

# PROCEEDINGS OF SPIE

[SPIDigitalLibrary.org/conference-proceedings-of-spie](https://spiedigitallibrary.org/conference-proceedings-of-spie)

## A fast wavefront reconstructor for the nonlinear curvature wavefront sensor

Codona, Johanan, Mateen, Mala, Hart, Michael

Johanan L. Codona, Mala Mateen, Michael Hart, "A fast wavefront reconstructor for the nonlinear curvature wavefront sensor," Proc. SPIE 10703, Adaptive Optics Systems VI, 107035Q (17 July 2018); doi: 10.1117/12.2313760

**SPIE.**

Event: SPIE Astronomical Telescopes + Instrumentation, 2018, Austin, Texas, United States

# A Fast Wavefront Reconstructor for the Nonlinear Curvature Wavefront Sensor

Johan L. Codona<sup>1</sup>, Mala Mateen<sup>2</sup>, Michael Hart<sup>1</sup>

<sup>1</sup>Hart Scientific Consulting International, LLC, 2555 N. Coyote Dr. #114, Tucson, AZ 85745

<sup>2</sup>Air Force Research Laboratory, Kirtland AFB, NM 87117

## ABSTRACT

The Nonlinear Curvature Wavefront Sensor (nCWFs), first proposed by Guyon,<sup>[1]</sup> determines wavefront shape from images of a reference beacon in a number of planes between the pupil and focal plane of a telescope. We describe a new algorithm that rapidly recovers the low-order aberrations accurately enough to allow practical use of the nCWFs in an adaptive optics (AO) system. The algorithm was inspired by refractive strong scintillation in the interstellar medium<sup>[2]</sup>, which behaves similarly to near-pupil linear curvature focusing, but over larger scales. The refractive component is extracted from the speckled images by binning with the lowest-order aberrations being additionally estimated through the use of first and second distribution moments. The linearity of the refractive scintillation process allows us to use a reconstructor matrix to compute an estimate of the pupil wavefront. The resulting wavefront estimate is then applied in reverse to a deformable mirror (DM), reducing the nonlinearity to the point that a single update phase retrieval algorithm such as a multi-plane version of Gerchberg-Saxton<sup>[3]</sup> (GS) can be used to estimate the remaining wavefront error (WFE). An AO simulation of a 1.5 m telescope, a 16x16 actuator DM, and four image planes show that the scintillation algorithm works, reducing ~800 nm rms WFE to ~40 nm, well below the fitting error (~90 nm) in closed loop. Once corrected to this level, the image planes still show a great deal of information that can then be used with a single-update wavefront retrieval algorithm. A couple simple variants of GS are suggested, including one that can be parallelized for each camera and run in parallel with the scintillation algorithm. A Monte Carlo study will be required to determine the best approach.

**Keywords:** Adaptive Optics; Wavefront Sensing; Nonlinear; Curvature; Scintillation; Strong Scattering; Refractive Scintillation.

## 1. INTRODUCTION

The *Nonlinear Curvature Wavefront Sensor* (nCWFs)<sup>[1,4]</sup> determines the pupil plane wavefront aberrations by examining the irradiance from a beacon in a number of defocused image planes located between the pupil and focal plane of a telescope. Since the irradiance patterns have structure at the full diffraction limit of the telescope and use light from the entire pupil, it is possible to obtain very high accuracy and efficiency. The potential advantages of the technique are independent of the algorithm used for extracting the wavefront from the images, so long as all of the diffraction-limited information is included.

We are motivated to explore wavefront estimation algorithms because the nCWFs overcomes limitations of wavefront sensors in general use in adaptive optics (AO) systems today. The Shack-Hartmann WFS and the modulated pyramid WFS, for example, are rather insensitive to low spatial frequency aberrations because they do not exploit information at the diffraction limit of the full aperture. On the other hand, the fixed pyramid WFS takes advantage of the spatial coherence of an unresolved beacon to achieve high sensitivity but suffers from limited dynamic range. By contrast, the nCWFs seeks to “have it all.” The concept stems from recognizing that the efficiency with which pupil plane phase aberrations on a given spatial scale are converted into measurable intensity variations is a function of the distance from the pupil. By placing several detectors at a range of distances, high sensitivity measurements may be made across the spectrum of aberrations. Furthermore, the morphology of the irradiance distributions is affected by the full optical path error rather than, for example, just the error modulo  $2\pi$ . High dynamic range is therefore also a feature of the nCWFs.

