D reconstructions, but the known-data region is usually very large, often




Figure 4.2. If the support and missing data region (indicated by the vertical lines) are both small enough it is possible to construct functions that are very weakly constrained by both the real-space and Fourier-space constraints. The Gaussian in this picture is such an example.

above 1 000 000 pixels, and this makes the problem computationally very heavy. We will now show that we can perform an equivalent calculation on the much smaller matrix \mathfrak{F}_{SM} .

The singular values of \mathfrak{F}_{SM} are also the square-root of the eigenvalues of $\mathfrak{F}_{SM}\mathfrak{F}_{SM}^\dagger$ [16] and the right singular vectors are the corresponding eigenvectors. We can use the fact that \mathfrak{F} is a unitary matrix i.e. $\mathfrak{F}\mathfrak{F}^\dagger = E$ where E is the unit matrix. It is then also true that

$$\mathfrak{F}_{SM}\mathfrak{F}_{SM}^\dagger + \mathfrak{F}_{SM}\mathfrak{F}_{SM}^\dagger = E \quad (4.17)$$

The eigenvectors of $\mathfrak{F}_{SM}\mathfrak{F}_{SM}^\dagger$ can then be written in the following way:

$$\mathfrak{F}_{SM}\mathfrak{F}_{SM}^\dagger \mathbf{V}_i = \Sigma_i^2 \mathbf{V}_i \quad (4.18)$$

$$\left(E - \mathfrak{F}_{SM}\mathfrak{F}_{SM}^\dagger \right) \mathbf{V}_i = \Sigma_i^2 \mathbf{V}_i \quad (4.19)$$

$$\mathfrak{F}_{SM}\mathfrak{F}_{SM}^\dagger \mathbf{V}_i = (1 - \Sigma_i^2) \mathbf{V}_i = \Sigma_i'^2 \mathbf{V}_i \quad (4.20)$$

This result means that we can calculate both the singular vectors and the singular values by working on the much smaller matrix \mathfrak{F}_{SM} . The weakly constrained modes will then instead be characterized by a strong singular value. More precisely the sought after singular value can be calculated from the singular value Σ_i' calculated with the efficient method as

$$\Sigma_i = \sqrt{1 - \Sigma_i'^2} \quad (4.21)$$

4.6.3 Relation to noise

In the above calculation, we have treated every pixel of the diffraction pattern equally, meaning that a contribution from a mode to one pixel is equally constraining as a contribution to any other pixel. If we know the noise for each pixel we can obviously




Figure 4.3. The figure shows the experimental setup used for all experiments described in **paper VI**. The sample is suspended on a thin membrane and the scattered light is reflected down to the CCD detector. The role of the mirror is to split the diffracted light from the direct beam that is let through a hole in the mirror.

do better since we know that a pixel with strong noise will constrain the mode less than a pixel with weak noise.

Surprisingly, this assumption of equal noise is a fairly good approximation when the noise is dominated by Poisson noise. The standard deviation of the noise is then given by $\sigma_{I,i} = \sqrt{I_i + 1}$ where I_i is the intensity in the i th pixel given in number of photons. The amplitudes used in the reconstruction relate to the intensity as $A_i = \sqrt{I_i}$. Error propagation of the standard deviation $\sigma_{I,i}$ to the standard deviation of the noise in the amplitudes gives

$$\sigma_{A,i} = \frac{1}{2} \sqrt{1 + \frac{1}{A_i^2}} \quad (4.22)$$

This shows that unless the intensity is very small, the noise is constant at $\frac{1}{2} \sqrt{\text{photon}}$.

This result is relieving since a more proper handling of the effects of noise requires a reweighing of the individual rows of \mathfrak{F}_{SM} which unfortunately means that the efficient calculation described in 4.6.2 doesn't work since that breaks the unitarity of \mathfrak{F} .

4.7 Reconstructing biological samples from experimental data: Imaging of cells on solid supports

In **paper VI** we used the techniques described in this section to perform some of the first single particle X-ray imaging experiments of biological samples, at the free-electron laser FLASH in Hamburg. The samples were suspended on a 20 nm thick silicon-nitride membrane that was hit by a pulse from the FEL. The scattered light was reflected down to the CCD detector by a graded multi-layer mirror. A central hole in the mirror let the direct beam through, thus preventing it from damaging the detector, see figure 4.3.

Figure 4.4 shows a reconstruction of a *Prochlorococcus marinus* cell from the paper. What looks like a hole in the center of the image is most likely a ruptured membrane that caused some of the cell content to escape, leaving the rest of the cell sunken in. This kind of damage was also seen in electron-microscopy (EM) images of the same cell type and are probably caused by a long exposure to vacuum, the cells spent several hours in the chamber before being imaged.

In addition to the *Prochlorococcus marinus* cell showed in figure 4.4, **paper VI** also shows reconstructions of *Synechococcus elongatus* and *Spiroplasma melliferum* cells. Even though no biologically relevant claims about the samples were produced in the paper, the result was significant in showing that cells can be imaged by the diffract and destroy technique.

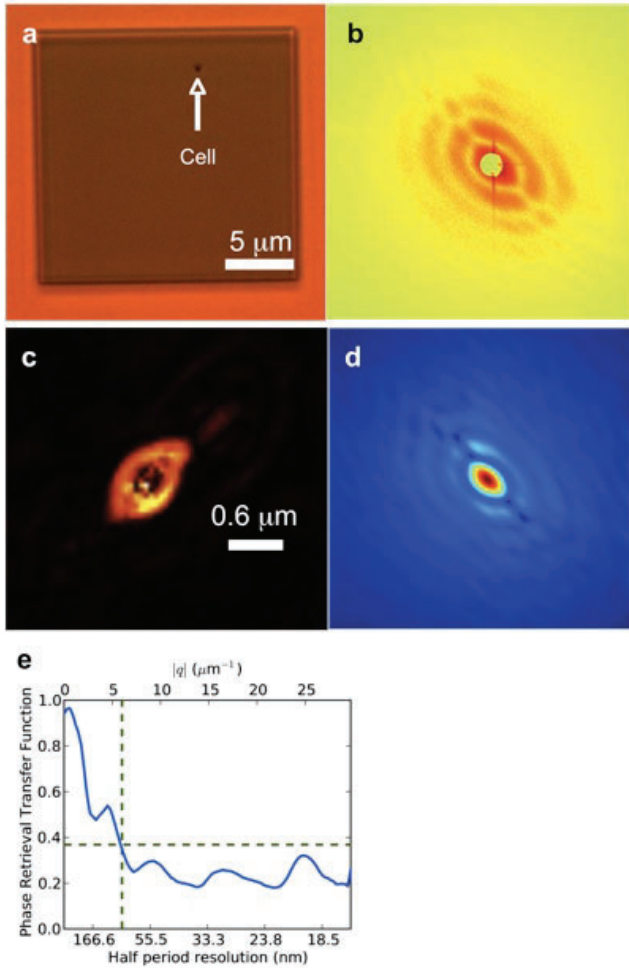


Figure 4.4. The figure shows a picture of the membrane with the cell taken with an optical microscope (a) and the measured diffraction pattern (b). The recovered pattern (d) matches closely the experimental one. The reconstructed image (c) has the expected dimensions but shows sign of damages in its central region, probably from an being exposed to vacuum for several hours. The PRTF (e) gives us a half-period resolution of roughly 80 nm.

5. The giant mimivirus

5.1 Background

Mimivirus (*Acanthamoeba polyphaga* mimivirus) is one of the largest viruses known today[26]. With a diameter of 450 nm the size is comparable to that of the smallest living cells and its name is actually a short for microbe mimicking virus. The virus has a pseudo-icosahedral capsid with a possible five-fold symmetry[49], but there is no known symmetry of the internals of the virus.

The double-stranded DNA genome has 1.2 million base pairs[41] which is one of the largest genomes found in any virus. It was the first virus found to have a larger genetic complexity than some cellular organisms[11] and even contains central parts of the protein translation apparatus. Studies of the mimivirus have sparked new debates about the boundaries between viral and cellular life.

The capsid is covered by a layer of thin fibrils[49, 10] and the total size of the particle including these fibrils is about 750 nm. The size and the fibrils make the virus impossible to study with crystallographic methods. The size also makes it hard to study intact virus particles with EM because of the limited penetration depth of electrons. This makes it a suitable target for demonstrating single-particle X-ray imaging.

