

Fast Reconstruction Method for Diffraction Imaging

Eliyahu Osherovich, Michael Zibulevsky, and Irad Yavneh

Technion — Israel Institute of Technology, Computer Science Department
`{oeli,mzib,irad}@cs.technion.ac.il`

Abstract. We present a fast image reconstruction method for two- and three-dimensional diffraction imaging. Provided that very little information about the phase is available, the method demonstrates convergence rates that are several orders of magnitude faster than current reconstruction techniques. Unlike current methods, our approach is based on convex optimization. Besides fast convergence, our method allows great deal of flexibility in choosing most appropriate objective function as well as introducing additional information about the sought signal, e.g., smoothness. Benefits of good choice of the objective function are demonstrated by reconstructing an image from noisy data.

1 Introduction

Determination of the structure of molecules is an important task in a variety of disciplines including biology, chemistry, medicine, and others. Currently, one of the major techniques for such an imaging is X-ray crystallography. However, its applicability is limited to the cases when a good quality crystals are available. Unfortunately, biological specimens such as cells, viruses, and many important macromolecules are difficult to crystallize. One of the promising alternatives to obtain the structure of single biomolecules or cells is the rapidly developing technique of X-ray diffraction microscopy also known as coherent diffraction imaging (CDI) [1]. In CDI a highly coherent beam of X-rays is scattered by electrons of the specimen, generating a diffraction pattern which, after being captured by a CCD sensor, is used to reconstruct the electron density of the specimen. It can be shown that under certain conditions the diffracted wavefront is approximately equal (within a scale factor) to the Fourier transform (FT) of the specimen. However, due to the physical nature of the sensor, the phase of the diffracted wave is lost. Hence, the problem is equivalent to reconstruction of an image (denoted by x) from the magnitude of its FT (denoted by r). The same problem is met in astronomy and X-ray crystallography where it is often called the phase (retrieval) problem.

In its classical form the phase retrieval problem assumes that the phase information is lost completely. However, current reconstruction methods do require additional information about the specimen support. Therefore, they either obtain the support information from other sources, e.g., low-resolution images or

attempt to reconstruct the support along with the phase retrieval using some sort of bootstrapping technique [2]. In this work we consider a situation where the additional information available is the phase uncertainty limits. Discussion on possible ways to obtain such information is postponed till Sect. 5. At this moment we assume that a rough estimation of the phase is available, i.e., the phase uncertainty is slightly less than π radians. In this case we give an efficient convex optimization method of solution.

The rest of the paper is organized as follows. In Sect. 2 we give a short overview of the previous work. This overview is divided into two main subsections. Firstly, the theoretical foundations are given in Sect. 2.1 where we chiefly consider the uniqueness of the solution. Secondly, we review current reconstruction methods in Sect. 2.2. In Sect. 3 we present our convex optimization procedure. Quantitative results of our simulations are presented in Sect. 4, followed by discussions and conclusions in Sect. 5.

2 Previous Work

Since a very important part of the information is lost it is not obvious that the image can be reconstructed from the FT magnitude alone. Therefore, several researches were devoted to the theoretical and practical aspects of the problem. In the subsequent sections we give a short review of major results and methods.

2.1 Theoretical Foundations

It has been shown [3] that, under mild restrictions, a finite-support one-dimensional or multidimensional image is uniquely specified to within a scale factor by the tangent of its FT phase. The FT magnitude, in contrast, does not uniquely specify a one-dimensional image. For multidimensional images, the FT magnitude is almost always sufficient to specify the image to within a trivial transformation. Namely, the set of images for which it is insufficient is of measure zero [3]. And by the trivial transformations we mean a combination of coordinate reversal, sign change, and translation.

As follows from the above results, the uniqueness properties of the reconstruction of one-dimensional and multidimensional images are quite distinct. In addition, the phase and the magnitude of the FT contribute differently to the uniqueness of reconstruction. By combining the FT magnitude with some phase information we can further restrict the uncertainty in the reconstruction result. It has been shown [4] that, under mild restrictions, the signed FT magnitude, i.e., the magnitude plus one bit of phase information per wavelength, (i.e., phase uncertainty of π radians) is sufficient to specify one-dimensional and multidimensional images uniquely.