Guyon<sup>[1]</sup> and Mateen<sup>[4]</sup> considered an algorithm for processing the images that was a round-robin extension of the Gerchberg-Saxton (GS) algorithm<sup>[3]</sup>, where the field in the pupil was numerically propagated to each camera plane in turn. At each plane, the amplitude was replaced by the square root of the corresponding camera image, retaining the computed phase. This propagation and amplitude replacement process was repeated until the most distant image, at which point it was back-propagated to the pupil where the phase was retained as the starting point for the next iteration.

While this algorithm converged for some placements of the camera planes, it would typically require many iterations and was hundreds of times too slow for a real-time implementation. To implement a practical sensor for AO, a different approach leading to a much faster algorithm is required.

Inspired by the physical effects observed in strong scattering interstellar scintillation<sup>[2]</sup>, along with the refractive modulation character of the low spatial frequency portion of the irradiance fluctuation power spectrum<sup>[5]</sup>, it appeared to be possible to smooth the images to reveal large “refractive scintillation” features resulting from larger-scale aberrations in the pupil modulating smaller-scale speckles and diffraction features. The smaller-scale scintillation is related to the Fried length,  $r_0$ , while the larger-scale modulation is due to focusing across what is best described as the seeing disk projected back to the pupil plane. These large modulating patterns, taken by themselves, could be processed in a manner similar to a Linear Curvature WFS, but a solution involving all of the image planes, while still linear, would need to be more general. By smoothing or binning the irradiance, we can see a band-limited signal with a reasonably linear relationship to the pupil wavefront. By subtracting and binning the ideal irradiance patterns from the instantaneous scintillating versions, we can reasonably expect an approximately linear relationship between pixel values and wavefront displacements, usually evaluated at the actuator positions of a deformable mirror (DM). We use Singular Value Decomposition<sup>[6]</sup> (SVD) to find a reconstructor in a manner analogous to the Shack-Hartmann slopes in a conventional adaptive optics (AO) system<sup>[7]</sup>. As an optional additional aid in determining the lowest-order aberrations, the centroids and second moments of the irradiance in the image planes are included, minus their ideal values. The number of SVD modes are selected by finding the smallest rms residual wavefront error (WFE) on a set of test wavefronts.

The scintillation algorithm was tested in a closed-loop adaptive optics (AO) simulation and found to perform remarkably well. The difference between the actual scintillation images and those computed assuming the scintillation-based estimates become the starting point for a single-iteration GS-like update. We do not cover this part in depth here other than to note that it is feasible to make a useful update with a single iteration at the full resolution of the images.

## 2. INSIGHTS FROM SCINTILLATION THEORY

The information gathered for a nCWFS consists of irradiance images with different wavelength bands and distances from the pupil. Since, for our purposes, different distance and wavelength combinations are equivalent if they have the same Fresnel scale,  $R_f = \sqrt{\lambda z}$ , we will refer to the image planes as “Fresnel planes.” We will assume that the beam past the pupil is collimated, with the telescope magnification such that the various cameras have pixels, actual or binned, between one half to one Fried length measured in the science band. This allows us to use the full-resolution nCWFS images to give diffraction-limited correction for the science camera. Accessing the performance advantages of the nCWFS requires that we make use of the full resolution of the images; but we will start by extracting and using the much lower-resolution refractive scintillation signal from the Fresnel plane images in strong scattering. Utilizing these larger scintillation patterns is easy with simple processing and works well enough that a more sophisticated full-resolution algorithm can give us a useful final correction with only one iteration. Even more important is that the scintillation algorithm, used in closed loop, reduces the residual wavefront error to the point where the final full-resolution correction is at least mostly linear, allowing us to use a much simpler algorithm than required to retrieve the full uncorrected wavefront.