5.2 Two-dimensional imaging of injected mimivirus particles

In **paper II** projection images of mimivirus particles were reconstructed from single diffraction patterns. In these studies the previous method of placing the sample on a membrane was replaced by a method of injecting a stream of particles into the FEL pulse train without any container. This has several benefits. First, it minimizes the time the sample is exposed to the vacuum, thus reducing the risk of drying out. Secondly, the diffracted signal is now free from scattering from a membrane that would otherwise reduce the quality of the data. Last, only a small number of samples can be suspended on one membrane and only a small number of membranes can be kept in the vacuum chamber at one time. Since opening the chamber is a lengthy process, this is rather avoided. With the injection technique the chamber doesn't need to be opened to change sample and in addition, there is no sample-carrying membrane that has to be realigned for each new shot. The injection technique has therefore increased the data rate by several orders of magnitude.




Figure 5.1. The mimivirus particles were injected through an aerodynamic lens (left) and intercepted by the X-ray pulses. The diffracted signal was captured on a pnCCD detector with a hole to let the direct beam through.

Purified mimivirus particles were transferred into a volatile buffer and were then aerosolized with helium in a gas dynamic nebuliser[15]. The aerosolized mimivirus particles were then injected into the pulse train of the FEL using an aerodynamic lens[4]. The electron bunches were measured to be 70 fs long at FWHM which correspond to a pulse length of between 20 fs and 40 fs[50]. Simulations of the damage process in **paper VI** and for nanocrystallography in [1] assures that we will see no effects of radiation damage at resolutions above 1 nm.

The X-ray energy was 1.8 keV and the pulse was focused to a spot with a diameter of $10\ \mu\text{m}$ (FWHM) with a peak intensity of $1.6 \cdot 10^{10}$ photons per square micrometer. The pnCCD detector[23] was placed 564 mm away from the interaction region which gives a theoretically achievable full-period resolution of 10.2 nm at the edge of the detector.

The experiment was performed at LCLS in 2010 and although many images were collected only two were selected for analysis. The selected images had fairly strong scattered signal but without a too large missing-data region. The missing data was, in addition to a central hole in the detector, caused by too high signal in the central part of the diffraction pattern which saturated the detector. In addition, the detector was built to handle saturation by letting the extra charge spill over to nearby pixels, especially in the vertical direction. This can cause a small highly illuminated area to render a large part of the image unusable.

The diffraction patterns and the corresponding autocorrelations are shown in figure 5.2 (a) and (b) and (d) and (e) respectively. Both autocorrelations show pseudo-hexagonal shapes which is characteristic for a projection of an icosahedron. Even though the images were selected for having a small region of missing data, we can see from the images in figures 5.2 (a) and (b) that such regions do exist.

The phases were recovered using the RAAR algorithm described in section 4.3.3 and the support was handled by a constant area shrinkwrap. The PRTF of the reconstruction is shown in figures 5.2 (h) and (i).

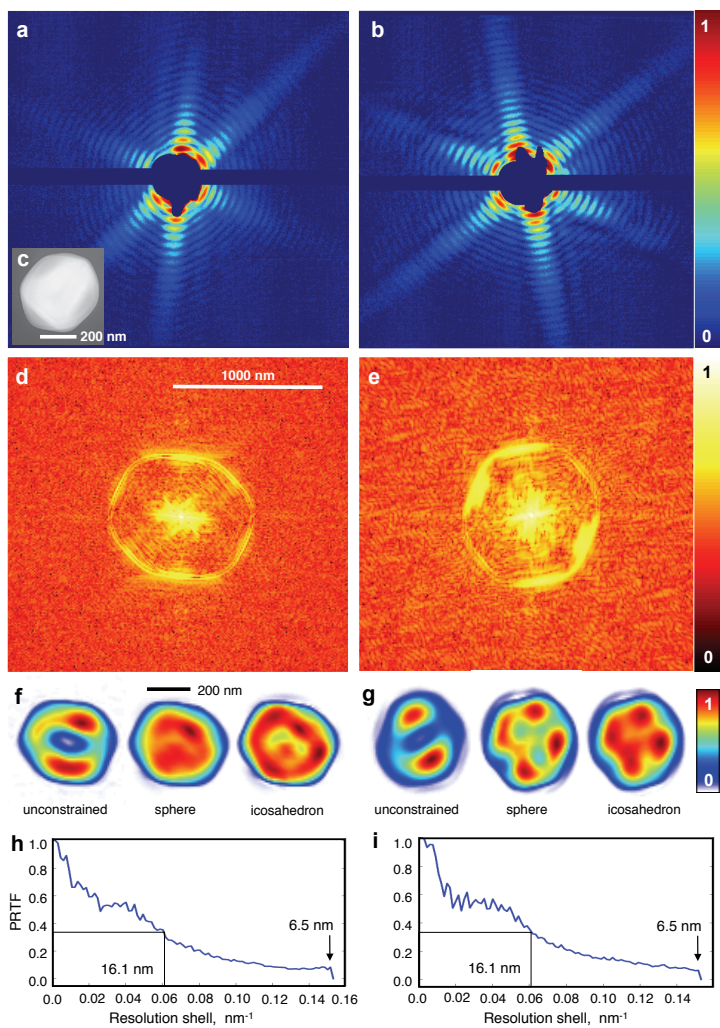


Figure 5.2. The figure shows the collected diffraction patterns from the mimivirus particle (**a**, **b**), an EM picture of the virus (**c**), the autocorrelations showing a pseudo-hexagonal shape (**d**, **e**), the reconstructed images with no mode correction, and correction based on assuming a spherical and icosahedral density respectively (**f**, **g**) and the PRTF of the reconstructions showing a full-period resolution of 32 nm in both reconstructions. (**h**, **i**)

Several weakly constrained modes exist in both images and were identified by the methods outlined in section 4.6. This essentially means that the low-resolution information in the reconstructions is unreliable, so to recover these modes we had to provide more assumptions about the sample. We chose to publish reconstructions based on two different assumptions. In the first one the strength of the unconstrained modes were fitted to make the total density of the virus fit a sphere. In the second one, the modes were fitted to make the total density fit an icosahedron. The icosahedron here had its orientation determined so that scattering simulated from it fit the experimentally measured patterns. Both results are shown together with the recovered particle without mode correction in figures 5.2 (f) and (g).

The reason for publishing both fits was to give a fair impression of the reliability of the reconstruction. The features that are common to both the spherical and the icosahedral fit are fairly trustworthy while the differing parts are likely to be biased by the assumptions. The reconstructions showed an inhomogeneous inside which is very uncommon for viruses.

6. Aligning diffraction patterns

In section 3.5 we learned that a diffraction image samples the Fourier transform of the scattering potential of the sample at the Ewald sphere. If the same sample is imaged multiple times at different orientations, the diffraction images will sample different spheres cutting through the same Fourier space, and by combining many such images the entire Fourier space can be stitched together. Figure 6.1 shows two such Ewald spheres.

In single particle CXI the sample is destroyed by the pulse so getting more than one diffraction image per sample is impossible¹. Some samples are however reproducible so that multiple images taken from different samples can be treated as if they were from the same sample. Especially proteins and viruses are often reproducible even at high resolution.

This does however introduce a new problem. Since every particle is injected into the experiment at an unknown orientation, we don't know the orientation of the sampled Ewald-spheres. Several methods have been proposed to solve this problem, most of them rely solely on the information that the scattered intensities carry about the orientation.

6.1 Common lines

Imagine two diffraction patterns of the same object, they will each sample an Ewald sphere from the same Fourier space. The two Ewald spheres will intersect and this intersection will have the shape of a curved line as shown in figure 6.1. This implies that to find the relative orientation between the two Ewald spheres, all we need to do is to find these lines. This is the basis for an algorithm called *common lines* or sometimes *common arcs*[6].

This method is also used for aligning samples in cryo EM[20]. There the very short wavelength of the electrons makes the Ewald-sphere sections essentially flat. This also happens in CXI when the wavelength is much shorter than the target resolution. In these cases, the identification of one line is not enough to obtain the relative orientation, and at least three patterns have to be combined for their relative orientations to be determined[46].

¹It has been suggested that by splitting the beam and illuminating the sample from multiple sides at once, more than one image can be collected. Even with this method, we can't however get enough data to fill the entire Fourier space.