2.2 Current Reconstruction Methods

Current reconstruction methods date back to the pioneering work of Gerchberg and Saxton (GS) [5]. Later, their method was significantly improved by




Fig. 1. Schematic of the Gerchberg-Saxton algorithm

Fienup [6] who suggested the Hybrid Input Output (HIO) algorithm. Since the later algorithm (HIO) is significantly better than the GS method we limit ourselves to this method in our simulations and comparisons.

The GS algorithm utilizes alternating projections on the constraints sets. The algorithm performs iterative Fourier transformations back and forth between the object and Fourier domains. In each domain it applies the known data (constraints). In our case, it is the modulus of a complex number in the Fourier domain and non-negativity in the object domain. Schematic of the GS algorithm is given in Fig. 1. The HIO algorithm is identical to the GS method in the Fourier domain, however, their behavior in the object domain differs. In the GS algorithm we have

$$x_{k+1}(t) = \begin{cases} x'_k(t), & t \notin \mathcal{V} \\ 0, & t \in \mathcal{V} \end{cases}. \quad (1)$$

While in the HIO the update rule is given by

$$x_{k+1}(t) = \begin{cases} x'_k(t), & t \notin \mathcal{V} \\ x_k(t) - \beta x'_k(t), & t \in \mathcal{V} \end{cases}, \quad (2)$$

where $\beta \in [0, 1]$, \mathcal{V} is the set of points at which $x'(t)$ violates the object domain constraints (positivity). x' is obtained from x by enforcing the FT magnitude as shown in Fig. 1.

It has long been known that the GS algorithm is equivalent to the gradient projection method with constant step length for the problem defined by (3) [6]. Recently, it has been shown that in this case the gradient is an eigenvector of the Hessian with the corresponding eigenvalue equal to two [7]. Thus, the GS method is, in fact, the Newton method combined with projections. To our knowledge, Fienup's HIO algorithm has no such counterpart in the convex optimization land. Despite the remarkable success of these two algorithms, they have several drawbacks. First, they require a good estimation of the support. Second, their flexibility is limited, e.g., it seems that the algorithms cannot be adapted to a particular noise distribution in the measurements nor can additional prior knowledge be easily integrated into the scheme.

3 Convex Optimization Methods

First, we must clarify that the problem is not convex. Thus, the term “convex optimization” is probably misleading. We do perform a convex relaxation at some stage and it turns out to be crucial for successful reconstruction. However, the problem itself remains non-convex. Hence, in this paper by “convex optimization methods” we mean the classical gradient and Newton-type methods.

Let us start by formulating the optimization problem. The very common formulation is as follows

$$\begin{cases} \min & \|Fx - r\|^2 \\ \text{s.t.} & x \geq 0 \end{cases}, \quad (3)$$

where F denotes the FT operator (a matrix in the discrete case). Of course, there is an endless number of ways to choose the objective function. A particular choice of it may affect the convergence speed and numerical stability. However, in our view, it is more important to choose the objective function that properly reflects the underlying physics phenomena. For example, the choice of Equation (3) is suitable when the measured quantity is r and the noise in the measurements has Gaussian distribution with zero mean. In practice, however, we measure r^2 and not r and noise distribution is Poissonian rather than Gaussian. The corresponding objective function and its influence on the reconstruction quality is shown in Sect. 4.

Unfortunately, the phase retrieval problem turns out to be particularly tough for convex optimization methods. For example, it was reported that naive application of Newton-type algorithms to the problem formulated by Equation (3) fails, except for tiny problems [8]. This failure is due to the high non-linearity and non-convexity of the problem. To our knowledge, no globally convergent convex optimization method exists at the moment. Despite this general failure, we demonstrate in this work, that the situation can change dramatically if additional phase information is available.