Figure 1 shows the irradiance inside of a telescope, measured at a number of distances from the pupil. In the upper left, we see the instantaneous irradiance at the pupil plane, shown here as uniformly illuminated. In reality, there will be scintillation in the incident light, but we will ignore it here. The justification is that under most astronomical observing circumstances the incident scintillation will be weak, while the wavefront aberrations will quickly lead to strong scintillation that will overwhelm the incident contribution. As we move away from the pupil, a varied irradiance pattern develops with darker regions separated by relatively bright boundaries. The spatial scale of this cell-like mottled pattern is close to the Fried length,  $r_0$ , and increases in contrast as we move farther from the pupil. Due entirely to the local slope of the wavefront and geometry, a simple geometric lag builds in the tilted wavefront as it propagates, where the leading edge is in the darker regions and is spreading the energy out and rays in the brighter regions inevitably move toward each other, increasing the irradiance. If we only consider the situation before the rays cross, we can see that the irradiance fluctuations are linearly related to the curvature of the wavefront since the change in energy density is related to the change in ray slope between two bounding rays. This is where conventional linear curvature wavefront sensing (CWFS) could be used. If the aberrated wavefront normals vary over an angular range of  $\theta_s = \lambda/r_0$ , the spacing between dark center cells is approximately  $r_0$ . We can expect rays from adjacent cells to cross when  $z \theta_s \approx r_0$ , or when the Fresnel scale matches the Fried length  $R_f = \sqrt{\lambda z} \approx r_0$  which occurs at a pupil distance of approximately