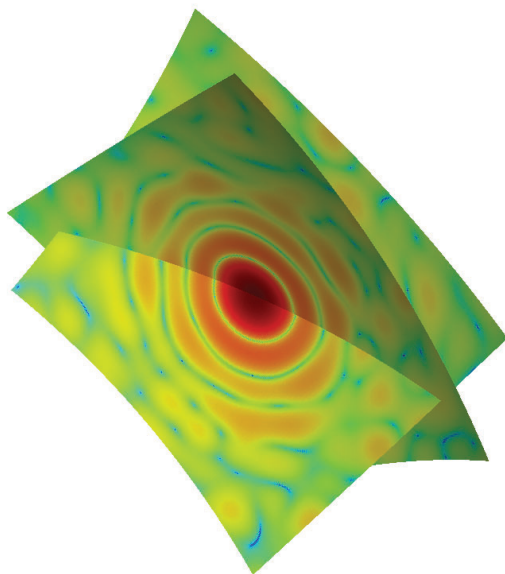


Figure 6.1. Two diffraction images of the same sample will sample parts of two different Ewald spheres. The two spheres will intersect along an arc.

The implementation is fairly straight forward. For every possible relative orientation, the supposedly intersecting arcs are identified and compared, the relative orientation that gives the best match is taken as the true one. Or, in the case of flat Ewald-spheres, this gives two of the three angles defining the orientation.

There is no consensus on how to compare lines. In [6], a Pearson correlation factor was used after the patterns were normalized by the total radial average of the entire dataset. The reason for the normalization is to prevent the very intense central region from dominating. This method has shown positive results on simulated data but other measures such as the Euclidean distance are also used.

To go from the pairwise relative orientations to absolute orientations, the simplest way is to fix the orientation of one diffraction pattern and then use the pairwise orientations relating this pattern to every other pattern to determine the absolute orientation of the respective pattern. This method is however very sensitive to potential misalignments. It is therefore more common to include two- and even three-step routes to a specific pattern. This gives a more accurate orientation and allows for identifying non-matching pairwise-orientations that can then be assumed to be false[6].

This method is extraordinarily simple and computationally both fast and parallelizable but does suffer from some serious drawbacks. Most notably is the sensitivity to noise. Since only a few pixels on a line determines every orientation, the method doesn't use the data very efficiently. The algorithms described below are all methods that use

all the data at once instead of taking the route through pairwise orientations, which allows for much better robustness to noise.

6.2 Manifold embedding

Let us introduce a new way to think of a dataset of diffraction patterns[22]. Start by defining each pattern as a vector \mathbf{a}_i where each element a_{il} is the value of the l th pixel. A 1000 by 1000 pixel diffraction pattern will then be described by a single vector with 1000000 elements.

All patterns in a dataset will be vectors in a very high-dimensional vector-space. Still we know that except for noise, the only thing different between one pattern and another is the orientation. And since any orientation can be described by three numbers (for example three Euler angles[18]) there can be only three degrees of freedom in the data. This means that there is a three-dimensional manifold on which all these vectors reside.

The above statement is only true in the noise-free case, but this is still a very important result. It means that if we can find the manifold that best fits the data, this will represent a denoised version of the data and if we can find a way to map the corresponding orientations to points on this manifold, it will also give us the orientations.

We have now reformulated the orientation problem into the more general problem of fitting a low-dimensional manifold to high-dimensional data, which is a well known problem called dimensionality reduction. Many algorithms exist that solve it, for example the Self Organizing Map (SOM)[25], Generative Topographic Mapping (GTM)[3] and diffusion maps[12]. Both of the two latter methods have been applied to simulated diffraction data and for the GTM, seemingly positive results have been reported[22].

It is hard to draw general conclusions on the whole family of methods, we will therefore concentrate on the only method that has yet showed positive results, the GTM. While the method rests on a firm theoretical base of Expectation Maximization², it suffers from problems with convergence to local minima. To solve these problems the authors of [22] have introduced an extension of the method where the manifold, during the course of the iterative algorithm, is split into patches that are then joined in a more favorable way[36]. This deviation from the original Expectation Maximization algorithm seems to improve the result but we loose the benefit of using a theoretically well understood algorithm.

Symmetric particles are common in biology and they do present a problem to the GTM and also generally to manifold embedding methods. It can probably be handled if the symmetry group is known and the algorithm is adapted accordingly, but it has yet to be demonstrated. If it is unknown whether a symmetry exist or not there are however no known way of handling it, except repeatedly testing the algorithm with

²Often abbreviated EM, which is avoided here to avoid confusion with Electron Microscopy.




Figure 6.2. The orientation of every pattern is determined under the constraint of being self consistent with all other patterns in the dataset.

different symmetries imposed, and still no method exists to evaluate the result of such a comparison.

A strength of the method is that it uses all of the data together and is therefore much more robust to noise than the common-lines method. It has been demonstrated to work for very low Signal to Noise Ratios (SNR)[22] but there are no tests that show its performance under other types of noise than Poisson noise.

Another strength is that it can easily be extended to handle sample diversity. For example, a data set where the sample differs along one or several degrees of freedom could likely be both oriented and sorted with the method[44]. This still lacks a demonstration though. This topic will be discussed further in section 8.

6.3 Expansion Maximization Compression (EMC)

The *Expansion Maximization Compression* (EMC) algorithm is another attempt at solving the orientation problem[28]. As opposed to the manifold embedding techniques it makes direct use of the knowledge that the diffraction patterns must fit together in a self-consistent three-dimensional Fourier-space, see figure 6.2.

The algorithm was used in **paper I** and we will therefore explain it in more detail. The algorithm is iterative and each iteration consists of three steps: Expansion (E), Maximization (M) and Compression (C).

The algorithm is seeded with an estimate of the Fourier-space that will be refined by each iteration. The start can be based on any prior knowledge we have about the particle, but random starts such as random intensities or patterns entered at random orientations also work.

Expansion

In the first step of each iteration the current Fourier-space is expanded into all possible Ewald-sphere slices. In other words, we calculate the diffraction pattern for every particle orientation, given the current 3D Fourier model.

Maximization

Next, the experimental images are compared to the expanded model calculated in the previous step. This results in a matrix, R_{jk} , describing how well each image fits in every orientation, where j is the index of the model slice and k is the index of the experimental image. The metric used is the probability of observing the experimental diffraction pattern given some assumed type of noise. Poisson noise gives the following metric

$$d_{\text{Poisson}}(K, W) = \prod_i \frac{e^{-W_i} W_i^{K_i}}{K_i!} \quad (6.1)$$

where K is the experimental pattern, W is the Fourier slice and i is the respective pixel index. If we believe that the dominant form of noise follows a Gaussian distribution, which could be for example detector-noise, the appropriate distance metric would be

$$d_{\text{Gaussian}}(K, W) = \prod_i e^{-\frac{(W_i - K_i)^2}{2\sigma^2}} \quad (6.2)$$

where σ is the standard deviation of the noise. The distance metric should be adapted to whatever noise we expect in the data, hence it is important to have a good understanding of the noise.

Compression

In the last step of the iteration, a new model is created based on the experimental patterns. The patterns are assigned orientations based on where they fit the best. In the original version of the algorithm[28] the patterns are distributed over all orientations with a weight proportional to how well they fit in that particular orientation. Other implementations of the algorithm use only the orientation with the best fit[47].

Scaling

Normally, the strength of the X-ray field interacting with the sample is not known. Even though it is possible to measure the total strength of the pulse, we don't know where in the pulse the sample was located. In [27], a method to recover these intensities alongside with the orientations was introduced. It estimates the scaling ϕ after the maximization step based on the formula:

$$\phi_k = \frac{\sum_j R_{jk} \sum_i K_{ik}^2}{\sum_j R_{jk} \sum_i K_{ik} W_{ij}} \quad (6.3)$$

where K_{ik} is the value of pixel i of the k th experimental pattern and W_{ij} is the value of pixel i of the j th model slice.

Identifying outliers

Many problems require a large number of diffraction patterns. This is the case for example for high resolution reconstructions or when the noise in a single pattern is very high. It is then not possible to hand pick the patterns to be included in the reconstruction. Automated filtering works fairly well but is far from perfect at the moment and even with good algorithms there are some types of "bad" patterns that are virtually indistinguishable from good patterns. One such case is when one pulse hits two particles that are both aligned with the beam axis. The pattern is then virtually indistinguishable from that from a single particle.

For these reasons it is important for the algorithm to be able to identify patterns that don't fit in the recovered Fourier-space. We have implemented this feature as a special bin alongside the rotational bins, that doesn't contribute to the Fourier volume. This bin also has the feature of forcing patterns to be either fully in or out, as opposed to the fuzzy classification used in the rotational bins. An important feature of this type of filtering is that patterns can have their classification changed throughout the progress of the algorithm and we are therefore not hurt by early misclassifications.

