

Focal-plane wavefront sensing with high-order adaptive optics systems

Visa Korkiakoski^{a,b}, Christoph U. Keller^b, Niek Doelman^{c,b}, Matthew Kenworthy^b, Gilles Otten^b and Michel Verhaegen^a

^aDelft Center for Systems and Control, Mekelweg 2, 2628CD Delft, The Netherlands

^bLeiden Observatory, Niels Bohrweg 2, 2333CA Leiden, The Netherlands;

^cTNO Science and Industry, Stieltjesweg 1, 2628CK Delft, The Netherlands;

ABSTRACT

We investigate methods to calibrate the non-common path aberrations at an adaptive optics system having a wavefront-correcting device working with an extremely high resolution (larger than 150×150 correcting elements). We use focal-plane images collected successively, the corresponding phase-diversity information and numerically efficient algorithms to calculate the required wavefront updates. Different approaches are considered in numerical simulations, and laboratory experiments are shown to confirm the results. We compare the performances of the standard Gerchberg-Saxton algorithm, Fast & Furious (use of small-phase assumption to take advantage of linearisation) and recently proposed phase-retrieval methods based on convex optimisation. The results indicate that the calibration task is easiest with algorithms similar to Fast & Furious, at least in the framework we considered.

1. INTRODUCTION

Future adaptive optics (AO) systems, for instance in the high-contrast imaging with the extremely large telescopes, will correct the wavefront in increasingly high resolution. The current high-density deformable mirrors (DM) have actuator arrays of up to 64×64 ¹ while the largest envisioned AO system at the E-ELT has a DM of 200×200 actuators.^{2,3}

To successfully use such high-order AO systems, it is necessary to accurately calibrate the non-common path aberrations caused by the differences in the optical paths between wavefront sensor and the focal-plane scientific camera.

The most popular method for the focal-plane wavefront sensing is perhaps the Gerchberg-Saxton (GS) error reduction algorithm^{4,5} and their variations, for instance.^{6,7} These are numerically very efficient algorithms, and it is easy to modify them for different applications. However, they suffer from accuracy, in particular because their iterative improvement procedure often stagnates at a local minimum. Various alternatives have been proposed, and a popular approach is to use general numerical optimization techniques to minimize an error function; examples include.⁸⁻¹⁰ However, when the number of optimization parameters is increased, the computational requirements generally rise unacceptably fast.

The numerical issues can be significantly reduced, if the unknown wavefront is sufficiently small. The assumption is valid, for example, when calibrating the non-common path aberrations. This has been exploited in numerous recent works, for instance.¹¹⁻¹⁵ It has also been shown that the Fast & Furious approach is suitable for effectively controlling wavefront correctors with 20 000–30 000 degrees of freedom.¹⁶

In this paper, we investigate different options to do focal-plane wavefront sensing for AO systems with an extremely high-resolution wavefront corrector – having dimensions larger than 150×150 correction elements. Section 2 discusses the options: FF-GS, GS and phase-retrieval based on convex optimisation. Sections 3 and 4 show the results and conclusions respectively.

Further author information: (Send correspondence to V.K.)
V.K.: E-mail: v.a.korkiakoski@tudelft.nl

2. ALGORITHMS

The algorithms we consider here can be divided into two groups. First, we consider algorithms that take advantage of the feedback loop when applying the correction: FF and FF-GS. Their performance will be compared to the conventional phase-retrieval setting where the unknown wavefront is retrieved from a set of images – but with no emphasis of iterative nature of the wavefront correction. The second group includes the traditional Gerchberg-Saxton algorithm and its Fienup variation^{4,5} and a newer class of phase retrieval based on convex optimization.

2.1 Fast & Furious

The Fast & Furious algorithm is based on iteratively applying a weak-phase approximation of the wavefront. The main principle of the weak-phase solution is presented in,¹¹ but we found slight modifications¹⁷ leading to significantly better performance. The algorithm uses focal-plane images and phase-diversity information to solve the wavefront, and the estimated wavefront is corrected with a wavefront correcting device. The correction step produces phase-diversity information and a new image that are again used to compute the following phase update. The schematic illustration of the algorithm is shown in Fig. 1.

The obvious limitation of the FF algorithm is the assumption of the pupil amplitudes being even. This holds reasonably well for most of the optical systems having a circular shape, possibly with a central obstruction. However, to achieve the optimal focal-plane wavefront sensing with a high-order system not suffering from other limiting factors, it is necessary to consider imaging models where the pupil amplitudes can have an arbitrary shape. We have approached the problem by combining the FF-style weak-phase solution and a version of the Gerchberg-Saxton (GS) algorithm. The new algorithm is referred to as FF-GS in the following.

