Fast Tomographic Reconstruction of Atmospheric Turbulence from Micro-Lens Imagery

James G. Nagy

Mathematics and Computer Science
Emory University
Atlanta, GA 30322, USA
nagy@mathcs.emory.edu

Michael Hart

Center for Astronomical Adaptive Optics
Steward Observatory
University of Arizona
Tucson, AZ 85721, USA
mhart@as.arizona.edu

Stuart Jefferies

Institue for Astronomy
University of Hawaii
Pukalani, HI 96768, USA
stuartj@ifa.hawaii.edu

Douglas Hope

Department of Physics
United States Air Force Academy
Colorado Springs, CO 80840, USA
douglas.hope@usafa.edu

ABSTRACT

We consider using data acquired from a micro-lens array through which multiple images of the full field-of-view of an astronomical target are formed to attempt to reconstruct the 3-D wave front for the observations. This opens the door for both a beacon-less wave front sensor and imaging of fields-of-view substantially larger than the isoplanatic angle. The reconstruction problem can be modeled as a large-scale linear inverse problem, but standard algorithms used for 3-D computed tomography (CT) reconstruction cannot be applied because measured data are only taken from limited angular range, leaving entire regions of the frequency space un-sampled. However, we show that there is substantial structure in the mathematical model that can be exploited to obtain a robust algorithm that is amenable to efficient implementations.

1. INTRODUCTION

In many imaging situations, such as when ground based telescopes are used to observe objects in space, the observed image is degraded by blurring and noise. Although the blurring can be partially removed through sophisticated (and expensive) imaging devices, such as adaptive optics telescopes, computational postprocessing techniques are also often needed to further improve the resolution of the image. This

computational postprocessing, which is referred to as deblurring, restoration, or deconvolution [1, 5, 10, 8, 11, 18], requires solving an ill-posed inverse problem

$$g(x,y) = \int_{\mathcal{R}^2} k(x-\xi, y-\eta) f(\xi,\eta) d\xi d\eta + e(x,y)\,, \tag{1}$$

where f is the true object, g is the observed image, and e is additive noise. The kernel function k models the blurring operation, and is called the *point spread function* (PSF). Using a standard Fourier optics model for atmospheric turbulence [15], the PSF can be expressed in terms of the wavefront phase error, ϕ, of the light that reaches the telescope mirror,

$$k(x,y) = \left| \mathcal{F}^{-1} \left\{ \mathcal{P}(x,y) e^{\iota\phi(x,y)} \right\} \right|^2 , \tag{2}$$

where $\iota = \sqrt{-1}$, $\mathcal{P}$ is a characteristic function that models the shape of the telescope aperture (e.g., a circle or annulus), and $\mathcal{F}^{-1}$ is the 2-dimensional inverse Fourier transform. In perfect seeing conditions, $\phi = 0$, and in good seeing conditions (low turbulence), ϕ is a smooth function.

The effectiveness of image restoration algorithms depends strongly on how well one can approximate the PSF, or, equivalently, how well one can estimate the phase, ϕ. In some cases, a *wavefront sensor* (WFS) can be used to measure gradients of the wavefront, from which the phase can be reconstructed. A WFS is standard technology in adaptive optics systems, and many papers have been written about efficiently reconstructing the wavefront from the gradient measurements; see, for example [3, 9]. The phase reconstruction problem, from measured gradients, depends on the sensor geometry [6, 9] but it essentially amounts to solving a discrete finite difference problem [7, 14, 4]. Probably the most commonly used WFS is the Shack-Hartmann, found in optical systems across a very broad range of applications in astronomy, medical imaging, space situational awareness, directed energy weapons, and secure communications [21, 22, 23]. But the Shack-Hartmann requires a well-defined and preferably unresolved source on the far side of the turbulence to act as a beacon. In space imaging applications, this is either a natural guide star or a laser guide star.