To identify bad patterns, for every pattern the summed likelihood of the pattern to be in any orientation is calculated:

$$\ell_k = \sum_j R_{jk} \quad (6.4)$$

Now, a fixed threshold of ℓ_k can be used to identify patterns that don't fit, or a fixed ratio of the images with lowest ℓ_k can be selected. For future updates of the algorithm, a better approach would be to identify a gap in ℓ_k between good patterns and bad patterns, and thus identify a suitable threshold on the fly.

Comparison to manifold embedding

Both the EMC and manifold embedding algorithms make use of the data in an efficient way. The key difference is that EMC forces the patterns to be consistent as a 3D Fourier-space whereas the manifold embedding techniques only require the dimensionality to be three. This makes the EMC algorithm very strong, but comes at the expense of having no theory proving convergence, which exists for many of the manifold embedding techniques, for example the GTM. In practice the EMC algorithm is however the only algorithm that has been successful on experimental data so far[27].

We mentioned that GTM had problems with symmetric particles. This is not the case for EMC due to the ability to distribute a single diffraction pattern over multiple orientations. This allows for the EMC algorithm to identify the symmetry by itself. This is an important benefit when studying objects with unknown symmetry, like many protein complexes and viruses.

Possible issues

While the probability-based distance metrics make sense when the algorithm is close to the solution, they can hurt the algorithm before it has converged. Before convergence, the main difference between the diffraction patterns and the model slices is due to the inaccuracy of the model and not the noise. To the noise model, all orientations will seem super improbable which can cause unexpected behaviors such as misalignments and numerical errors.

A more proper way to handle this would be to introduce a model-confidence term into the distance metric. This is a neat solution but requires careful tuning to assure that the distances are reasonable for all stages of the algorithm.

A quick-and-dirty remedy could be to use the Gaussian metric with a large σ to handle the large model errors at early iterations and then decrease the value as the algorithm progresses. The decrease has to be very slow to make sure that the model has time to adapt to the new σ at every stage of the algorithm.

6.4 Estimating the required number of diffraction patterns

Independent of the ability to find the correct orientation of each pattern, we need enough diffraction patterns to ensure a sufficiently dense sampling of Fourier space up to the target resolution. Since the process of orienting diffraction patterns is computationally very hard it is usually done post beam-time, so currently there is no way of knowing if the collected data is covering enough of Fourier space while still collecting data. What we can do is to calculate the number of diffraction patterns statistically

required to get a dense enough sampling, with a certain probability. This problem was covered in **paper III**.

We define a full coverage here as where every pixel in the pattern is sampled at least once. The required sampling rate, that defines the size of these pixels, is denoted s , where $s = 1$ is the Nyquist rate. A sampling of $s = 2^{1/3}$ is a common requirement for phase recovery.

We note that if every pixel in the outermost shell of pixels are covered then every pixel inside will be covered as well. We will therefore only analyze the occupancy of this outer shell and we define $R = D/d$ where D is the object size and d is the resolution. The number of pixels in the outer shell is then given by $K_s = 4\pi \left(sR - \frac{1}{2}\right)^2$ and the number of pixels in the outer shell covered by a single diffraction pattern is $K_s = 2\pi \left(sR - \frac{1}{2}\right)$. This can be further reduced when there is a flat Ewald geometry since every covered pixel has a centrosymmetric partner that will be covered by the same patterns and can therefore be treated as the same:

$$K = \frac{K_s}{2} = 2\pi \left(sR - \frac{1}{2}\right)^2 \quad (6.5)$$

$$k = \frac{k_s}{2} = \pi \left(sR - \frac{1}{2}\right) \quad (6.6)$$

Making the simplifying assumption that the coverage of each pixel is independent of the covering of other pixels we can use K and k to calculate the probability of achieving a full coverage with N patterns:

$$p = \left(1 - \left(1 - \frac{k}{K}\right)^N\right)^K \quad (6.7)$$

If we instead set a target probability p we can estimate the number of patterns required to reach our target resolution:

$$N = \frac{\ln(1 - p^{1/K})}{\ln(1 - k/K)} \quad (6.8)$$

An example is plotted in figure 6.3 for $p = 0.95$ and $s = 2^{1/3}$.

If the scattered signal is very weak the number of patterns will instead be limited by overcoming the noise. In that case the number of patterns will scale with increasing resolution as.

$$N \propto (2sR - 1)(2sR)^{d_f} \approx (2sR)^{d_f+1} \propto R^{d_f+1} \quad (6.9)$$

where d_f is the fractal dimension of the particle. For proteins, this number is usually in the range of 1.5 to 2.0[14] and the number of required patterns therefore increases with at most resolution to the power three.

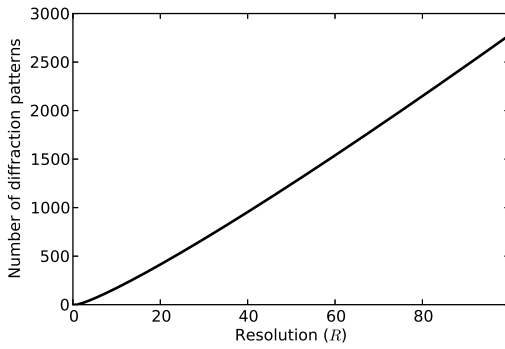


Figure 6.3. The number of images required to completely fill Fourier space with a probability of 0.95 is plotted against R for $s = 2^{1/3}$. As the resolution gets higher, the required number of images increases. For high resolutions the increase is close to linear.

7. Experimental three-dimensional imaging of the mimivirus

In **paper I** we perform a three-dimensional reconstruction of the mimivirus. As a proof of concept, this result is very important since it is the first particle to be imaged in 3D with this technique. The experiment was performed at the AMO endstation at LCLS. Just like for the two-dimensional imaging in **paper II** (see section 5.2) the samples were aerosolized in a volatile buffer and injected into the vacuum through an aerodynamic lens.

The pnCCD detectors[23] were placed 740 mm from the interaction region which was the largest distance permitted by the geometry of the vacuum chamber. The reason for the large distance was to reduce the effect of the missing data due to the central hole in the detector. Further, the X-ray energy was chosen to 1.2 keV which, together with the distance, assures no weakly constrained modes for a sample of 450 nm size (see section 4.6). Each pulse contained about $1.2 \cdot 10^{12}$ photons which corresponds to a total pulse energy of 0.24 mJ. The pulse was focused to a spot with a diameter of $10 \mu\text{m}$ (FWHM).

Several tens of thousands of images were captured and many hundreds of them contained scattered signal, implying that a particle was intercepted by the pulse. To identify diffraction patterns that originate from mimiviruses, a program sorted the patterns based on the existence of the characteristic six-streak star shape common to diffraction from an icosahedron. 307 diffraction patterns were identified in this way.

A second level of sorting was performed by hand to get rid of images containing saturation. Not having saturation assures that all images can be compared to the model based on the same pixels, and therefore makes the comparisons more reliable. In this step 261 patterns were identified as suitable for using in a 3D reconstruction.

The pnCCD detectors have a dynamic range of only about 500 photons at 1.2 keV. This means that the selection of unsaturated diffraction patterns will severely limit the resolution by avoiding the stronger pulses. Since the experiment was mainly intended as a proof-of-principle experiment we did however prioritize the robustness of the algorithm above resolution.

The detector output data of 1024×1024 pixels was cropped to 256×256 pixels which corresponds to a maximum achievable resolution of 79 nm which, for the 450 nm sized mimivirus corresponds to $10 \times 10 \times 10$ independent pixels in the sample. We also binned the patterns 2×2 resulting in final patterns of 128×128 pixels. This makes the computation significantly less heavy. Based on the results in section 6.4, 261 patterns are sufficient to achieve a resolution of 79 nm.




Figure 7.1. The 261 slices oriented with the EMC algorithm. This image plots the patterns in logarithmic scale in the best-fitting orientation. A section of the Fourier volume has been cut away to reveal the central parts.

A slightly varying beam position meant that the center of the diffraction pattern fluctuated slightly from shot to shot. Before combining the patterns, the correct center therefore had to be identified for each pattern. The center position that maximized the Friedel symmetry of the image was chosen. The variation of the center position was found to have a standard deviation of 2.5 pixels in both the horizontal and vertical direction. The maximum difference was 9 pixels in the horizontal and 10 pixels in the vertical direction.