A detailed presentation and analysis of FF and FF-GS is shown in.¹⁶

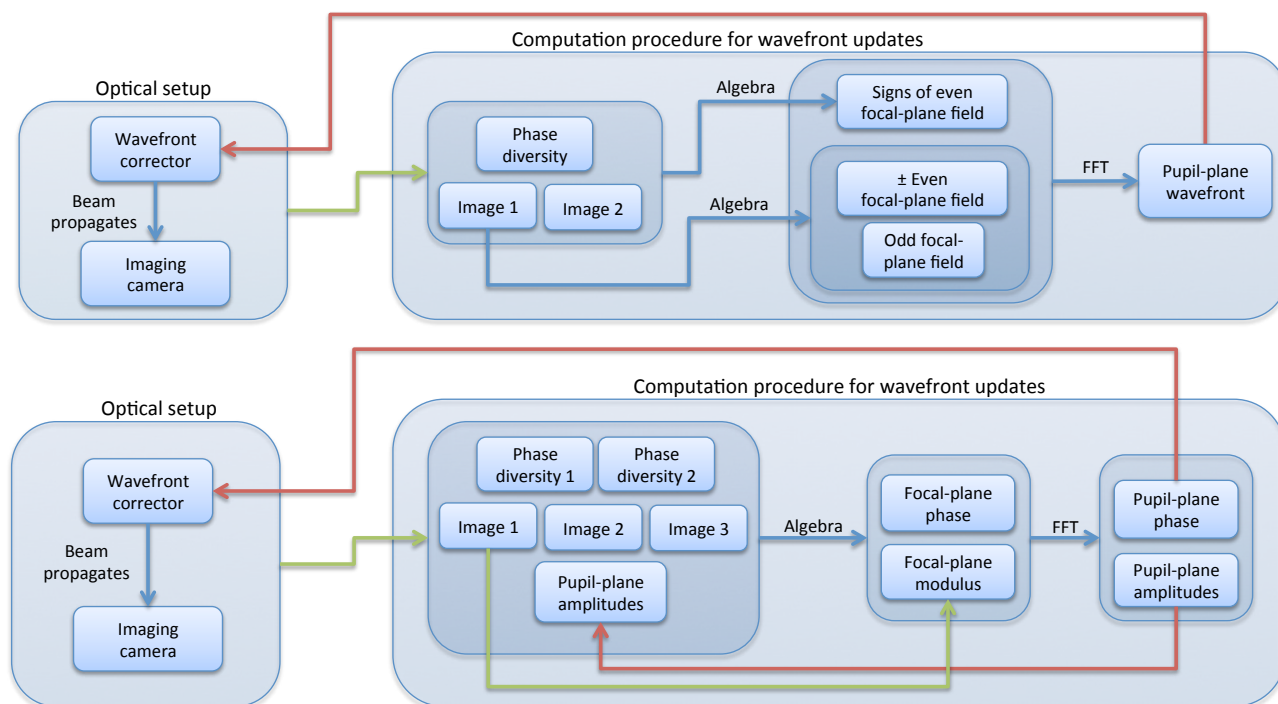


Figure 1. Schematic illustration of the FF and FF-GS algorithms.

2.2 Phase retrieval with convex optimization

A wave of new research¹⁸⁻²¹ has emerged since it was shown in 2011 that a lifting technique called PhaseLift¹⁸ can be used as a tool for phase retrieval. It uses tools from optimization theory to trade ("lift") a non-linear

optimization problem into a well-known convex problem. However, the amount of unknown parameters increases dramatically.

The original idea of PhaseLift is to formulate the standard phase-retrieval problem as rank-minimization problem and then approximate it by minimizing the trace. However, if there exist a unique solution, the trace minimization is not necessary, and it is possible to solve the lifted problem as a convex feasibility search. Such a feasibility problem can be solved using POCS-style algorithms (projection onto convex sets), and the convergence is faster than PhaseLift, if the optimization is done using standard gradient search.²¹

The current literature, however, hardly says anything about results on data sets of any realistic size. To apply the lifting techniques for sufficiently large sets, it is necessary to introduce additional approximations to reduce the computation costs to feasible levels. Here, we show how a practical implementation of POCS-style codes can be done for a lifted phase-retrieval case. The used notations are listed in Table 1.

Table 1. Notations	
\mathbf{x}	unknown complex field (phase and amplitude), a matrix of $N \times N$
\mathbb{b}_i	i th measurement array, a matrix of $M \times M$
x	stacked vector of \mathbf{x}
b_i	stacked vector of \mathbb{b}_i
A_i	matrix modelling the linear relationship (FFT) from x to b_i
\mathbb{X}	lifted unknown variables, a matrix of $N^2 \times N^2$
X	stacked vector of \mathbb{X}
Φ_i	i th phase diversity, a matrix of $N \times N$
ϕ_i	stacked vector of Φ_i
\mathcal{A}_i	matrix modelling the linear relationship from X to b_i

The traditional phase retrieval problem can be written as

$$b_i = |A_i x|^2 = |A \langle \phi_i, x \rangle|^2, \quad (1)$$

where b_i ($i = 1, \dots, M$) is the i th intensity measurement vector, x is the unknown field and A_i is the matrix that models the linear relationship between the phase and measurement domains. ϕ_i is the electric field corresponding to the known phase diversity – it is introduced between the measurements. In many imaging applications, the situation can be very accurately modeled by the Fraunhofer diffraction,

$$\mathbb{b}_i = |\mathcal{F} \{ \langle \Phi_i, \mathbf{x} \rangle \}|^2, \quad (2)$$

where $\mathcal{F} \{ \cdot \}$ is a two dimensional Fourier transform and $\langle \cdot, \cdot \rangle$ denotes an inner product. In this work, we focus on the Fraunhofer diffraction case, and therefore the matrices A_i are defined through the use of point-wise multiplication of the phase diversities and the application of FFTs.

PhaseLift formulates the nonlinear inverse problem as a convex optimization problem:

$$\min \text{Tr}(\mathbb{X}) \quad (3)$$

$$\alpha(\mathbb{X}) = b \quad (4)$$

$$\mathbb{X} \succeq 0, \quad (5)$$

where \mathbb{X} is a complex unknown, defined by lifting of x : $\mathbb{X} = x x^H$. $\alpha(\cdot)$ is a linear mapping from positive semidefinite matrices to a vector matching the measurement dimensions.

The PhaseLift problem is identical to the original phase retrieval problem, if a unique solution exist for x .²¹ The trace minimization was introduced to approximate the rank minimization, which makes intuitively sense. However, given the unique solution, it is not necessary. PhaseLift can be solved using general convex optimization tools. However, if the trace minimization is omitted, the problem is reduces to a feasibility search, which can be solved faster.