The wavefront phase gradient information measured by the WFS is an integration of the turbulence effects along a line of sight through the 3-dimensional atmosphere. If imaging over a narrow field of view (FOV), then it is generally not necessary to model the 3D turbulence profile; a 2D integrated phase provides sufficient information to construct a spatially invariant PSF. However, when imaging over a wide FOV, modeling the 3D atmospheric structure is important because a spatially invariant assumption on the PSF may not be appropriate, especially in extreme turbulence situations.

Atmospheric tomography techniques have been developed to reconstruct a pseudo-3D representation of atmospheric turbulence. We use the phrase *pseudo-3D* because it is not feasible to obtain a full 3D reconstruction of the atmosphere. Instead, the aim of atmospheric tomography is to reconstruct information at a discrete set of dominant layers (slices) of atmospheric turbulence. In order to do this, it is necessary to use multiple measurements, where, loosely speaking, each measurement is obtained by viewing the turbulence layers from a different perspective. The typical, and perhaps most obvious, approach to do this, is by collecting WFS measurements of guide stars distributed throughout the FOV, and use standard limited angle tomography techniques to reconstruction the turbulent layers [2, 12, 13, 16, 17, 19, 20]. This is the approach taken in the Gemini multi-conjugate adaptive optics system [20], which employs five sodium laser guide stars to correct a field of view of about an arcminute.

Because it may be impractical or undesirable to obtain multiple observations of guide stars, we take an alternative approach. Specifically, we consider using WFS data acquired from a micro-lens array in which the full FOV of an astronomical target is imaged through each micro-lens. Gradients of the local wavefront phase aberration over each subaperture are used to reconstruct phase gradients at a discrete set of turbulence layers. With known phase gradients, it is then straightforward to reconstruct phases at each layer.

We show that the problem of reconstructing phase gradients at each layer, given phase gradients in each subaperture, can be modeled as a large-scale linear inverse problem. The problem is severely ill-conditioned, and can have distinct null space components. However, in certain situations, and with appropriate regularization, we show that phases can be reconstructed at multiple layers.

The rest of this paper is organized as follows. In Section 2 we describe the problem setup, discuss the properties of the mathematical model, and our approach to solve the resulting inverse problem. In Section 3 we show the results of some experiments on simulated data to illustrate the effectiveness of our approach. Some concluding remarks are given in Section 4.

2. PROBLEM SETUP

We consider a situation where the WFS divides the telescope pupil into a set of discrete lenses; see Figure 1 for an illustration of 5 lenses (one should really think of this as a discrete grid of 5×5 lenses, but the one-dimensional grid, or side view of a two-dimensional grid, shown in Figure 1, simplifies the notation). Each subaperture captures light reflected off of objects in the FOV, but from slightly different perspectives. In particular, as light travels through the atmosphere, it is affected by the dominant layers of turbulence in a cone region defined by the diameter of the FOV and the diameter of the subaperture. These cones are illustrated in Figure 1 with different colored lines; for example, red lines depict the cone region corresponding to the first (far left) subaperture.

Equation (2) can be used to model local blurring in the images captured by each subaperture, but because of the slightly different perspectives of the subapertures, each local blurring is defined by a different phase. Specifically, the phase corresponding to the i-th subaperture is defined by integrating through the atmospheric layers within the cone region from the i-th subaperture to the FOV. Note that the WFS actually measures gradients, and not phases directly, but since differentiation is a linear operation, we can assume the WFS measured gradients are an integration of gradients in the same cone regions.

To describe a mathematical model that relates the measured phase gradients to phase gradients at the various layers of turbulence, let $\bar{\phi}_x(i)$ denote measured phase gradients in the x, or horizontal direction, and $\bar{\phi}_y(i)$ denote measured phase gradients in the y, or vertical direction, in the i-th subaperture. We will assume an $n_s \times n_s$ square grid of subapertures, and define $N_s = n_s^2$, and thus $i = 1, 2, \ldots, N_s$.