The patterns were oriented using the EMC algorithm. This algorithm was chosen mainly because of its proven effectiveness with experimental data and its ability to handle unknown symmetries. This was particularly important since the symmetry of the entire mimivirus particle was still unknown. The noise was modeled as a Gaussian with a fixed width since detector noise and miss-alignment error were expected to dominate over poisson noise. The oriented images in their most likely orientation i.e. the one with largest R_{jk} , are plotted in figure 7.1.

The resulting Fourier-volume was phased by Hawk using the RAAR algorithm with area-constrained shrinkwrap. 50 reconstructions with independent random starts were averaged to produce the final density map which is plotted in figure 7.2 as five iso-surfaces. High density is represented by dark blue and low density by light blue.

Figure 7.3 shows the PRTF from the 50 reconstructions. It shows reproducible phases almost to the edge of the data and the criterium of a PRTF higher than e^{-1} gives a full-period resolution of 83 nm. This resolution gives approximately 10 independent resolution elements along the particle diameter. Therefore only features of lower resolution than that are reliable. Apparent higher-resolution features in 7.2 are an effect of the iso-surfaces.




Figure 7.2. The reconstructed density of the mimivirus. Dark-blue iso-surfaces indicate high density and light blue iso-surfaces indicate low density. We see a strong shift of density to one side and a mushroom shaped low-density region inside the virus. At the current full-period resolution of 83 nm, details smaller than a tenth of the particle size are not reliable.




Figure 7.3. The PRTF shows that phases are recovered almost to the edge of the data. The resolution where the PRTF drops below e^{-1} is 83 nm. The periodic structure coincides with the fringes in the diffraction patterns and are caused by the fact that phases are generally harder to recover where the signal is weak.

The map shows a density distribution that is clearly not icosahedrally symmetric but does have a pseudo five-fold symmetry. It also shows two very interesting low-resolution features: a strong density that is shifted to one side and a mushroom-shaped lower-density region in the center. The high-density region could be attributed to the DNA in which case some mechanism to confine the DNA to a volume smaller than the capsid volume has to exist, which has so far not been known to exist in viruses. If the dense region is not the DNA it must be some other unknown dense structure.

Future studies of the mimivirus at FELs have great prospects of improving on the current resolution and further increasing our understanding of this internal structure.

8. Recovering conformations

8.1 Effects of structural variability on phase recovery

In **paper IV** we investigated the effects that sample variation has on the reconstruction process. In the orientation recovery process the samples were assumed to be identical. What are the effects if this assumption is not true?

The problems this cause turns out to be more severe than for example in crystallography. The reason is that we are averaging diffracted intensities while in crystallography the averaging is an effect of the scattered waves from many unit cells coinciding, hence the waves themselves, and not the intensity, is averaged. The averaging in crystallography is called a coherent averaging while our phase-free case is referred to as incoherent averaging. The difference is illustrated by the following equations.

$$U_{\text{coherent}} = \frac{1}{2} |U_1 + U_2| \quad (8.1)$$

$$U_{\text{incoherent}} = \frac{1}{2} (|U_1| + |U_2|) \quad (8.2)$$

where U_1 and U_2 are two diffracted waves and U_{coherent} and $U_{\text{incoherent}}$ is the coherent and incoherent average respectively.

In the paper we calculate the Fourier transform of a protein along many different states of a Molecular Dynamics (MD) trajectory. Under coherent averaging it contained variations up to 2.7 Å while under incoherent averaging it was only reconstructable up to 3.7 Å. This discrepancy arises since a coherent average of the scattered waves is equivalent to the scattering from an averaged sample while the incoherently averaged waves do not correspond to anything intuitively understandable. And worse, it doesn't correspond to anything interpretable to the reconstruction algorithm. This is the reason why the errors are larger than the structural differences.

Many proteins have several different conformations, and in addition to that there might be flexible parts adding to the diversity of the structures. This diversity is often central to the function of the protein. In a crystal, the packing can force the proteins to obtain fairly similar conformations and the heterogeneity of the sample is reduced.

Initially there was much skepticism about whether proteins retain their native structure when injected into vacuum. Extensive MD simulations have since shown that the structures stay preserved as long as a few water molecules are still present to occupy key hydrogen bonding sites. Without these water molecules, the protein stays folded but diverges more significantly from the native structure[39][34].

The lack of crystal constraints can increase the heterogeneity and could, in principle, reduce the resolution in an imaging experiment. But, it also means that the diffraction patterns represent the protein in many of its possible structures, not just the one most favored by crystal packing.

If we were able to separate the patterns into different conformational groups before averaging them we could not only circumvent the problem of heterogeneity but also image the full conformational space of the protein and this could provide information about the dynamics and function of the protein. In the next section we will discuss how current algorithms for orienting patterns could be extended to also do this type of sorting.

8.2 Extensions of alignment algorithms

Manifold embedding

If a protein exist in distinct conformations, the most straightforward extension of the algorithm is to simultaneously fit several manifolds to the data, one per conformation. The algorithm will likely struggle if the conformations are fairly similar since the manifolds will then lie close together and could fairly well be described by only one of the manifolds. A possible remedy would be to solve the problem assuming a single conformation and then duplicate the manifold when it is close to the solutions.

If the heterogeneity of the particle being imaged has one or several full degrees of freedom, the extension of the manifold embedding algorithms is conceptually fairly straightforward. It is done by extending the dimensionality of the manifold to include both the three rotational degrees of freedom and the additional conformational ones.

Extending the algorithm like this does however make the problem computationally harder and requires more data which will limit how large amounts of heterogeneity that could be handled. The technique was suggested in [44] but has not yet been backed up by demonstrations.

EMC

The extension of EMC to handle distinct conformations is as straight forward as for manifold embedding. We simply create one Fourier model per conformation and thus have one set of rotational bins per conformation. We will then not only determine the best orientation for each pattern but in the same way also the best conformation.

The extension to finding full degrees of freedom is not as simple as for manifold embedding. The seemingly equivalent way of adding a full set of Fourier models along the degree of freedom is a good idea but without some additional constraints, the ordering along the coordinate will lack meaning and this case will be similar to having

multiple unrelated conformations. One simple solution is to artificially add some overlap between the conformations. We know that models that are close together along the degree of freedom must also have similar Fourier models. The compression step should therefore be extended to also include a mixing between neighboring models. Further, this mixing should be large for low q and small, or even non-existent for high q . The reason is that at low resolution, the model changes little even over larger conformational changes while at high resolution the effect of even a small conformational change could be considerable.

9. Outlook

We are just beginning to utilize the potential held by free-electron lasers. Method development is rapid and accompanied by an equally fast development of the hardware. Technologies like seeding and improved detectors could push already existing sources to a point where almost any protein could be studied. In addition, new sources are being built and planned, for example the European XFEL is scheduled to start operation in 2015. It will have the capability to deliver 27000 pulses per second which could improve the rate of data collection by orders of magnitude.

I expect a rapid stream of new reconstructions and new biological results over the next few years. In particular I look forward to seeing development in the following three areas.

Two-dimensional imaging of live samples

The two dimensional images presented in this thesis were not simply a step on the way towards three-dimensional imaging but is a unique way to image irreproducible samples, like living cells. The method of injecting samples is very gentle and both cells and viruses have been shown to be active after injection. This, combined with the very short time they spend in vacuum allow us to image cells at high resolution while they are still alive.

The resolution of the 2D images have so far been limited by the dynamic range of the detector. Increased dynamic range, or a detector gain that scales radially could increase the resolution by up to an order of magnitude. Further, better focusing optics and seeded lasers could allow us to image irreproducible objects at truly unprecedented resolutions.

Three-dimensional imaging of reproducible samples

The resolution when imaging in 3D is not limited by the dynamic range. Many images are combined and it is enough that some images cover the low-resolution region of diffraction space and others cover the high-resolution region. Even with already existing data, much higher resolution could therefore be achieved. What is needed is a strong effort to evaluate and improve the algorithms for data analysis. I am convinced that this will lead to an explosion of new structures both at higher resolution and of smaller objects.

Separation of conformational states

Imaging one particle at a time has the potential of separating conformational states and thus to image the entire ensemble of structures of a protein. Combined with the possibility to collect data at a very quick rate, this could become a powerful technique to study not only unknown proteins but also understand the dynamics of already known proteins .

The ability to image without the need for crystallization combined with the revolutionizing ability to separate conformational states promises a very interesting future for structural biology.