If the PhaseLift is simplified into a feasibility problem, it can be solved using general algorithms like Projection Onto Convex Sets (or its faster extensions). The general idea is to project the current estimate (\mathbb{X}) alternately to the sets that satisfies

$$b_i = \mathcal{A}_i X \quad (6)$$

and

$$\mathbb{X} \succeq 0. \quad (7)$$

To solve the problem with a practical dimension, some approximations are needed. For practical reasons, it is not possible to store the whole matrix of \mathbb{X} . Therefore, we represent it as

$$\mathbb{X} = \sum_{j=1}^K \lambda_j x_j x_j^*, \quad (8)$$

where λ_j is the j th eigenvalue and x_j is the j th eigenvector.

Thus, we maintain a factorized version of \mathbb{X} at all times (K largest eigenvalues are kept), and therefore the projection onto $\mathbb{X} \succeq 0$ becomes trivial: set the negative eigenvalues to zero. Also the recovery of the original unknown (x) is (in principle) easy: it is the eigenvector corresponding to the largest eigenvalue. The projection to satisfy the measurements could be done using the matrix equations as:

$$X_{k+1} = X_k - \mathcal{A}_i^* (\mathcal{A}_i^* \mathcal{A}_i)^{-1} (\mathcal{A}_i X_k - b_i), \quad (9)$$

where X_k is the k th estimate and i is the measurement index, $i = 1, \dots, K$. Due to the properties of the Fourier transform, this simplifies to

$$X_{k+1} = X_k - \mathcal{A}_i^* (\mathcal{A}_i X_k - b_i). \quad (10)$$

However, the matrices \mathcal{A}_i are too large to be explicitly modeled, and we would like to use the FFTs to make the computation faster. To implement this efficiently, the latter part can be calculated as series of FFTs (as many FFTs as there are non-zero eigenvalues in \mathbb{X}_k),

$$u_i = \mathcal{A}_i X_k - b_i = \text{vec} \left(\sum_{j=1}^K \lambda_j |\mathcal{F} \{ \langle \text{mat}(x_j), \Phi_i \rangle \} |^2 \right) - b_i, \quad (11)$$

where λ_j is the j th eigenvalue of \mathbb{X}_k and Φ_i is the field corresponding to the phase diversity of the i th image. The multiplication by \mathcal{A}_i^* is more difficult since it projects the vector of size N (number of measurements) to the dimension of N^2 . To find the approximation of this, we apply power methods on the quantity

$$Z = \mathbb{X}_{k+1} = \mathbb{X}_k - \text{mat}(\mathcal{A}_i^* u_i). \quad (12)$$

to find its K highest eigenvalues. This requires finding a fast way to implement a matrix operation

$$Zv = \mathbb{X}_k v - \text{mat}(\mathcal{A}_i^* u_i) v. \quad (13)$$

This can be achieved by two consecutive FFTs and simple algebra,

$$p_1 = \mathcal{F} \{ \langle \text{mat}(v), \Phi_i \rangle \}, \quad (14)$$

$$p_2 = \langle \mathcal{F}^{-1} \{ \langle p_1, \text{mat}(u_i) \rangle \}, 1/\Phi_i \rangle, \quad (15)$$

and

$$Zv = \mathbb{X}_k v - \text{vec}(p_2) \quad (16)$$

$$= \sum_{j=1}^K \lambda_j x_j (x_j^H v) - \text{vec}(p_2). \quad (17)$$

This linear operation (Zv) is then done sequentially as a part of power iteration to find the K highest eigenvalues of Z , which is the next iterate for \mathbb{X} .

In practice, as pointed out by many others, the projection algorithms get often significantly faster, if some memory is introduced in the process. We define $Y = 2\mathbb{X}_k - \mathbb{X}_{k-1}$, and projection is applied on Y instead of \mathbb{X}_k . We also applied this modification.

3. RESULTS

Our earlier publications^{16,22,23} show in detail – both simulations and laboratory experiments were carried out – how the FF and FF-GS work. Here, we illustrate, with the help of numerical simulations, how the alternative approaches would perform.

This section shows a comparison of the FF-GS algorithm with both the traditional GS algorithms and the convex optimization. As a test case, we use data from the setup we have used with our earlier publication.¹⁶ It has optical aberrations typical for high-contrast adaptive optics; a spatial light modulator is used to correct those aberrations with the help of a static focal-plane camera.

3.1 FF-GS compared with GS

We use numerical simulations that are shown to model the reality rather accurately.^{16,23} We model eight different sources of errors explicitly:

1. SLM quantification. We use only 6 bits to control the wavefront.
2. PSF sampling. The wavefront and the resulting PSF are sampled internally by a factor of two higher than what the hardware controls or observes. The control algorithms use re-binned PSFs, and the simulated wavefront correction is interpolated bilinearly from the reconstruction at a resolution of 170×170 .
3. Image noise and dynamic range. We estimate the read-out noise of the HDR images to be at a level of $2.2 \cdot 10^{-6}$ of the image maximum. Gaussian random noise is added to the simulated PSFs. The HDR images have maximum values $\sim 4 \cdot 10^8$, corresponding to about 29 bits, and this is also modeled in the simulations.
4. Background level. Standard background subtraction is performed on the PSF images, but a small error will still remain. Therefore, we add a constant background level, $2.7 \cdot 10^{-6}$ of the image maximum, to the simulated PSFs.
5. Non-perfect pupil. Instead of the perfect top-hat function, we use pupil amplitudes shown in Fig. 2. See also.¹⁶
6. Amplitude aberrations. We simulate the coupling of the wavefront and the transmission of the SLM as illustrated in.¹⁶
7. Alignment errors. Although the dOTF calibration is rather accurate, some error could still be present in the affine transform that we use to map the wavefront to the SLM pixels. The simulations indicate that if the transform has a mismatch corresponding to a rotation larger than 0.4° , FF and FF-GS would be unstable. In practice, with the used hardware, we saw no hints these of instabilities. Therefore, a rotation error of 0.4° represents the maximum misregistration that the wavefront control algorithms are likely to experience.
8. Tip-tilt error. Internal turbulence in the optical setup causes frame-to-frame wavefront variations, which can be approximated to a degree as small shifts of the recorded images. We measured the difference of the center-of-gravity between two consecutive PSFs recorded with the HDR method, and it was found to be on average 0.025 pixels. This error cannot be taken into account by the phase-diversity approach, and we model its impact on the performance.