Let us consider one pixel in the Fourier domain. In case the phase is known to lie within a certain interval $[\alpha, \beta]$, the correct complex number must belong to the arc defined by α and β as depicted in Fig. 2a. Despite this additional information, the problem still remains non-convex and cannot be solved efficiently. However, we perform *convex relaxation*. That is, we relax our requirements on the modulus and let the complex number lie in the convex region \mathcal{C} as shown in Fig. 2b.

The problem now becomes convex. Its formal definition is as follows

$$\begin{cases} \min & \|d(Fx, \mathcal{C})\|^2 \\ \text{s.t.} & x \geq 0 \end{cases} \quad (4)$$

where $d(a, \mathcal{C})$ denotes the Euclidean distance from point a to the convex set \mathcal{C} . In our experience, several dozen iterations are sufficient to solve this convex problem (see Fig. 5a). Of course, the solution does not match the original image because both the phase and magnitude may vary significantly. However, we suggest the following method for the solutions of the original problem. Stage 1:




Fig. 2. Convex relaxation

starting at random x_0 , we solve the problem defined by Equation (4). Stage 2: the solution obtained in Stage 1 (denoted by x_1) is used as the starting point for the minimization problem that combines both the convex and non-convex parts, as defined below

$$\begin{cases} \min & \||Fx| - r\|^2 + \|d(Fx, \mathcal{C})\|^2 \\ \text{s.t.} & x \geq 0 \end{cases} \quad (5)$$

More precisely, in our implementation we use the unconstrained minimization formulation, i.e., instead of Equations (4) and (5) we minimize the following functionals

$$E_c(x) = \|d(Fx, \mathcal{C})\|^2 + \|[x]_-\|^2 \quad (6)$$

$$E(x) = \||Fx| - r\|^2 + \|d(Fx, \mathcal{C})\|^2 + \|[x]_-\|^2 \quad (7)$$

where $[x]_-$ is defined as follows

$$[x]_- = \begin{cases} 0, & x \geq 0 \\ x, & x < 0 \end{cases} \quad (8)$$

Results of our simulations are given in the next section.

4 Simulations and Results

Due to the high dimensionality of the problem (especially in the 3D case) we shall limit our choice to the methods that do not keep the Hessian matrix or its approximation. Hence, in our implementation we used a modified version of the SESOP algorithm [9] and the LBFGS method [10]. Both algorithms demonstrated very close results. The main difference is that SESOP guarantees that there are two Fourier transforms per iteration just like in the GS and HIO methods. The LBFGS method, on the other hand, cannot guarantee that. However, in practice the number of the Fourier transforms is very close to that of SESOP and HIO.