Suppose further that we partition the unknown phase gradients at the various layers of turbulence into an $n_p \times n_p$ grid, and denote these partitioned regions of phase gradients as $\phi_x(j, \ell)$ and $\phi_y(j, \ell)$, where $j = 1, 2, \ldots, N_p = n_p^2$, and $\ell = 1, 2, \ldots N_L$, with N_L defining the number of layers of turbulence that we want to reconstruct. Figure 2 illustrates our notation, which again shows a side view of a two-dimensional grid with $N_p = 4$, $N_L = 3$, and $N_s = 25$.

Using this notation, a mathematical model relating the measured phase gradients to the unknown phase gradients at each atmospheric layer can be given by

$$\begin{aligned} \bar{\phi}_x(i) &= \sum_{\ell=1}^{N_L} \sum_{j=1}^{N_p} w(i, j, \ell) \phi_x(j, \ell) + \varepsilon_x(i) \\ \bar{\phi}_y(i) &= \sum_{\ell=1}^{N_L} \sum_{j=1}^{N_p} w(i, j, \ell) \phi_y(j, \ell) + \varepsilon_y(i)\,, \end{aligned} \tag{3}$$

where $w(i, j, \ell)$ is a weight associated with the contribution of $\phi(j, \ell)$ to $\bar{\phi}(i)$, and $\varepsilon_x(i)$ and $\varepsilon_y(i)$ represent unknown noise and other measurement errors. The weight $w(i, j, \ell)$ is the fraction of the region in which $\phi(j, \ell)$ is defined, which is contained within the i-th cone. For example, the situation depicted in Figure 2 shows that $w(2, 2, 1) = 0$. In general, it is expected that for a given subaperture, i, most of the weights $w(i, j, \ell)$ will be zero.