At first, we show how the conventional GS algorithm detects the iteratively corrected aberrations. We let FF algorithm to run 3–5 initial iterations, and the resulting new images and phase-diversities are fed to the GS algorithm to retrieve the phase. The retrieved WF is then iteratively used to correct the aberrations. To optimize the GS performance, we consider the following parameters:

1. Loop gain in the iterative correction. A value of 0.3 of optimal for FF-GS, and a value of 0.9 or higher is optimal for GS.

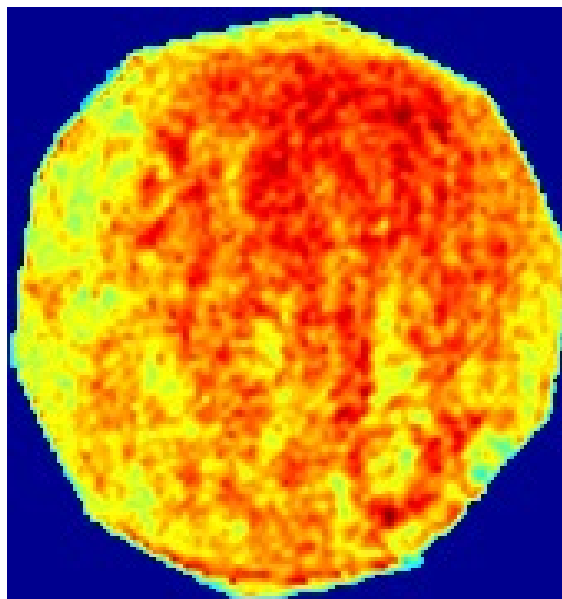


Figure 2. Pupil amplitudes (170×170 pixels) used in the simulations.

2. Number of used images. We test the algorithm with 3 and 5 images.
3. Number of GS iterations (alternate projections to all the images in the measurement set). A rough estimate of the WF is obtained with only 10 GS iterations. We also test a more accurate reconstruction with 50 iterations.
4. Restriction of the pupil amplitudes. We test the algorithm both in its traditional form (enforce even pupil amplitudes over the aperture) and in an unrestricted more where no pupil-plane amplitude restrictions are enforced.

Fig. 3 shows the results. The upper row has the case where simulations are down with no error sources. The lower row has the case with all the error sources modeled.

It appears to be very difficult to make the GS algorithm stable. The best results are obtained with 5 images and 50 GS iterations (possibly further images and projection iterations would still improve the results). A lower amount used images or GS iterations works to some extent without error sources if the pupil amplitudes are restricted. If pupil amplitudes are not restricted (thus also an estimate for the pupil amplitudes is obtained), five images are needed also in the error-free case – but no significant gain is obtained by using 50 instead of 10 GS iterations. The more practical simulations show successful results only with the higher number of GS iterations and five used images.

In total about ten wavefront correcting iterations are needed to reach close to the optimum. Without error sources, ~ 15 iterations produce the best results, with the error sources 20–30 iterations are needed. At its optimum, the GS algorithm with five used diversity images has $\sim 20\%$ higher residual WF error compared to FF-GS, but in the error-included simulations the WF error is $\sim 20\%$ lower. However, the comparison to FF-GS is unfair in the sense that it uses only three images as implemented in.¹⁶

After the optimum, the GS algorithm becomes unstable, and the wavefront error starts to creep upwards. This happens since the phase-diversity becomes too small to give useful extra information. In addition, unlike in the FF-GS, the phase-diversity information is not enforced when applying the image-plane restrictions.

The creep appears as an amplification of certain higher spatial frequencies. Those frequencies are visible as single speckles, and each iterations makes them brighter. We illustrate this in Fig. 4. The top-left image shows the ideal image without any wavefront aberrations (the visible aberrations are caused by non-uniform