10. Sammanfattning på svenska

I denna avhandling försöker vi öppna nya dörrar inom mikroskopin genom att använda en ny typ ljuskälla, fri elektronlasern. Att kunna se och avbilda små saker har inte bara länge fascinerat människor utan har även lett till några av de största vetenskapliga genombrotten så som upptäckten av celler, bakterier och DNA.

Hur små saker man kan se med ett mikroskop begränsas av våglängden på det ljus man använder. Med synligt ljus går det inte att se detaljer mindre än några hundra nanometer (en nanometer är en miljondels millimeter). Av denna anledning används istället ofta röntgenstrålar med vilka det går att se detaljer mindre än en tiondels nanometer, vilket är tillräckligt för att kunna urskilja enskilda atomer. Nackdelen är dock dels att de inte går att tillverka effektiva linser för röntgenstrålar samt att det mesta av strålningen går rakt igenom det prov man studerar och alltså väldigt lite reagerar med provet och bidrar till att ge oss information om det.

Trots dessa begränsningar så har forskare i mer än 50 år använt röntgenstrålar i metoden röntgenkristallografi för att studera atomstrukturen hos molekyler med många tusen atomer (så kallade makromolekyler) som exempelvis proteiner. En kristall skapas av flera miljarder identiska molekyler och trots att varje molekyl bara reagerar väldigt svagt med röntgenstrålarna så ger de tillsammans en tillräckligt stark signal för att göra det möjligt att räkna ut positionen hos varje atom. Istället för en lins används kraftiga datoralgoritmer som kan räkna ut hur molekylen såg ut utifrån röntgenstrålarnas diffraktion, alltså hur de sprids när de träffar molekylen. Över 70000 proteiner har avbildats med den här metoden och vår förståelse för hur livet fungerar har ökat enormt tack vara denna information.

Den största svårigheten inom röntgenkristallografi är att skapa kristallen. Det tar ofta lång tid och många proteiner har hittills ingen lyckats kristallisera. Det innebär att det finns ett gap i vår förståelse av biologin på grund av att vi inte kan se strukturen hos dessa okristalliserbara molekyler.

En lovande lösning på det här problemet ges av en ny typ av röntgenkälla; fri elektronlasern. En fri elektronlaser producerar röntgenljuspulser som är 10 miljarder gånger kraftigare än någon tidigare röntgenkälla. Med denna extremt kraftiga strålning kan det gå att få en tillräckligt stark signal från ett enda protein och vi skulle därför kunna kringgå behovet av kristaller. När en molekyl träffas av så kraftig röntgenstrålning så kommer den mycket snabbt att förvandlas till ett plasma och explodera, så frågan är hur vår bild av molekylen påverkas av detta. Ser vi molekylen som den såg ut innan, efter eller under explosionen? Nyckeln till svaret är att röntgenpulserna från fri elektronlasern är extremt korta, under 100 femtosekunder (alltså under en tiomiljondel av en miljondels sekund). Denna tid är kortare än den tid det tar för molekylen att explodera och eftersom molekylen då exploderar först efter att röntgenljuset har passerat

så kommer bilden vara opåverkad av explosionen. Denna princip kallas diffraktion före destruktions.

Då det naturligt är svårare att se små saker än stora så testades inte metoden direkt på de proteiner som inspirerat den. Istället gjordes de första försöken på konstgjorda föremål ca 4000 nanometer stora. Dessa gjordes 2006 och bevisade att metoden fungerar.

I den här avhandlingen ingår de första avbildningarna av biologiska prover, dels av bakterier som är ca 1000 nanometer stor och viruset mimivirus som med sina 450 nanometer är ett av världens största virus. De datoralgoritmer som används inom röntgenkristallografi fungerar inte i vårt fall och därför har vi i Uppsala utvecklat ett nytt programpaket, Hawk, för analys av diffraktionsdata från enskilda partiklar som användes för dessa avbildningar.

Dessa avbildningar gav tvådimensionella bilder och för att kunna avbilda något i 3D behövs många bilder ur många olika vinklar. Eftersom provet exploderar när det träffas av röntgenpulsen är det omöjligt att ta flera bilder av samma prov. Om det finns flera identiska kopior av provet så går det dock att behandla bilder av dem som om de kom från samma prov. Proteiner och virus har ofta den egenskapen och går därför att avbilda i 3D.

För att lyckas återstår dock svårigheten att ta reda på från vilket håll bilden av provet är taget. Proverna sprutas in i en vakuumkammare där de träffas av en röntgenpuls och vi vet inte vilken rotation varje prov hade när det träffades. Rotationen måste alltså räknas fram endast genom att studera diffraktionsdatan. Vi lyckades med detta för data från mimiviruset där 261 diffraktionsbilder kombinerades till en tredimensionell bild. Detta har gett oss ny information om viruset, bland annat så ser vi en oväntad förskjutning av massa mot ena sidan av viruset. Denna massa är troligen virusets DNA och det betyder att det måste finnas en okänd mekanism som hindrar DNAt från att sprida sig till andra delar av viruset. Resultatet är även viktigt eftersom det visar att rotationsproblemet kan lösas i praktiken.

Jag är övertygad om att vi tack vare fri elektronlasrar snart kommer att se avbildningar i 3D av mindre virus och även av proteiner. Utvecklingen av både lasrarna själva och algoritmerna som används för att analysera datan går mycket fort och forskningsfältet fortsätter att attrahera duktiga forskare. Detta kan äntligen fylla det gap inom strukturell biologi som kravet på kristallisation har skapat.

Acknowledgements

Many persons have helped and supported me in my work with this thesis, both directly and indirectly. In particular I want to thank:

Janos, for accepting me as a student and for supporting me all the way through. Thank you for inspiring me to tackle seemingly impossible problems and trusting me on the ones that were truly impossible.

David, my co-supervisor, for the encouragement and for constant computer support.

Filipe, for being an amazing officemate, discussion partner and mentor through the early years.

Bianca for supporting me ever since the detector building days in a SLAC basement and **Daniel** for sharing all these years in Uppsala.

Max and **Gijs**, for each of you writing an awesome master thesis. It's been great sharing this office with you and I apologize for the noise.

Duane and **Andrew**, for some of the most interesting discussions I've had, both over Skype and during countless hours of beamtime nightshifts. **Anton** was not on the Skype calls but deserves to be thanked nonetheless.

Nic for great discussions and making my tea breaks so much more enjoyable.

Marvin, Malin, Erik, Gunilla, Gösta, Martin, Inger, Jessica, Calle, Olof, Asawari, Daniel, Jakob, Dirk, Michiel, Karin and all other past and present members of the lab. Thank you for trying to teach me biology, checking my equations and making BMC an awesome place to spend my days.

All my family and friends, for reminding me there is a life outside the academic one.

Ann and **Örjan**, for encouraging my curiosity, for teaching me about quarks at the age of 12, for giving me my first set of lego and for having a really awesome bathroom tiling.

And **Maria**, for the million things I love about you.

Authors contribution

Paper I

I performed the entire data analysis, except for the automated screening of diffraction patterns. My work included preprocessing, software development, determining orientations and phase recovery.

Paper II

I performed the data analysis including pattern selection, phase recovery, identification and handling of missing modes.

Paper III

I wrote the paper and performed the entire analysis and theoretical work. Filipe Maia contributed by very fruitful discussions.

Paper IV

Me Filipe Maia did the theoretical calculations and demonstrations together.

Paper V

Over a number of years I have made multiple contributions to the Hawk software package.

Paper VI

I performed the data analysis for the *Prochlorococcus marinu* and *Synechococcus elongatus* samples.