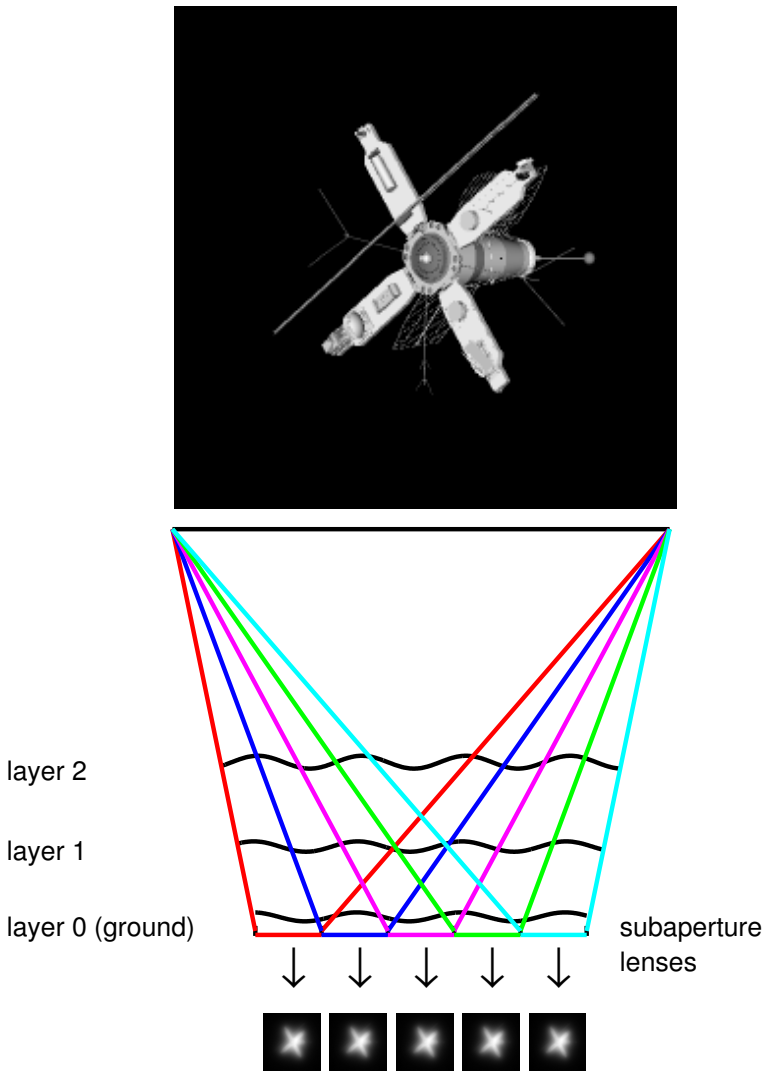

Figure 1: This figure illustrates a situation in which there are three dominant atmospheric layers of turbulence (one at the ground, and two at higher altitudes), and 5 lenslets in the aperture. In reality, there should be a grid of lenslets (e.g., 5×5), but to simplify the illustration, we show only a one-dimensional side view of the aperture.

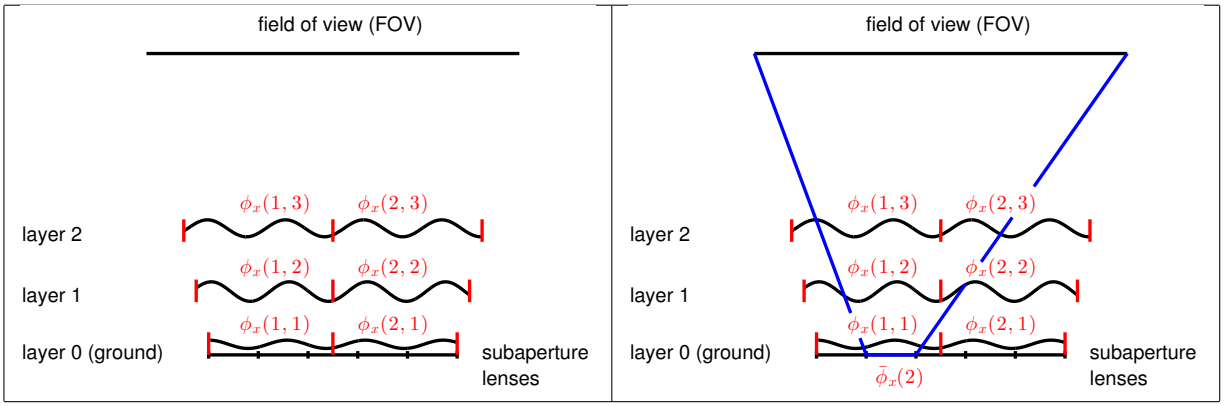

Figure 2: This figure illustrates discretization of the phase gradients.

Note that the equations in (3) can be written in matrix-vector form,

$$\bar{\boldsymbol{\phi}}_x = \boldsymbol{W}\boldsymbol{\phi}_x + \boldsymbol{\varepsilon}_x \quad \text{and} \quad \bar{\boldsymbol{\phi}}_y = \boldsymbol{W}\boldsymbol{\phi}_y + \boldsymbol{\varepsilon}_y, \tag{4}$$

where $\boldsymbol{W}$ is an $N_s \times N_p N_L$ matrix whose coefficients are given by $w(i,j,\ell)$. It is important to observe that to avoid an explicitly under-determined system, we should require that the partitioning of the phase gradients on each layer satisfy $N_p \leq N_s/N_L$. Even with such a restriction, the matrix W can be highly ill-conditioned, with a possibly nontrivial null space; this is illustrated in Section 3. Thus, solutions can be very sensitive to noise and other data errors, and it is necessary to incorporate regularization in the reconstruction methods. In this paper, we use standard Tikhonov regularization; that is, we solve the least squares problems:

$$\min_{\boldsymbol{\phi}_x} \left\| \begin{bmatrix} \boldsymbol{W} \\ \alpha_x \boldsymbol{I} \end{bmatrix} \boldsymbol{\phi}_x - \begin{bmatrix} \bar{\boldsymbol{\phi}}_x \\ \boldsymbol{0} \end{bmatrix} \right\|_2 \quad \text{and} \quad \min_{\boldsymbol{\phi}_y} \left\| \begin{bmatrix} \boldsymbol{W} \\ \alpha_y \boldsymbol{I} \end{bmatrix} \boldsymbol{\phi}_y - \begin{bmatrix} \bar{\boldsymbol{\phi}}_y \\ \boldsymbol{0} \end{bmatrix} \right\|_2 . \tag{5}$$

We remark that although the two problems in equation (5) have the same coefficient matrix, the fact that they have different data vectors means that the "optimal" level of regularization (defined by scalars α_x and α_y) might be different. However, currently we are using the same level of regularization for both problems, and we leave for future work a more complete analysis of the two problems, as well as the design of optimal regularization methods.

3. NUMERICAL EXPERIMENTS

In this section we provide experimental results on simulated data to test the potential of our proposed approach to atmospheric tomography. In each simulation, we assume that the FOV is at 330 km (the altitude of, for example, the International Space Station), coverage area of 40 arcsecs, two dominant layers of turbulence ($N_L = 2$), and an aperture diameter of 3.6 m.

We use the forward models given in equation (4), where we first use $\boldsymbol{\varepsilon}_x = \boldsymbol{\varepsilon}_y = 0$, and then take these terms to be 1% Gaussian white noise; that is, $\boldsymbol{\varepsilon}_x$ and $\boldsymbol{\varepsilon}_y$ are vectors with normally distributed random entries, mean 0, standard deviation 1, and scaled to that

$$\frac{\|\boldsymbol{\varepsilon}_x\|_2}{\|\boldsymbol{W}\boldsymbol{\phi}_x\|_2} = \frac{\|\boldsymbol{\varepsilon}_y\|_2}{\|\boldsymbol{W}\boldsymbol{\phi}_y\|_2} = 0.01\,.$$

Note that even without additional noise, the data contains downsampling approximation errors, which can have a deleterious effect on the computed reconstructions.

We consider two possible configurations:

$$n_s = 32, \quad n_p = 16 \quad \text{and} \quad n_s = 64, \quad n_p = 32\,.$$

The first configuration, where we assume the aperture is partitioned into 32×32 subapertures, reflects current capabilities of a typical WFS.

Figure 3 shows the singular values of the reconstruction matrix, $\boldsymbol{W}$, for both configurations. From these plots we see that the matrix has a one-dimensional (numerical) null space, and except for the one tiny singular value, the remaining singular values are relatively large. Therefore, we expect that a reasonable reconstruction of the layers should be possible, provided we incorporate a modest amount of regularization, and provided that the null space component of the solution does not contribute significant information to the reconstructed phases.

Reconstructions for the configuration with $n_s = 32$ and $n_p = 16$, using both noise free and noisy data, are shown in Figure 4 . Reconstructions for the configuration with $n_s = 64$ and $n_p = 32$, using both noise free and noisy data, are shown in Figure 5 . For both configurations we used regularization parameters $\alpha_x = \alpha_y = 10^{-5}$ for the noise free data, and $\alpha_x = \alpha_y = 10^{-2}$ for the noisy data.