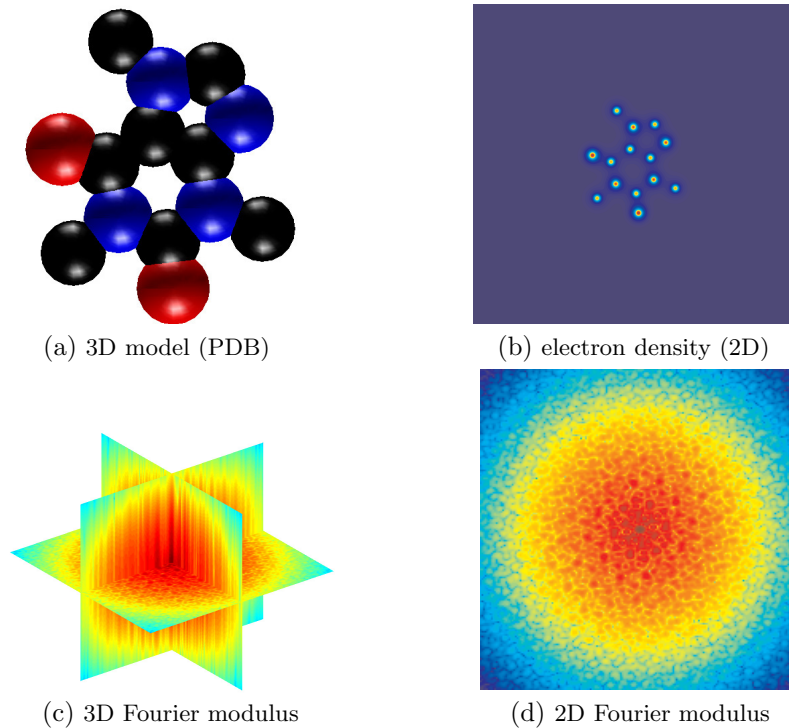


Fig. 3. Caffeine molecule

The method was tested across a variety of data. In this section we present some of these examples. The first example is a molecule of caffeine whose 3D model, a 2D projection of its electron density, and the corresponding FT modulus are shown in Fig. 3. This information was obtained from a PDB¹ (protein database) file. In addition, we use a “natural” image which represents a class of images with rich texture and tight support. Moreover, it may be easier to estimate the visual reconstruction quality of such images. This image and its Fourier modulus are given in Fig. 4.

Note that we assume that a rectilinear sampling is available in the 3D case. In practice, however, the sensors measure a two-dimensional slice of the 3D volume. Provided that a sufficient number of such slices were measured an interpolation can be used to form a rectilinear array of measurements [11]. However, the slices can be incorporated directly into our minimization scheme. It will be addressed in future work.

In our experiments we used the phase uncertainty of 3 radians. The bounds were chosen at random such that the true phase has uniform distribution inside the interval. The starting point (x_0) was also chosen randomly. Of course, there

¹ see <http://www.pdb.org> for more information.

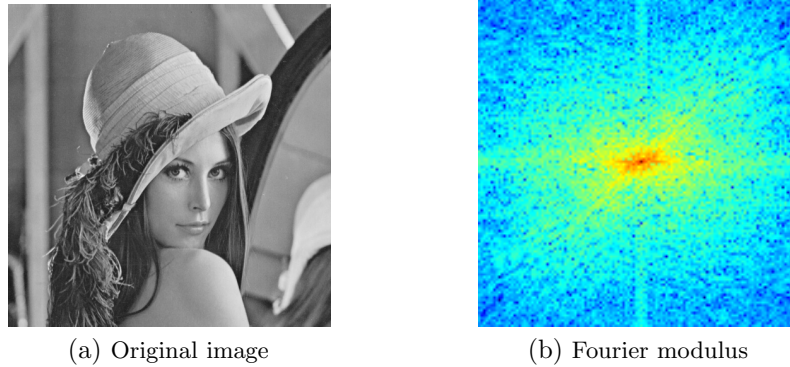


Fig. 4. A natural image (Lena)

is an obvious way to make a more educated guess: by choosing the middle of the uncertainty interval. However, our experiments indicate that the starting point has little influence. In all cases the reconstructed images obtained with our method were visually indistinguishable from the original. Therefore, we present the values of $E_c(x)$ and $E(x)$ as defined in Equations (6) and (7) to visualize the progress of the first and the second stages respectively. The second stage is compared with the HIO algorithm for which the error term is $E(x)$ without the phase bounds constraint, i.e.,