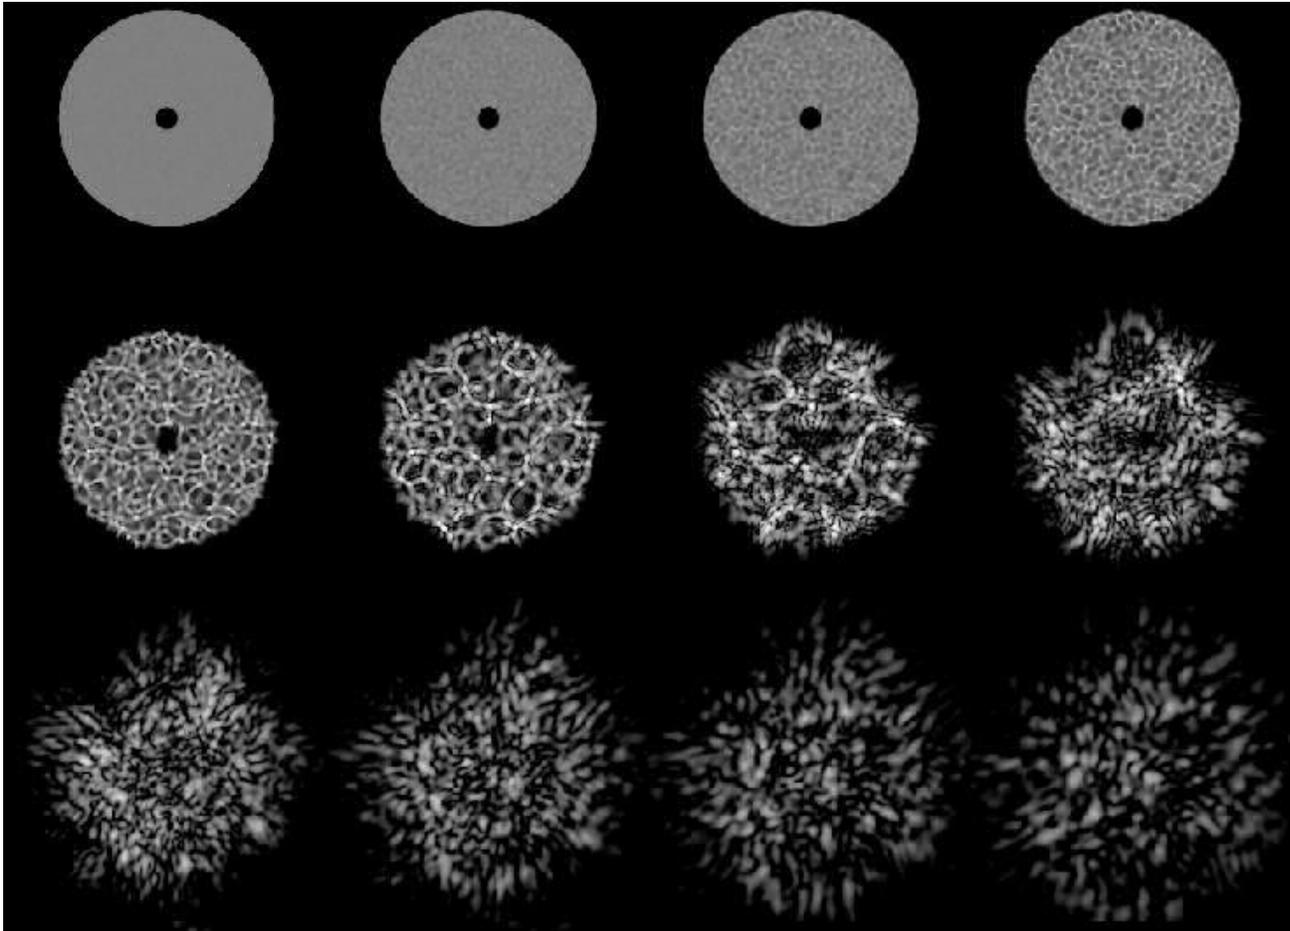


Figure 1. Scintillation inside of a telescope in a series of planes from the pupil to a point where the seeing disk nearly fills the pupil.

$z \approx r_0^2 / \lambda$ . Measuring relative to the telescope's exit pupil where both the pupil and the Fried length have been reduced by the magnification  $M$ , the usable distance for using CWFS is reduced by a factor of  $M^2$ . Once the rays from adjacent cells begin to cross, the irradiance begin to exhibit caustics that increase in intensity as the dark regions grow. Since the rays in the caustics have geometrically-related path lengths, the caustics exhibit diffraction fringes with scales related to the Fresnel scale. Also, when caustics cross, the irradiance fluctuations can become quite large, causing what is called the strong-focusing regime. Beyond this distance we enter into the strong scattering regime, where the irradiance at a given point results from the pupil field over multiple  $r_0$  patches, causing random interference fringes with size  $r_0$ , and a larger-scale focusing process that is possible over the size of the angular seeing disk, projected back to the pupil:  $L_s = \theta_s z = \lambda z / r_0$ . Note that  $L_s r_0 = R_f^2$  or that the Fresnel scale is the geometric mean of the smallest speckles and the largest refractive focusing modulation patterns. The refractive focusing scale continues to increase in size with increasing range until the seeing disk appears to fill the pupil aperture, after which there are no more ray crossings and the pattern never changes other than a simple proportional scaling with the distance from the pupil. These spatial scales and their sensitivities to wavelength are summarized graphically in Fig. 2. The horizontal axis is the log propagation distance beyond the pupil mask with scintillation regimes labeled. The widths of the various lines indicate their chromatic sensitivity, the dependence on wavelength being shown in red. The projected seeing disk is the least chromatic scale, which is consistent with it refractive in nature. The other scales depend in one way or another on waves interfering with waves, which directly depends on the wavelength. The weak dependency on wavelength in this case depends on the exponent in the turbulence spectrum's power law. If the exponent is different, as with turbulence that is not in equilibrium, the turbulence spectrum will be different and consequently show a different dependence on wavelength. There is no advantage to placing more than one Fresnel plane cameras to the right of the vertical gray bar denoting the seeing disk filling the pupil. Similarly, placing cameras to the left of where the Fresnel scale matches the