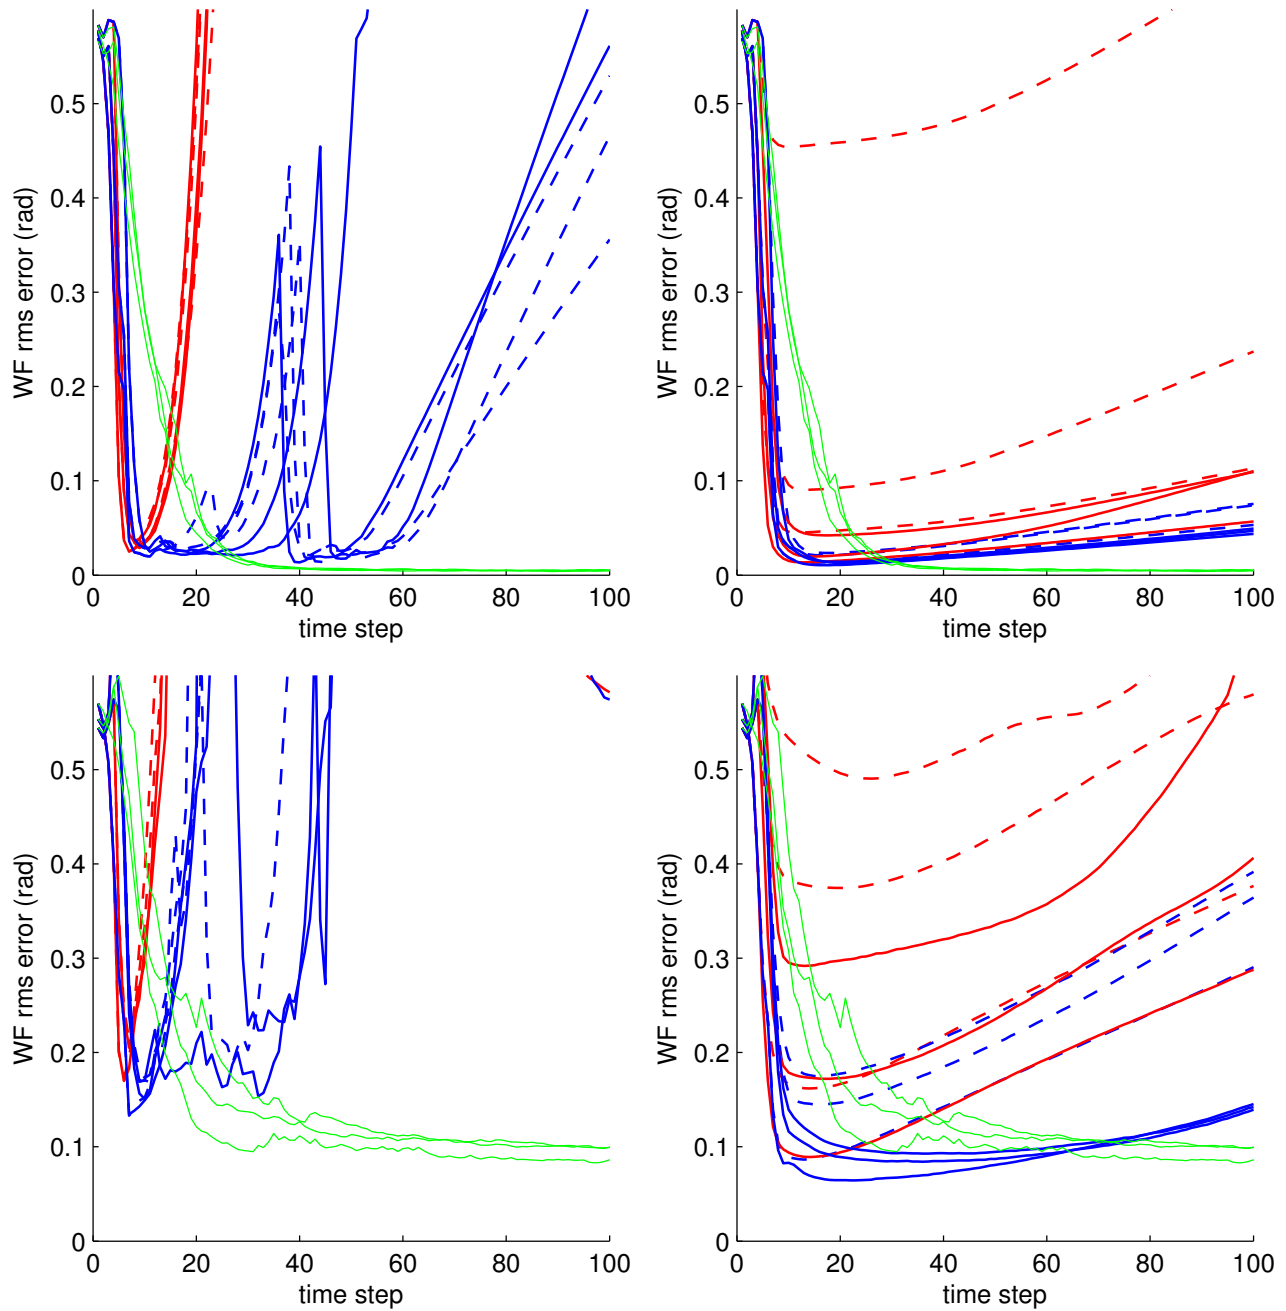


Figure 3. Comparison of GS and FF-GS algorithms: WF rms error as a function of time step. Upper row: error free case. Lower row: all errors simulated. Left column: GS restrict pupil amplitudes. Right column: GS lets pupil amplitudes have arbitrary values. Green lines: FF-GS for comparison, data from.¹⁶ Blue lines: GS uses 5 images. Red lines: GS uses 3 images. Dotted lines: 10 GS iterations. Solid lines: 50 GS iterations.

transmission). The optimal GS algorithm (5 images, 50 GS iterations, 25 correction iterations) has corrected well all the lower spatial frequencies, but a few uncorrected speckles remain further on the halo. After total 100 correction iterations those uncorrected speckles have become brighter and many more have emerged.

A visual inspection also reveals that most of the problematic speckles appear also in the initial uncorrected image. Thus, the GS algorithm fails to be reliable at the high spatial frequencies. The last column shows results of the FF-GS algorithm for comparison. FF-GS has no problematic speckles that would grow when repeating the correction iterations, which makes it significantly more robust compared of the GS approach. The overall WF error, however, is $\sim 20\%$ higher compared to the ideal GS performance – but this could likely be improved if the algorithm was updated to use as many images as the GS method is using. FF-GS used three images while the best GS realization here used five.