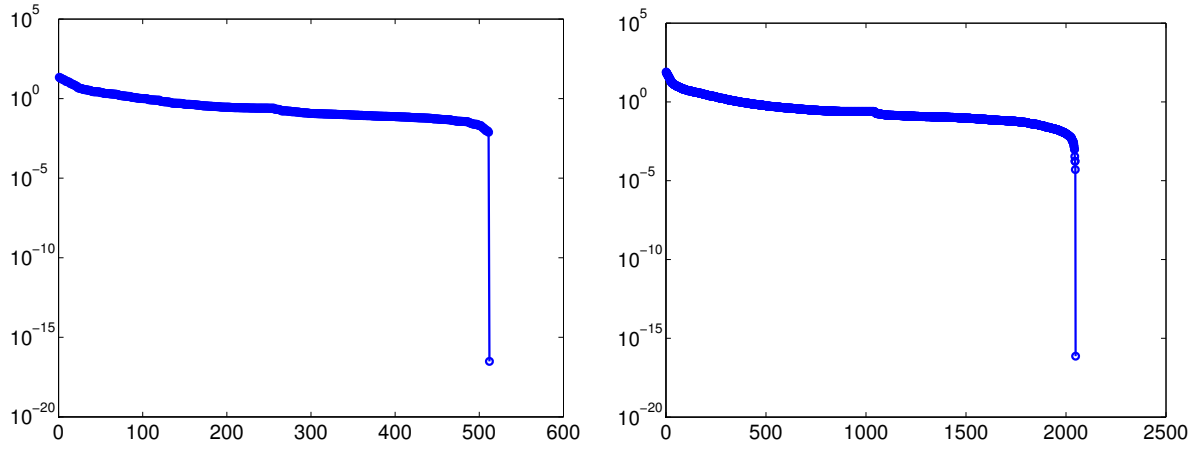

Figure 3: This figure shows the singular values for the matrix $\boldsymbol{W}$; the left plot shows singular values for the case $n_s = 32$ and $n_p = 16$, and the right figures shows singular values for the case $n_s = 64$ and $n_p = 32$.

These results illustrate that our model and computational approach to solve the atmospheric tomography problem, while not perfect at reconstructing atmospheric layers, still provides relatively good reconstructions considering we use only one set of collected data of a scene of interest. In particular, we do not do not collect multiple images, and we do not use laser guide stars. It appears that the null space component is important, and further work is needed to understand its significance, and how to incorporate its information into the solution.

4. CONCLUDING REMARKS

Previously proposed techniques for atmospheric tomography typically use a distribution of guide stars to obtain information from which to obtain a pseudo-3D reconstruction of atmospheric turbulence layers. Because this may not always be practical, or feasible, we propose an alternative approach using a WFS that is partitioned into a grid of subapertures. Our approach requires the collection of only one set of data of a scene of interest, and does not require the use of any laser guide stars. We provided a mathematical and computational framework to use the phase gradients of the local blur in each subaperture