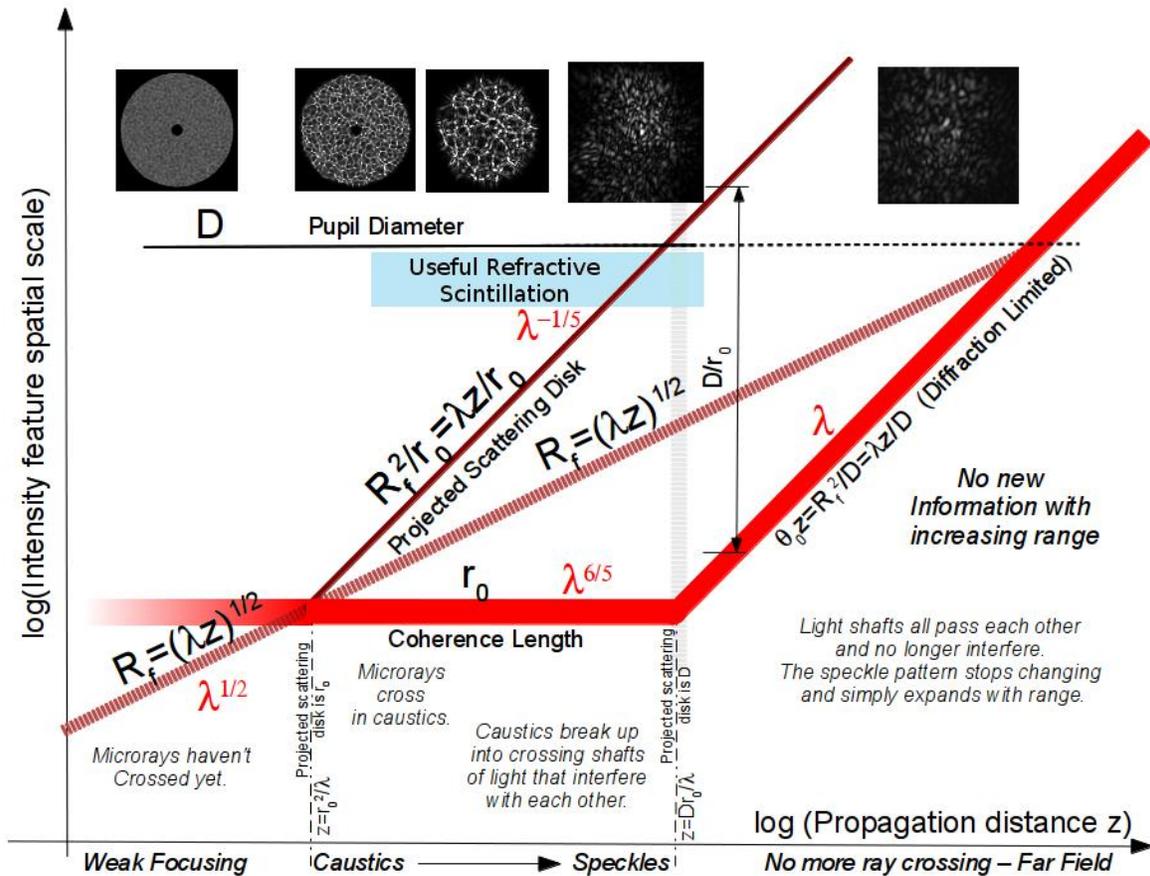


Figure 2. Spatial scales of scintillation beyond a pupil mask. Linear curvature wavefront sensing is only useful at the far left of the diagram, before caustics appear. The refractive scintillation effect is useful for our purposes from the onset of strong scattering ( $R_f=r_0$ ) to just beyond where the seeing disk appears to fill the pupil for an observer looking back (blue box).

Fried length, only gives useful information to roughly the Fried length at the wavefront sensor wavelength. If, as is common in astronomical AO, the science camera operates at wavelengths significantly longer than the WFS wavelength, with a correspondingly larger Fried length. Ideally for the scintillation algorithm, in contrast to the full nCWFS GS-like algorithms we will apply later, it makes most sense to use planes that are beyond where the caustics are breaking up and the Fried length and the projected seeing disk have had a chance to separate and work clearly as a modulation process. Between the nearest and most distant planes we chose to make the smoothing/binning size be proportional to the Fresnel scale, and equally space those scales. This also may not be the optimum answer, which will require a more complete analysis.

### Isolating the linear modulating focus signal

The first message from the above discussion is that linear CWS is simple, but limited in spatial scale to the “footprint” of the contributing rays on the pupil. Once the rays begin to cross, the relationship becomes too complicated to treat as a linear relationship between change in irradiance and wavefront curvature. The second message is that in strong scattering where the Fresnel scale is large compared with the Fried length, the scintillation tends to break up into a large focusing process modulating a smaller-scale pattern of diffraction speckles. If we could somehow isolate the effect of the large focusing process, we could treat it similarly to the linear CWFS. We can visualize the effect of this by running a Gaussian smoothing function over the irradiance with a spatial scale that lies between the small speckles and the focusing scale. We found that an effective choice is some fraction of the Fresnel scale. For the actual processing, we forego the smoothing and simply bin the Fresnel plane images to their respective smoothing scales. It is possible to smooth less, possibly leaving more of the linear modulation signal, with the cost of including more nonlinear signal. The nonlinear signal will appear to a linear reconstructor model as possibly-correlated noise, but it might be a win to push it somewhat. A Monte Carlo trade study will reveal the best answer.