References

- [1] Anton Barty, Carl Caleman, Andrew Aquila, Nicusor Timneanu, Lukas Lomb, Thomas A. White, Jakob Andreasson, David Arnlund, Sasa Bajt, Thomas R. M. Barends, Miriam Barthelmess, Michael J. Bogan, Christoph Bostedt, John D. Bozek, Ryan Coffee, Nicola Coppola, Jan Davidsson, Daniel P. DePonte, R. Bruce Doak, Tomas Ekeberg, Veit Elser, Sascha W. Epp, Benjamin Erk, Holger Fleckenstein, Lutz Foucar, Petra Fromme, Heinz Graafsma, Lars Gumprecht, Janos Hajdu, Christina Y. Hampton, Robert Hartmann, Andreas Hartmann, Guenter Hauser, Helmut Hirsemann, Peter Holl, Mark S. Hunter, Linda Johansson, Stephan Kassemeyer, Nils Kimmel, Richard A. Kirian, Mengning Liang, Filipe R. N. C. Maia, Erik Malmerberg, Stefano Marchesini, Andrew V. Martin, Karol Nass, Richard Neutze, Christian Reich, Daniel Rolles, Benedikt Rudek, Artem Rudenko, Howard Scott, Ilme Schlichting, Joachim Schulz, M. Marvin Seibert, Robert L. Shoeman, Raymond G. Sierra, Heike Soltau, John C. H. Spence, Francesco Stellato, Stephan Stern, Lothar Strueder, Joachim Ullrich, X. Wang, Georg Weidenspointner, Uwe Weierstall, Cornelia B. Wunderer, and Henry N. Chapman. Self-terminating diffraction gates femtosecond X-ray nanocrystallography measurements. *Nature Photonics*, 6(1):35–40, Jan 2012.
- [2] HH Bauschke, PL Combettes, and DR Luke. Hybrid projection-reflection method for phase retrieval. *Journal Of The Optical Society Of America A-Optics Image Science And Vision*, 20(6):1025–1034, Jun 2003.
- [3] CM Bishop, M Svensen, and CKI Williams. GTM: The generative topographic mapping. *Neural Computation*, 10(1):215–234, JAN 1 1998.
- [4] Michael J. Bogan, W. Henry Benner, Sebastien Boutet, Urs Rohner, Matthias Frank, Anton Barty, M. Marvin Seibert, Filipe Maia, Stefano Marchesini, Sasa Bajt, Bruce Woods, Vincent Riot, Stefan P. Hau-Riege, Martin Svenda, Erik Marklund, Eberhard Spiller, Janos Hajdu, and Henry N. Chapman. Single particle X-ray diffractive imaging. *Nano Letters*, 8(1):310–316, Jan 2008.
- [5] Max Born and Emil Wolf. *Principles of Optics*. Cambridge, 2006.
- [6] Gabor Bortel and Miklos Tegze. Common arc method for diffraction pattern orientation. *Acta Crystallographica Section A*, 67(Part 6):533–543, Nov 2011.
- [7] R.N. Bracewell. Strip integration in radio astronomy. *Australian Journal of Physics*, 9:198–217, 1956.
- [8] Henry N. Chapman, Anton Barty, Michael J. Bogan, Sebastien Boutet, Matthias Frank, Stefan P. Hau-Riege, Stefano Marchesini, Bruce W. Woods, Sasa Bajt, Henry Benner, Richard A. London, Elke Ploenjes, Marion Kuhlmann, Rolf Treusch, Stefan Duesterer, Thomas Tschentscher, Jochen R. Schneider, Eberhard Spiller, Thomas Moeller, Christoph Bostedt, Matthias Hoener, David A. Shapiro, Keith O. Hodgson, David Van der Spoel, Florian Burmeister, Magnus Bergh, Carl Caleman, Goesta Hultdt, M. Marvin Seibert, Filipe R. N. C.

- Maia, Richard W. Lee, Abraham Szoek, Nicusor Timneanu, and Janos Hajdu. Femtosecond diffractive imaging with a soft-X-ray free-electron laser. *Nature Physics*, 2(12):839–843, Dec 2006.
- [9] Henry N. Chapman, Petra Fromme, Anton Barty, Thomas A. White, Richard A. Kirian, Andrew Aquila, Mark S. Hunter, Joachim Schulz, Daniel P. DePonte, Uwe Weierstall, R. Bruce Doak, Filipe R. N. C. Maia, Andrew V. Martin, Ilme Schlichting, Lukas Lomb, Nicola Coppola, Robert L. Shoeman, Sascha W. Epp, Robert Hartmann, Daniel Rolles, Artem Rudenko, Lutz Foucar, Nils Kimmel, Georg Weidenspointner, Peter Holl, Mengning Liang, Miriam Barthelmeß, Carl Caleman, Sebastien Boutet, Michael J. Bogan, Jacek Krzywinski, Christoph Bostedt, Sasa Bajt, Lars Gumprecht, Benedikt Rudek, Benjamin Erk, Carlo Schmidt, Andre Hoemke, Christian Reich, Daniel Pietschner, Lothar Strueder, Guenter Hauser, Hubert Gorke, Joachim Ullrich, Sven Herrmann, Gerhard Schaller, Florian Schopper, Heike Soltau, Kai-Uwe Kuehnel, Marc Messerschmidt, John D. Bozek, Stefan P. Hau-Riege, Matthias Frank, Christina Y. Hampton, Raymond G. Sierra, Dmitri Starodub, Garth J. Williams, Janos Hajdu, Nicusor Timneanu, M. Marvin Seibert, Jakob Andreasson, Andrea Rocker, Olof Joansson, Martin Svenda, Stephan Stern, Karol Nass, Robert Andritschke, Claus-Dieter Schroeter, Faton Krasniqi, Mario Bott, Kevin E. Schmidt, Xiaoyu Wang, Ingo Grotjohann, James M. Holton, Thomas R. M. Barends, Richard Neutze, Stefano Marchesini, Raimund Fromme, Sebastian Schorb, Daniela Rupp, Marcus Adolph, Tais Gorkhover, Inger Andersson, Helmut Hirsemann, Guillaume Potdevin, Heinz Graafsma, Bjoern Nilsson, and John C. H. Spence. Femtosecond X-ray protein nanocrystallography. *Nature*, 470(7332):73–U81, Feb 2011.
- [10] Jean-Michel Claverie and Chantal Abergel. Mimivirus and its Virophage. *Annual Review Of Genetics*, 43:49–66, 2009.
- [11] Jean-Michel Claverie and Chantal Abergel. Mimivirus: the emerging paradox of quasi-autonomous viruses. *Trends In Genetics*, 26(10):431–437, Oct 2010.
- [12] Ronald R. Coifman and Stephane Lafon. Diffusion maps. *Applied And Computational Harmonic Analysis*, 21(1):5–30, Jul 2006.
- [13] John M. Cowley. *Diffraction Physics*. North-Holland Physics Publishing, 1981.
- [14] Marina de Leeuw, Shlomi Reuveni, Joseph Klafter, and Rony Granek. Coexistence of Flexibility and Stability of Proteins: An Equation of State. *Plos One*, 4(10), Oct 2009.
- [15] D. P. DePonte, U. Weierstall, K. Schmidt, J. Warner, D. Starodub, J. C. H. Spence, and R. B. Doak. Gas dynamic virtual nozzle for generation of microscopic droplet streams. *Journal Of Physics D-Applied Physics*, 41(19), Oct 2008.
- [16] G. Eckart, C.; Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–8, 1936.
- [17] V Elser. Phase retrieval by iterated projections. *Journal Of The Optical Society Of America A-Optics Image Science And Vision*, 20(1):40–55, Jan 2003.
- [18] L Euler. Formvlae generales pro translatione qvacvnqve corporvm rigidorvm. *Novi Commentarii academiae scientiarum Petropolitanae*, 20:189–207, 1776.
- [19] P Ewald. Introduction to dynamical theory of x-ray diffraction. *Acta Crystallographica Section A-Crystal Physics Diffraction Theoretical And*