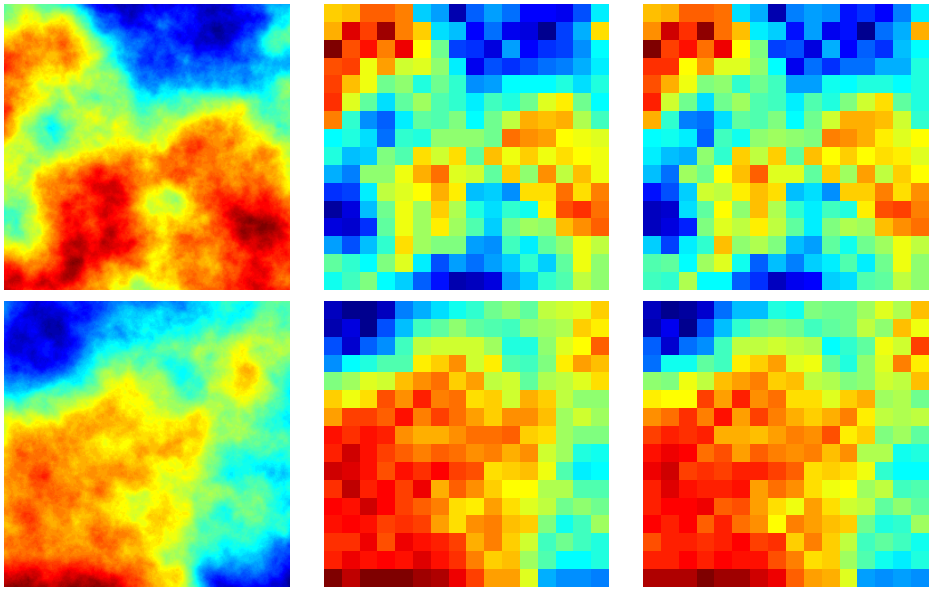

Figure 4: Comparison of true phase with reconstructed phases for the case $n_s = 32$ and $n_p = 16$. The top row is the phase the layer at 0 km and the bottom row is the layer at 5 km. The left column is the true phase, the middle column is the reconstructed phase using noise free data, and the right column is the reconstruction using noisy data.

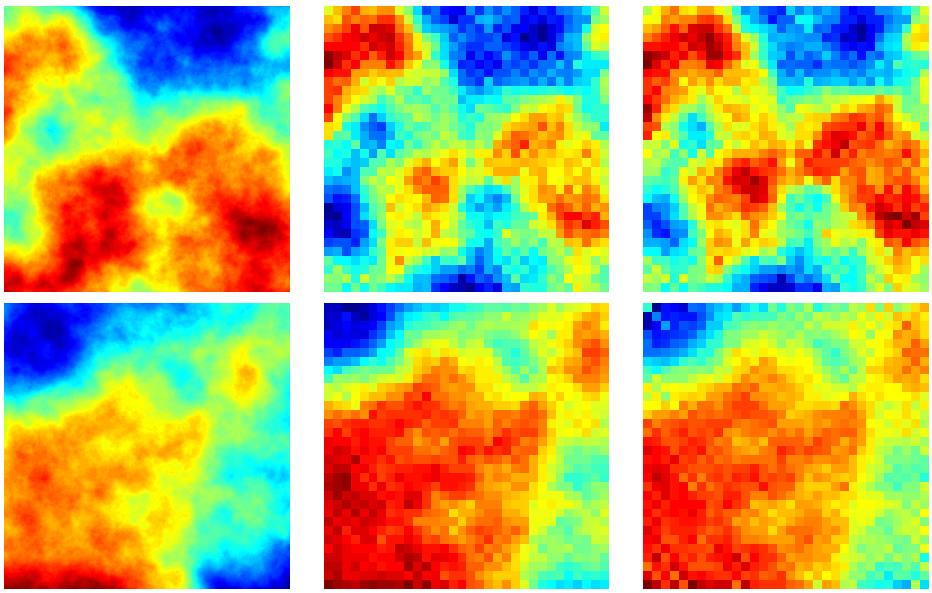

Figure 5: Comparison of true phase with reconstructed phases for the case $n_s = 64$ and $n_p = 32$. The top row is the phase the layer at 0 km and the bottom row is the layer at 5 km. The left column is the true phase, the middle column is the reconstructed phase using noise free data, and the right column is the reconstruction using noisy data.

to reconstruct phase gradients at a discrete set of turbulence layers, which is then followed by a phase reconstruction on each layer. The numerical experiments presented in this paper illustrate that our model can be an effective approach to solve the atmospheric tomography problem. There are still many open problems that we intend to address, including an analysis of the reconstruction matrix, $\boldsymbol{W}$, and more specifically, how to better regularize the null space components.

Acknowledgments

The work of J. Nagy is supported by grant no. DMS-1115627 from the US National Science Foundation, and grant no. AF9550-12-1-0084 from the US Air Force Office of Scientific Research.