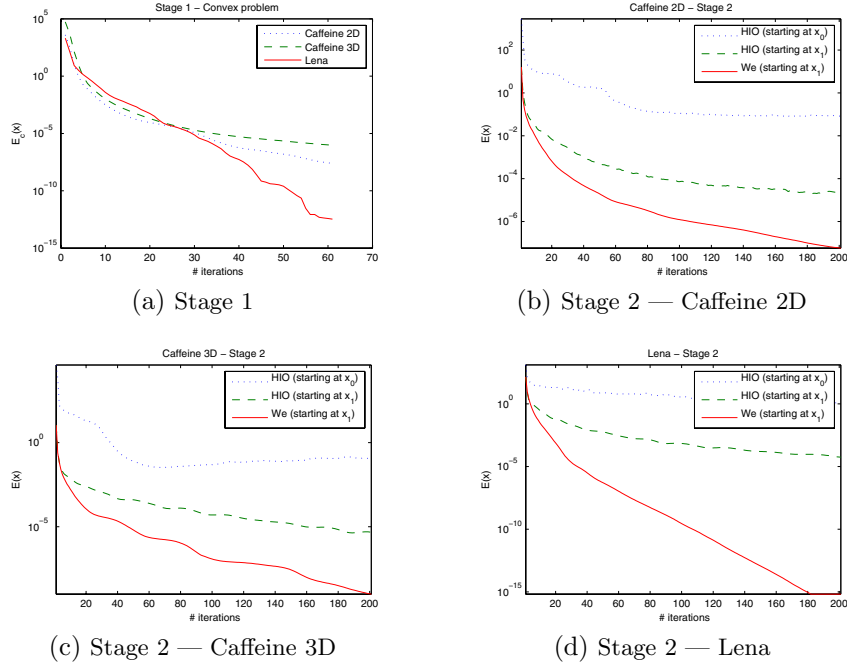
$$E_{\text{HIO}} = \||Fx| - r\|^2 + \|[x]_-\|^2 . \quad (9)$$

The first experiment is as follows. First, we ran 60 iterations of Stage 1. The progress of different images is shown in Fig. 5a. In the second stage we ran 200 iterations of our algorithm (SESOP) starting at the solution obtained in the previous stage (x_1). To compare the convergence rate with the current methods, we ran twice the HIO algorithm. Once, starting at x_0 and thus, the algorithm is unaware of the additional phase information. Another run was started at x_1 , hence, the phase information was made (indirectly) available to the algorithm. The results for 2D and 3D reconstruction of the caffeine molecule are shown in Figs. 5b and 5c respectively. The results of the natural image are shown in 5d.

It is evident from these results that our method significantly outperforms the HIO algorithm in all experiments. However, its superiority on the “Lena” image is tremendous.

In addition to the examples shown in this paper we have studied a number of other examples. Based on our observations we conclude that our algorithm demonstrates a significantly better convergence rate when the interval of phase uncertainty is not too close to π radians.

Besides the fast convergence rate, our method allows us to incorporate additional information about the image or the noise distribution in the measurements. Typical noise in photon/electron counting processes is distributed according to

**Fig. 5.** Reconstruction results

the Poisson distribution. In this case the maximum-likelihood criterion implies the functional for the error measure in the Fourier domain as given in Equation (10).

$$E_P(x) = \mathbf{1}^T (|Fx|^2 - r^2 \ln(|Fx|^2)). \quad (10)$$

To demonstrate the performance of our method we contaminated the measurements (r^2) of the “Lena” image with Poisson noise such that the signal to noise ratio (SNR) was 53.6 dB. The phase uncertainty was 3 radians as before. First, we started by solving the convex problem, as defined by Equation (4). The solution obtained was then used as the starting point for the second stage of our method using the non-convex functional defined in Equation (10). The HIO algorithm also started at this solution. In addition to using the objective function that fits the noise distribution we also included a regularization term in the object space. In this example, we used the total variation functional [12]

$$\text{TV}(x) = \int |\nabla x| \quad (11)$$

with a small weight. Total variation is a good prior for a broad range of images, especially, for those that are approximately piece-wise constant. In our case, introduction of this regularization added about 3 dB to the reconstruction SNR. The reconstruction results are shown in Fig. 6. Our method achieved the SNR of




Fig. 6. Reconstruction from noisy data

30dB, while the HIO algorithm produced a significantly inferior result. Its SNR was only 16.7dB. Note that the SNR in the reconstructed images is worse than the SNR in the data. This is a typical situation for ill-conditioned problems. It is also worth to note that Poisson noise of small intensity can be well approximated by Gaussian noise. However, if you use the objective function suitable for the Gaussian noise in r^2 , i.e.,

$$\| |Fx|^2 - r^2 \|^2 \quad (12)$$

the reconstruction results are few dB's worse than those we got with the proper choice of the objective function.