Making a signal that could be reliably used in an approximate linear model of the modulation would need to have two additional properties: 1. Ignoring noise, it needs to be independent of the brightness of the reference beacon. 2. It should go to zero when the wavefront is perfectly flat. We can do this by normalizing the image by the sum of its pixel values, and subtracting a similarly-normalized image of the irradiance pattern when the incident wavefront is perfect.

### Image distribution moments help estimate the lowest-order aberrations

An optional step, one that does not add any processing time in a parallel implementation, is to compute the first and second image moments

$$\langle x \rangle = \frac{\sum x_{nm} I_{nm}}{\sum I_{nm}}, \langle y \rangle = \frac{\sum y_{nm} I_{nm}}{\sum I_{nm}} \quad (1)$$

$$\langle x^2 \rangle = \frac{\sum x_{nm}^2 I_{nm}}{\sum I_{nm}}, \langle xy \rangle = \frac{\sum x_{nm} y_{nm} I_{nm}}{\sum I_{nm}}, \langle y^2 \rangle = \frac{\sum y_{nm}^2 I_{nm}}{\sum I_{nm}}. \quad (2)$$

The first moments give the centroid of each image and allow us to make a linear estimate of the full aperture wavefront tip and tilt. Subtracting off the desired centroids gives a linear deflection signal that can be used along with the binned image values. The second moments similarly give us an estimate of defocus and astigmatism, but need to be made into linear estimates to give the best response. We do this by computing

$$\begin{aligned} \sigma_x &= \sqrt{\langle x^2 \rangle - \langle x \rangle^2} \\ \sigma_y &= \sqrt{\langle y^2 \rangle - \langle y \rangle^2} \end{aligned} \quad (3)$$

and

$$\sigma_{xy} = \sqrt{\langle xy \rangle - \langle x \rangle \langle y \rangle}. \quad (4)$$

Again, we subtract the values computed for the ideal patterns and we have signals that will allow us to adjust focus and astigmatism.

## 3. WAVEFRONT RECONSTRUCTION ALGORITHM

The preparation of the Fresnel plane images are illustrated in Figure 3. The images are from a simulation of a 1.5 m telescope with a 30 cm central obstruction. The telescope exit pupil is 6.3 mm which gives 100 pixels across the pupil for a camera with 21  $\mu\text{m}$  pixels, binned  $3 \times 3$ . This corresponds to a magnification of 238. The camera images were  $128 \times 128$ , allowing us to see light scattered outside of the nominal bounds of the pupil in more distant Fresnel planes. These choices were made to allow for plausible distances to the various cameras, since they scale as the inverse square of the magnification. In a real optical design of the sensor, it would make more sense to use a much larger plate scale, but for the purposes of illustrating the algorithm and visualizing its effects, we used the smaller pixels. We consider four Fresnel planes at a WFS reference wavelength of 700 nm. The Fried length  $r_0$  at 500 nm was assumed to be 5.6 cm, or 7.7 cm at 700 nm. For comparison with scintillation scales shown in Fig. 2, the distance beyond the exit pupil where the Fresnel scale is equal to the Fried length, is 15 cm. We placed four cameras at post-pupil distances of 5, 20, 40, and 75 cm. Relative to the Fried=Fresnel distance, they are 0.33, 1.32, 2.65, 4.96. This places the first camera in the weak scintillation regime where we could imagine measuring small-scale structure using linear CWFS, while the remaining three cameras are in strong scintillation, the final two being well beyond the caustics and strong focusing range. The images in the first row (a) of Figure 3 show the unaberrated irradiance resulting from an incident plane wave propagating parallel with the optical axis. Each image is assumed to be properly prepared with background subtraction, flat fielding, etc. before continuing. Each image is also normalized by its sum to make it nominally independent of the source brightness. The normalized reference images, measured with a reference source, should be comparable with the images obtained with perfect correction. A Kolmogorov phase screen was placed in front of the telescope, resulting in the four images in row (b), shown in negative gray scale for better visible contrast. Note that the character of the speckle patterns is as expected: the first image showing small cellular regions surrounded by the beginnings of caustics, the second camera showing caustics, and the third and