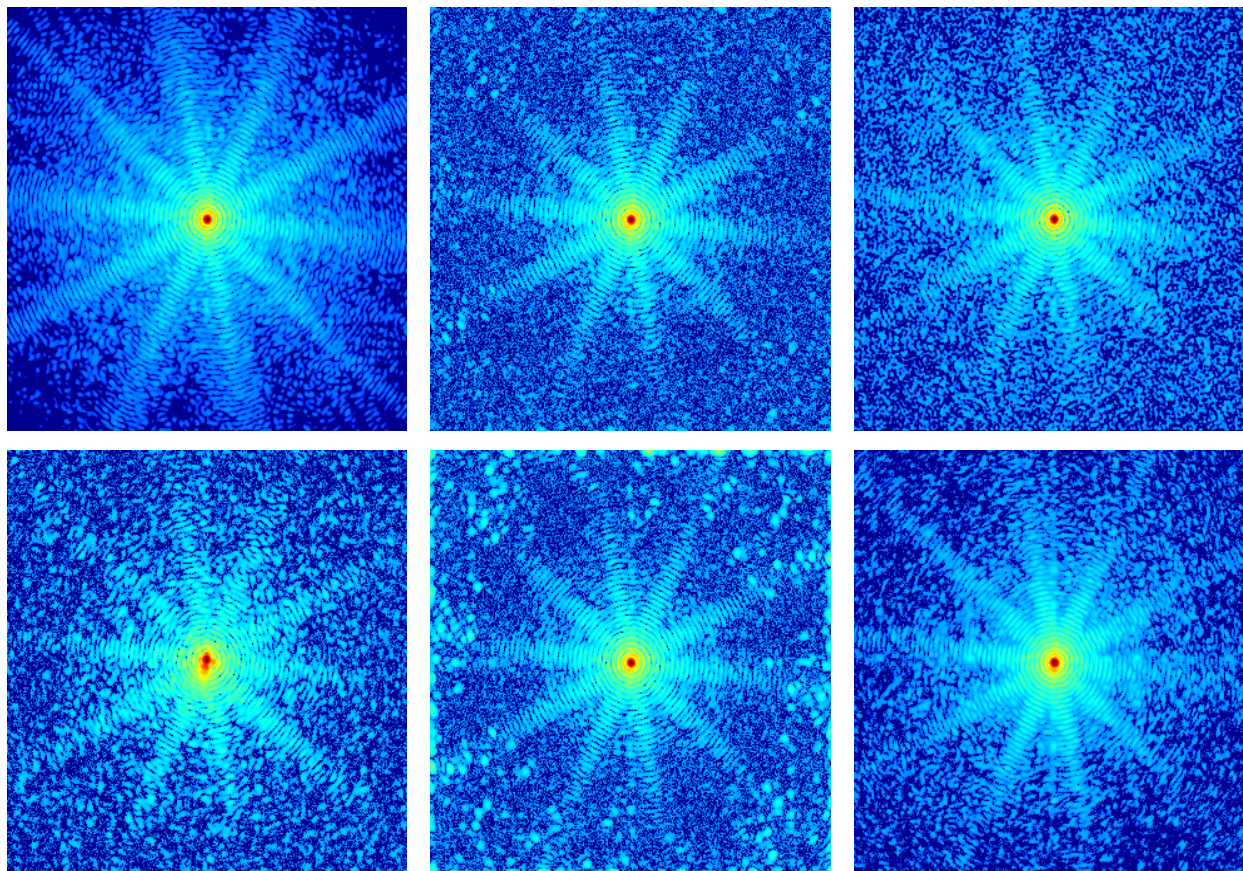


Figure 4. Examples of the focal-plane images raised to 0.1 power. Top row: perfect image (no wavefront aberrations), GS algorithm with optimal parameters after 25 iterations, FF-GS algorithm converged (simulated). Bottom row: initial image, GS algorithm with optimal parameters after 100 iterations, FF-GS algorithm converged (laboratory measurement¹⁶).

3.2 Results with convex optimization

Here, we test the new idea of convex optimization (see Section 2.2) applied in phase retrieval. The test case is the same as in the previous section, but we show the results only for the error-free case. We use five images and corresponding diversity information.

However, even with the fast implementation we wrote, the implementation is still too low to effectively use it as an iterative corrector. Therefore, we do only a single phase retrieval and compare the results to the expected values. The comparison to the previous sections is therefore not entirely fair since there we used several correction iterations to improve the performance.

Fig. 5 shows the original wavefront and the final reconstruction of the POCS algorithm. The figure also shows radial cuts of the wavefronts. The convergence properties are shown in Fig. 6. We see that 1000 iterations are not enough for a sufficient convergence since the convergence becomes extremely slow after 10–100 first iterations. Each iteration requires about 200 FFTs with a dimension of 640×640 . Total calculation time for the final reconstruction took several hours with a modern PC and standard Matlab FFT implementation.

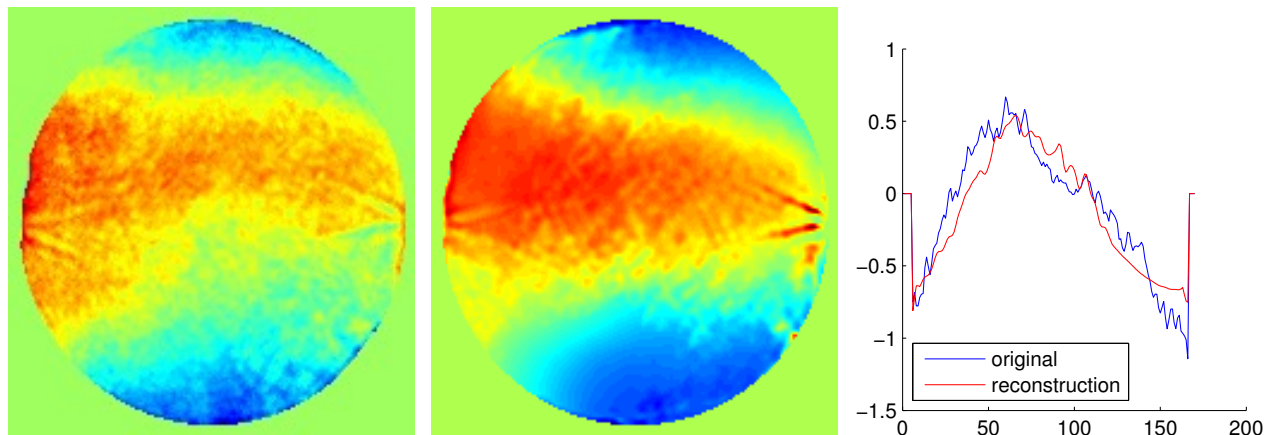


Figure 5. Results for phase retrieval with convex optimization. Left: the unknown WF to be reconstructed. Middle: the reconstruction. Right: radial cuts of the images (units are in radians).