- General Crystallography*, A 25(Part 1):103, 1969.
- [20] NA Farrow and FP Ottensmeyer. A-posteriori determination of relative projection directions of arbitrarily oriented macromolecules. *Journal Of The Optical Society Of America A-Optics Image Science And Vision*, 9(10):1749–1760, Oct 1992.
- [21] JR Fienup. Reconstruction of an object from modulus of its fourier-transform. *Optics Letters*, 3(1):27–29, 1978.
- [22] Russell Fung, Valentin Shneerson, Dilano K. Saldin, and Abbas Ourmazd. Structure from fleeting illumination of faint spinning objects in flight. *Nature Physics*, 5(1):64–67, Jan 2009.
- [23] Robert Hartmann, Sascha Epp, Sven Herrmann, Peter Holl, Norbert Meidinger, Christian Reich, Daniel Rolles, Heike Soltau, Lothar Strueder, Joachim Ullrich, and Georg Weidenspointner. Large Format PnCCDs as Imaging Detectors for X-ray Free-Electron-Lasers. In *2008 Ieee Nuclear Science Symposium And Medical Imaging Conference (2008 Nss/Mic), Vols 1-9*, Ieee Nuclear Science Symposium - Conference Record, pages 1865–1870, 345 E 47TH ST, New York, NY 10017 USA, 2009. IEEE Nucl & Plasma Sci Soc; Natl Nucl Secur Adm; US Def Threat Reduct Agcy; ICx Radiat GmbH; ORTEC; Hamamatsu; European Phys Journal; Hilger Crystals; SAFC Hitech; ATOMTEX; Canberra; SensL; ASP; Brookhaven Natl Lab; CEA; CERN; DESY; Forschungszentrum Julich GmbH; Int Atom Energy Agcy; Lawrence Livermore Natl Lab, IEEE. IEEE Nuclear Science Symposium/Medical Imaging Conference, Dresden, Germany, Oct 19-25, 2008.
- [24] BL Henke, EM Gullikson, and JC Davis. X-ray interactions - photoabsorption, scattering, transmission, and reflection at $e=50\text{-}30,000$ ev, $z=1\text{-}92$. *Atomic Data And Nuclear Data Tables*, 54(2):181–342, Jul 1993.
- [25] T Kohonen. The self-organizing map. *Proceedings Of The IEEE*, 78(9):1464–1480, Sep 1990.
- [26] B La Scola, S Audic, C Robert, L Jungang, X de Lamballerie, M Drancourt, R Birtles, JM Claverie, and D Raoult. A giant virus in amoebae. *Science*, 299(5615):2033, Mar 2003.
- [27] N. D. Loh, M. J. Bogan, V. Elser, A. Barty, S. Boutet, S. Bajt, J. Hajdu, T. Ekeberg, F. R. N. C. Maia, J. Schulz, M. M. Seibert, B. Iwan, N. Timneanu, S. Marchesini, I. Schlichting, R. L. Shoeman, L. Lomb, M. Frank, M. Liang, and H. N. Chapman. Cryptotomography: Reconstructing 3d fourier intensities from randomly oriented single-shot diffraction patterns. *Phys. Rev. Lett.*, 104:225501, Jun 2010.
- [28] Ne-Te Duane Loh and Veit Elser. Reconstruction algorithm for single-particle diffraction imaging experiments. *Physical Review E*, 80(2), Aug 2009.
- [29] DR Luke. Relaxed averaged alternating reflections for diffraction imaging. *Inverse Problems*, 21(1):37–50, Feb 2005.
- [30] S. Marchesini. Invited article: A unified evaluation of iterative projection algorithms for phase retrieval (vol 78, art no 011301, 2007). *Review Of Scientific Instruments*, 78(4), Apr 2007.
- [31] S Marchesini. Ab initio compressive phase retrieval. 2008. arXiv:0809.2006v1.
- [32] S. Marchesini, H. He, H. N. Chapman, S. P. Hau-Riege, A. Noy, M. R. Howells, U. Weierstall, and J. C. H. Spence. X-ray image reconstruction from a

- diffraction pattern alone. *Phys. Rev. B*, 68:140101, Oct 2003.
- [33] Stefano Marchesini. Phase retrieval and saddle-point optimization. *Journal Of The Optical Society Of America A-Optics Image Science And Vision*, 24(10):3289–3296, Oct 2007.
- [34] Erik G. Marklund, Daniel S. D. Larsson, David van der Spoel, Alexandra Patriksson, and Carl Caleman. Structural stability of electrosprayed proteins: temperature and hydration effects. *Physical Chemistry Chemical Physics*, 11(36):8069–8078, 2009.
- [35] A. V. Martin, F. Wang, N. D. Loh, T. Ekeberg, F. R. N. C. Maia, M. Hantke, G. van der Schot, C. Y. Hampton, R. G. Sierra, A. Aquila, S. Bajt, M. Barthelmess, C. Bostedt, J. D. Bozek, N. Coppola, S. W. Epp, B. Erk, H. Fleckenstein, L. Foucar, M. Frank, H. Graafsma, L. Gumprecht, A. Hartmann, R. Hartmann, G. Hauser, H. Hirsemann, P. Holl, S. Kassemeyer, N. Kimmel, M. Liang, L. Lomb, S. Marchesini, K. Nass, E. Pedersoli, C. Reich, D. Rolles, B. Rudek, A. Rudenko, J. Schulz, R. L. Shoeman, H. Soltau, D. Starodub, J. Steinbrener, F. Stellato, L. Strüder, J. Ullrich, G. Weidenspointner, T. A. White, C. B. Wunderer, A. Barty, I. Schlichting, M. J. Bogan, and H. N. Chapman. Noise-robust coherent diffractive imaging with a single diffraction pattern. *Opt. Express*, 20(15):16650–16661, Jul 2012.
- [36] Brian Moths and Abbas Ourmazd. Bayesian algorithms for recovering structure from single-particle diffraction snapshots of unknown orientation: a comparison. *Acta Crystallographica Section A*, 67(Part 5):481–486, Sep 2011.
- [37] R Neutze, R Wouts, D van der Spoel, E Weckert, and J Hajdu. Potential for biomolecular imaging with femtosecond X-ray pulses. *Nature*, 406(6797):752–757, Aug 2000.
- [38] G Oszlanyi and A Suto. Ab initio structure solution by charge flipping. *Acta Crystallographica Section A*, 60(Part 2):134–141, Mar 2004.
- [39] Alexandra Patriksson, Erik Marklund, and David van der Spoel. Protein structures under electrospray conditions. *Biochemistry*, 46(4):933–945, Jan 2007.
- [40] Hugh T. Philipp, Lucas J. Koerner, Marianne S. Hromalik, Mark W. Tate, and Sol M. Gruner. Femtosecond Radiation Experiment Detector for X-Ray Free-Electron Laser (XFEL) Coherent X-Ray Imaging. *IEEE Transactions On Nuclear Science*, 57(6):3795–3799, Dec 2010.
- [41] D Raoult, S Audic, C Robert, C Abergel, P Renesto, H Ogata, B La Scola, M Suzan, and JM Claverie. The 1.2-megabase genome sequence of mimivirus. *Science*, 306(5700):1344–1350, Nov 2004.
- [42] W Rudin. *Principles of Mathematical Analysis*. McGraw-Hill, 1976.
- [43] D Sayre. Some implications of a theorem due to shannon. *Acta Crystallographica*, 5(6):843, 1952.
- [44] P. Schwander, R. Fung, G. N. Phillips, Jr., and A. Ourmazd. Mapping the conformations of biological assemblies. *New Journal Of Physics*, 12, Mar 2010.
- [45] Paul Shannon. Communication theory of secrecy systems. *Bell System Technical Journal*, 28(4):656–715, 1949.
- [46] V. L. Shneerson, A. Ourmazd, and D. K. Saldin. Crystallography without crystals. I. The common-line method for assembling a three-dimensional diffraction volume from single-particle scattering. *Acta Crystallographica*

Section A, 64(Part 2):303–315, Mar 2008.

- [47] Miklos Tegze and Gabor Bortel. Atomic structure of a single large biomolecule from diffraction patterns of random orientations. *Journal of Structural Biology*, 179(1):41 – 45, 2012.
- [48] A. Vretblad. *Fourier analysis and its applications*. Graduate texts in mathematics. Springer, 2003.
- [49] Chuan Xiao, Yuri G. Kuznetsov, Siyang Sun, Susan L. Hafenstein, Victor A. Kostyuchenko, Paul R. Chipman, Marie Suzan-Monti, Didier Raoult, Alexander McPherson, and Michael G. Rossmann. Structural Studies of the Giant Mimivirus. *PLOS Biology*, 7(4):958–966, Apr 2009.
- [50] L. Young, E. P. Kanter, B. Kraessig, Y. Li, A. M. March, S. T. Pratt, R. Santra, S. H. Southworth, N. Rohringer, L. F. DiMauro, G. Doumy, C. A. Roedig, N. Berrah, L. Fang, M. Hoener, P. H. Bucksbaum, J. P. Cryan, S. Ghimire, J. M. Glownia, D. A. Reis, J. D. Bozek, C. Bostedt, and M. Messerschmidt. Femtosecond electronic response of atoms to ultra-intense X-rays. *Nature*, 466(7302):56–U66, Jul 2010.

Acta Universitatis Upsaliensis

*Digital Comprehensive Summaries of Uppsala Dissertations
from the Faculty of Science and Technology 960*

Editor: The Dean of the Faculty of Science and Technology

A doctoral dissertation from the Faculty of Science and Technology, Uppsala University, is usually a summary of a number of papers. A few copies of the complete dissertation are kept at major Swedish research libraries, while the summary alone is distributed internationally through the series Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Science and Technology.



ACTA
UNIVERSITATIS
UPSALIENSIS
UPPSALA
2012

Distribution: publications.uu.se
urn:nbn:se:uu:diva-179643