5 Discussion and Conclusions

We presented a convex optimization method for the phase retrieval problem when the phase is known to lie within a certain interval. Straight forward incorporation of this information does not lead to a successful reconstruction. Therefore, we designed a two-stages method. At the first stage we perform convex relaxation and solve the resulting convex problem. At the second stage the original objective function is introduced into the scheme and the reconstruction continues from the solution of the first stage. The algorithm demonstrates significantly better convergence rate compared to the current reconstruction methods. Moreover, in contrast to these methods, our technique is flexible enough to allow incorporation of additional information. Practical examples of such information include measurement noise distribution, knowledge of the reconstructing image being smooth or piece-wise constant.

There are several ways to obtain rough phase estimations. First, and probably most obvious, is to introduce into the scene an object whose Fourier transform is known. In this case the recorded data is the modulus of a sum of two complex numbers one of which is known. It is easy to show that in this case the phase of the unknown number can be reconstructed by solving two equations with

two unknowns. Since our method allows phase uncertainty to be as large as 3 radians, the Fourier transform of this additional object can be known only approximately. However, we believe that a rough phase estimation is possible without introducing a change into the physical setup. At this time we work on an algorithm that will perform some sort of bootstrapping for simultaneous phase retrieval and phase uncertainty estimation. For example, if we use one of the current reconstruction methods that are able to reconstruct the image. Then, at some point, the phase will be close enough the true one and if we manage to identify such a moment we can switch to our algorithm and get a faster reconstruction. As a last resort, the whole interval of 2π radians can be split into 2 or three intervals and each one can be tested separately. In this case, once a correct partitioning is found our method will recover the image.

References

1. Neutze, R., Wouts, R., van der Spoel, D., Weckert, E., Hajdu, J.: Potential for biomolecular imaging with femtosecond x-ray pulses. *Nature* 406(6797), 752–757 (2000)
2. Marchesini, S., He, H., Chapman, H.N., Hau-Riege, S.P., Noy, A., Howells, M.R., Weierstall, U., Spence, J.C.H.: X-ray image reconstruction from a diffraction pattern alone. *Physical Review B* 68(14) (October 2003); 140101 Copyright (C) 2009 The American Physical Society; Please report any problems to prola@aps.org
3. Hayes, M.: The reconstruction of a multidimensional sequence from the phase or magnitude of its fourier transform. *IEEE Transactions on Acoustics, Speech, and Signal Processing [IEEE Transactions on Signal Processing]* 30(2), 140–154 (1982)
4. Hove, P.V., Hayes, M., Lim, J., Oppenheim, A.: Signal reconstruction from signed fourier transform magnitude. *IEEE Transactions on Acoustics, Speech and Signal Processing* 31(5), 1286–1293 (1983)
5. Gerchberg, R.W., Saxton, W.O.: A practical algorithm for the determination of phase from image and diffraction plane pictures. *Optik* 35, 237–246 (1972)
6. Fienup, J.R.: Phase retrieval algorithms: a comparison. *Applied Optics* 21(15), 2758–2769 (1982)
7. Osherovich, E., Zibulevsky, M., Yavneh, I.: Signal reconstruction from the modulus of its fourier transform. Technical Report CS-2009-09, Technion (December 2008)
8. Nieto-Vesperinas, M.: A study of the performance of nonlinear least-square optimization methods in the problem of phase retrieval. *Journal of Modern Optics* 33, 713–722 (1986)
9. Narkiss, G., Zibulevsky, M.: Sequential subspace optimization method for Large-Scale unconstrained problems. CCIT 559, Technion, EE Department (2005)
10. Liu, D.C., Nocedal, J.: On the limited memory BFGS method for large scale optimization. *Mathematical Programming* 45(1), 503–528 (1989)
11. Miao, J., Hodgson, K.O., Sayre, D.: An approach to three-dimensional structures of biomolecules by using single-molecule diffraction images. *Proceedings of the National Academy of Sciences of the United States of America* 98(12), 6641–6645 (2001)
12. Rudin, L.I., Osher, S., Fatemi, E.: Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena* 60(1-4), 259–268 (1992)