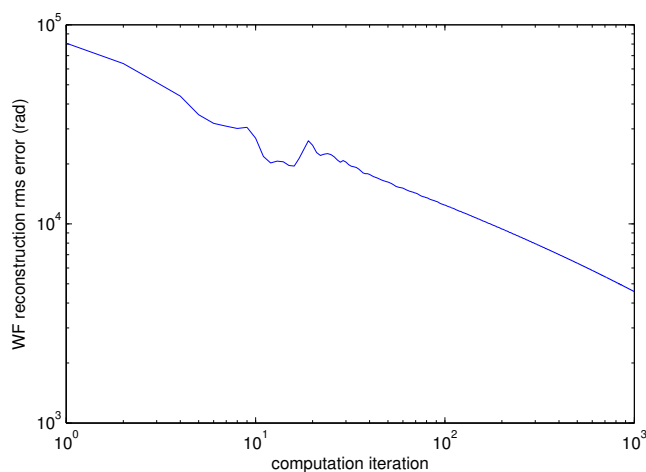


Figure 6. Convergence of the POCS algorithm for phase retrieval.

We note that the POCS implementation is well able to retrieve the low spatial frequencies of the wavefront, but it fails to correctly find the sharp high-order features. In theory, this might still happen if the algorithm is let to converge – but this option is still limited by the computational requirements. It is also possible that iterative correction, in the same way as with the FF-GS and GS algorithms, would help to correctly retrieve the unknown wavefront. However, the computational problems make also this approach difficult.

We also show here, in Fig. 7, the PSF image of the unknown wavefront and the PSF resulting from the estimated wavefront. The image also suggests that the high spatial frequencies are not correctly estimated, and the dataset should have enough information to yield a better wavefront estimate.

4. CONCLUSIONS AND DISCUSSION

We have studied different options to use focal-plane wavefront sensing to determine the unknown pupil-plane field at a very high spatial resolution (more than 150×150 pixels). With the help of numerical simulations,

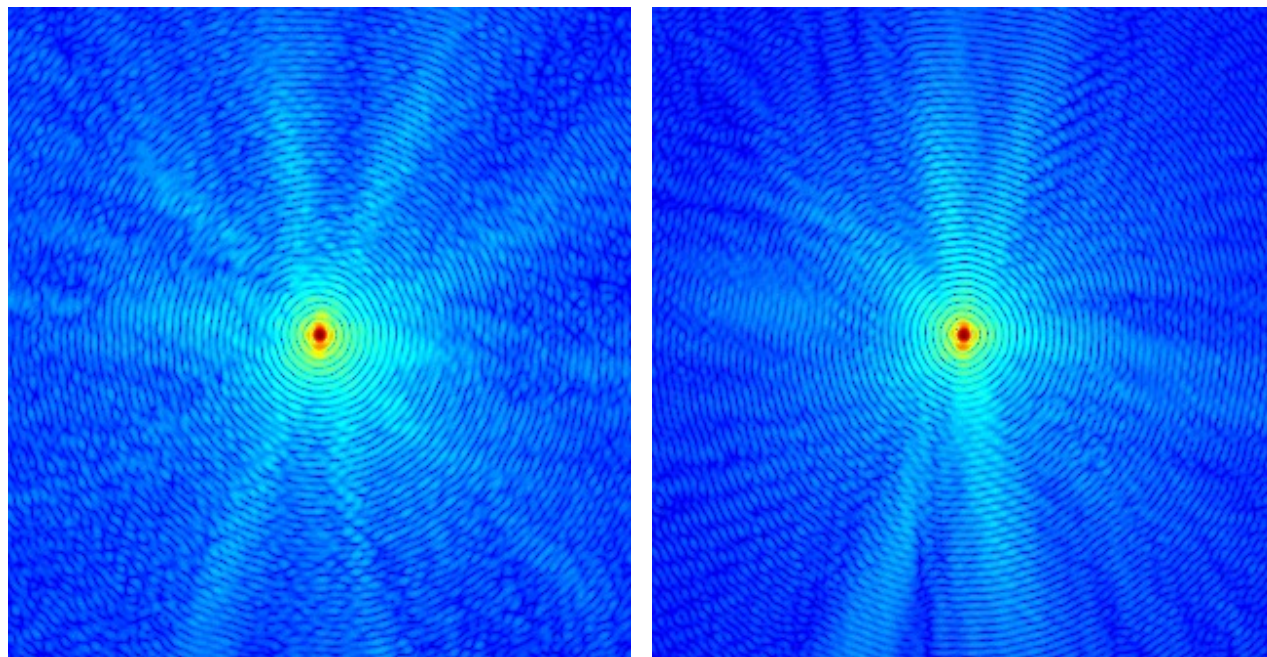


Figure 7. Illustration the performance of our POCS based phase retrieval. Left: recorded image (no diversity). Right: the image resulting from retrieved wavefront.

we tested the performance of traditional Gerchberg-Saxton, recently proposed Fast & Furious algorithm and a convex optimisation based phase retrieval.

Our results indicate that the methods based on fast iterative correction and linear approximations tend be easier to implement in a reliable way. The FF and FF-GS algorithms have been shown to an excellent choice for this type of phase-retrieval.¹⁶

The GS method works to some extent, but, in particular at high spatial frequencies, it is difficult to find reliable wavefront estimates. In addition, its computational requirements are ~ 100 times higher compared to the FF-GS.

The phase retrieval based on convex optimisation also works to some extent, but our numerical implementation is not fast enough to make it a competitive alternative to the traditional methods.