REFERENCES

1. H. C. Andrews and B. R. Hunt. *Digital Image Restoration*. Prentice-Hall, Englewood Cliffs, NJ, 1977.
2. Y. Baharav, E. Ribak, and J. Shamir. Atmospheric tomography using a fringe patter in the sodium layer. *Optics Letters*, 19:242–244, 1993.
3. J. M. Bardsley. Wavefront reconstruction methods for adaptive optics systems on ground-based telescopes. *SIAM J. Matrix Anal. Appl.*, 30:67–83, 2008.
4. J. M. Bardsley, S. M. Knepper, and J. G. Nagy. Structured linear algebra problems in adaptive optics. *Advances in Comp. Math.*, pages DOI:10.1007/s10444–011–9172–9, 2011.
5. M. Bertero and P. Boccacci. *Introduction to Inverse Problems in Imaging*. Institute of Physics Publishing, London, 1998.
6. D. L. Fried. Least-square fitting a wave-front distortion estimate to an array of phase-difference measurements. *J. Opt. Soc. Am*, 67(3):370–375, 1977.
7. L. Gilles. Order-N sparse minimum-variance open-loop reconstructor for extreme adaptive optics. *Optics Letters*, 28:1927–1929, 2003.
8. P. C. Hansen, J. G. Nagy, and D. P. O'Leary. *Deblurring Images: Matrices, Spectra and Filtering*. SIAM, Philadelphia, PA, 2006.
9. R. H. Hudgin. Wave-front reconstruction for compensated imaging. *J. Opt. Soc. Am*, 67(3):375–378, 1977.
10. A. K. Jain. *Fundamentals of Digital Image Processing*. Prentice-Hall, Englewood Cliffs, NJ, 1989.
11. R. L. Lagendijk and J. Biemond. *Iterative Identification and Restoration of Images*. Kluwer Academic Publishers, Boston/Dordrecht/London, 1991.
12. R. Ragazzoni, E. Marchetti, and F. Rigaut. Modal tomography for adaptive optics. *A&A*, 342:L53–L56, 1999.
13. R. Ramlau and M. Rosensteiner. An efficient solution to the atmospheric turbulence tomography problem using Kaczamarz iteration. *Inverse Problems*, 28:DOI:10.1088/0266–5611/28/9/095004, 2012.
14. H. Ren, R. Dekany, and M. Britton. Large-scale wave-front reconstruction for adaptive optics systems by use of a recursive filtering algorithm. *Applied Optics*, 44(13):2626–2637, 2005.
15. M. C. Roggemann and B. M. Welsh. *Imaging through Turbulence*. CRC-Press, 1 edition edition, 1996.
16. M. Tallon and R. Foy. Adaptive telescope with laser probe: isoplanatism and cone effect. *A&A*, 235:549–557, 1990.
17. A. Tokovinin, M. Le Louarn, E. Viard, N. Hubin, and R. Conan. Optimized modal tomography in adaptive optics. *A&A*, 378:710–721, 2001.
18. C. R. Vogel. *Computational Methods for Inverse Problems*. SIAM, Philadelphia, PA, 2002.
19. Q. Yang, C. R. Vogel, and B. L. Ellerbroek. Fourier domain preconditioned conjugate gradient algorithm for atmpospheric tomography. *Applied Optics*, 45:5281–5293, 2006.
20. B. Neichel et al. GeMS first science results. Proc. AO4ELT3, eds. S. Esposito & L. Fini, http://ao4elt3.sciencesconf.org/ 2013.

21. P. L. Wizinowich. First Light Adaptive Optics Images from the Keck II Telescope: A New Era of High Angular Resolution Imagery. *Publ. Ast. Soc. Pac.*, 112:315–319, 2000.
22. P. Godara, A. M. Dubis, A. Roorda, J. L. Duncan, and J. Carroll. Adaptive optics retinal imaging: emerging clinical applications. *Optom. Vis. Sci.*, 87:930–941, 2010.
23. R. Q. Fugate. The Starfire Optical Range 3.5-m Adaptive Optical Telescope. *Large Ground-Based Telescopes* (Proc. SPIE), eds. J. M. Oschmann & L. M. Stepp, 4837:934–943, 2003.