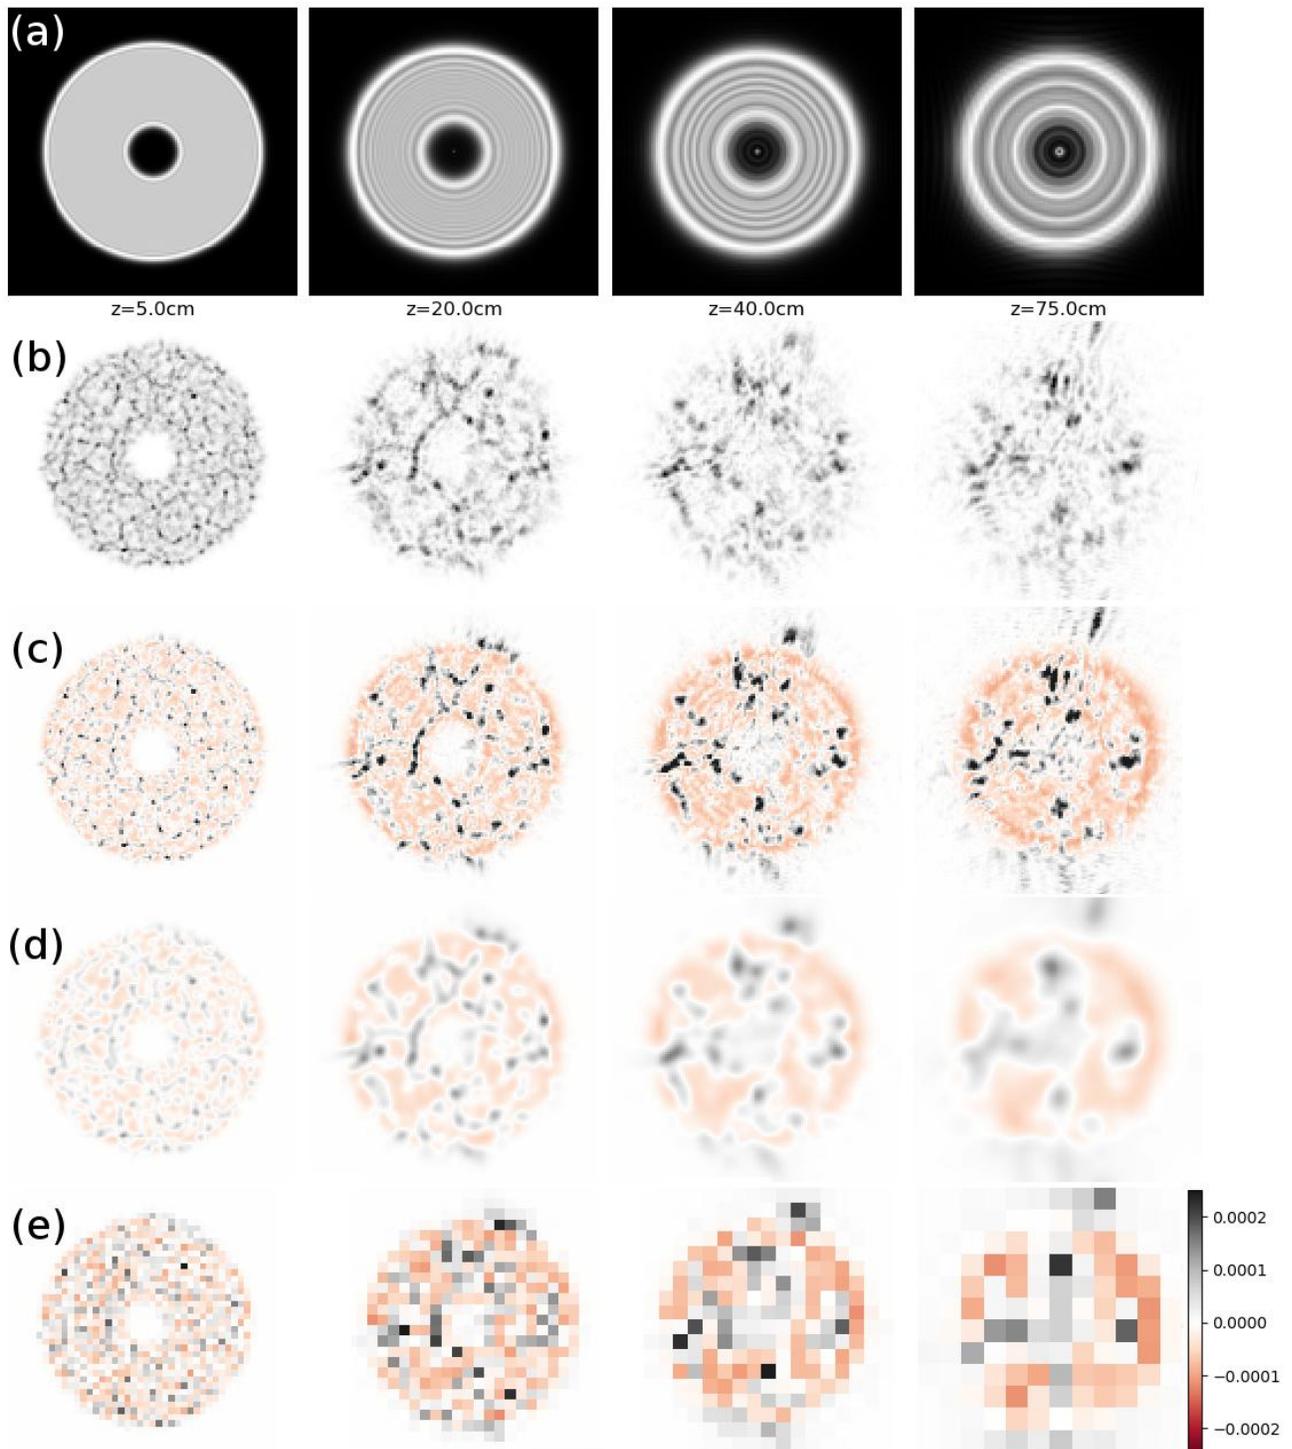


Figure 3. Image processing steps for the scintillation algorithm. Row (a) ideal reference images; (b) typical instantaneous aberrated images; (c) normalized instantaneous image minus normalized reference image; (d) difference images smoothed to reveal refractive scintillation modulation regions; (e) binning instead of smoothing the difference images prepares the data for use with the reconstructor matrix.

fourth cameras showing speckles with regions that are brighter or dimmer. The aberrated images were similarly processed and normalized, after which the ideal images were subtracted, shown in row (c). We call these images the delta irradiance. The remaining rows are displayed using a two-color negative gray scale with red indicating negative values. Even though it only really makes sense for the strong scintillation cameras 3 and 4, we smoothed the delta irradiance images with a Gaussian FWHM equal to half the Fresnel scale. This progressively shows