REFERENCES

- [1] L. A. Poyneer, B. Bauman, S. Cornelissen, J. Isaacs, S. Jones, B. A. Macintosh, and D. W. Palmer, "The use of a high-order mems deformable mirror in the gemini planet imager," in *SPIE MOEMS-MEMS, Proc. SPIE* **7931**, pp. 793104–793104, 2011.
- [2] C. Vérinaud, M. Kasper, J.-L. Beuzit, R. G. Gratton, D. Mesa, E. Aller-Carpentier, E. Fedrigo, L. Abe, P. Baudoz, A. Boccaletti, M. Bonavita, K. Dohlen, N. Hubin, F. Kerber, V. Korkiakoski, J. Antichi, P. Martinez, P. Rabou, R. Roelfsema, H. M. Schmid, N. Thatte, G. Salter, M. Tecza, L. Venema, H. Hanenburg, R. Jager, N. Yaitskova, O. Preis, M. Orecchia, and E. Stadler, "System study of EPICS: the exoplanets imager for the E-ELT," in *Adaptive Optics Systems II, Proc. SPIE* **7736**, pp. 77361N–77361N–12, July 2010.
- [3] V. Korkiakoski and C. Verinaud, "Simulations of the extreme adaptive optics system for EPICS," *Proc. SPIE* **7736**, pp. 773643–773643–14, July 2010.
- [4] R. W. Gerchberg and W. O. Saxton, "A practical algorithm for the determination of the phase from image and diffraction plane pictures," *Optik* **35**, pp. 237–246, 1972.
- [5] J. R. Fienup, "Phase retrieval algorithms: a comparison," *Applied Optics* **21**, pp. 2758–2769, Aug. 1982.

- [6] J. J. Green, D. C. Redding, S. B. Shaklan, and S. A. Basinger, "Extreme wave front sensing accuracy for the Eclipse coronagraphic space telescope," in *High-Contrast Imaging for Exo-Planet Detection*, A. B. Schultz, ed., *Proc. SPIE* **4860**, pp. 266–276, Feb. 2003.
- [7] R. S. Burruss, E. Serabyn, D. P. Mawet, J. E. Roberts, J. P. Hickey, K. Rykoski, S. Bikkannavar, and J. R. Crepp, "Demonstration of on sky contrast improvement using the modified Gerchberg-Saxton algorithm at the Palomar Observatory," *Proc. SPIE* **7736**, July 2010.
- [8] J. Sauvage, T. Fusco, G. Rousset, and C. Petit, "Calibration and precompensation of noncommon path aberrations for extreme adaptive optics," *Journal of the Optical Society of America A* **24**, pp. 2334–2346, Aug. 2007.
- [9] P. Riaud, D. Mawet, and A. Magette, "Nijboer-Zernike phase retrieval for high contrast imaging. Principle, on-sky demonstration with NACO, and perspectives in vector vortex coronagraphy," *Astron. Astrophys.* **545**, p. A150, Sept. 2012.
- [10] B. Paul, L. M. Mugnier, J.-F. Sauvage, M. Ferrari, and K. Dohlen, "High-order myopic coronagraphic phase diversity (COFFEE) for wave-front control in high-contrast imaging systems," *Optics Express* **21**, p. 31751, Dec. 2013.
- [11] R. A. Gonsalves, "Small-phase solution to the phase-retrieval problem," *Optics Letters* **26**, pp. 684–685, May 2001.
- [12] A. Give'On, R. Belikov, S. Shaklan, and J. Kasdin, "Closed loop, DM diversity-based, wavefront correction algorithm for high contrast imaging systems," *Opt. Express* **15**, p. 12338, 2007.
- [13] S. Meimon, T. Fusco, and L. M. Mugnier, "Lift: a focal-plane wavefront sensor for real-time low-order sensing on faint sources," *Opt. Lett.* **35**, pp. 3036–3038, Sep 2010.
- [14] F. Martinache, "The Asymmetric Pupil Fourier Wavefront Sensor," *Publ. Astron. Soc. Pac.* **125**, pp. 422–430, Apr. 2013.
- [15] C. S. Smith, R. Mariničá, A. J. den Dekker, M. Verhaegen, V. Korhikoski, C. U. Keller, and N. Doelman, "Iterative linear focal-plane wavefront correction," *Journal of the Optical Society of America A* **30**, p. 2002, Oct. 2013.
- [16] V. Korhikoski, C. U. Keller, N. Doelman, M. Kenworthy, G. Otten, and M. Verhaegen, "Fast & Furious focal-plane wavefront sensing," *Applied Optics* **accepted**, 2014.
- [17] C. U. Keller, V. Korhikoski, N. Doelman, R. Fraanje, R. Andrei, and M. Verhaegen, "Extremely fast focal-plane wavefront sensing for extreme adaptive optics," *Proc. SPIE* **8447**, pp. 844721–844721, 2012.
- [18] E. J. Candes, Y. C. Eldar, T. Strohmer, and V. Voroninski, "Phase retrieval via matrix completion," *SIAM Journal on Imaging Sciences* **6**(1), pp. 199–225, 2013.
- [19] F. Fogel, I. Waldspurger, and A. d'Aspremont, "Phase retrieval for imaging problems," *arXiv preprint arXiv:1304.7735*, 2013.
- [20] I. Waldspurger, A. d'Aspremont, and S. Mallat, "Phase recovery, maxcut and complex semidefinite programming," *Mathematical Programming*, pp. 1–35, 2012.
- [21] L. Demanet and P. Hand, "Stable optimizationless recovery from phaseless linear measurements," *Journal of Fourier Analysis and Applications*, pp. 1–23, 2012.
- [22] V. Korhikoski, C. U. Keller, N. Doelman, R. Fraanje, R. Andrei, and M. Verhaegen, "Experimental validation of optimization concepts for focal-plane image processing with adaptive optics," *Proc. SPIE* **8447**, pp. 84475Z–84475Z, 2012.
- [23] V. Korhikoski, N. Doelman, J. Codona, M. Kenworthy, G. Otten, and C. U. Keller, "Calibrating a high-resolution wavefront corrector with a static focal-plane camera," *Applied Optics* **52**, p. 7554, Nov. 2013.