focusing over larger areas of the pupil, largely unaffected by the now-subtracted reference diffraction pattern. The Gaussian smoothing is useful for visualizing the effect of binning, but is not necessary in practice. Instead we bin the delta irradiance into larger pixels which are then used in the linear wavefront reconstruction model. The binning values used here are 3, 5, 7, and 10. Finally, the down-sampled pixel data from all of the cameras is placed into a long column vector,  $\delta I$ . We also included our centroids and width estimation moments minus their ideal values as five more elements of the column vector.

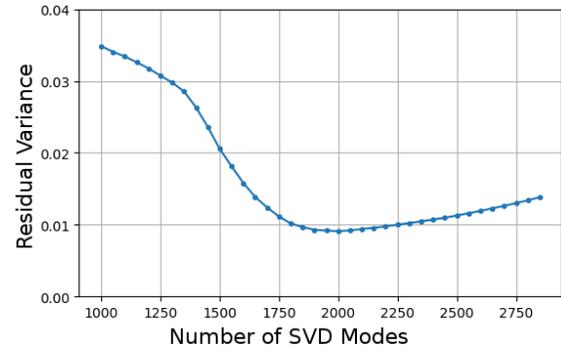


Figure 4. Residual variance of the reconstructed test wavefronts at the actuator locations. We used 1900 modes to build our reconstructor.

To use this as a wavefront sensor in an adaptive optics application, we can imagine a plane wave illuminating the DM with the actuator values set to the elements of the column vector  $\alpha$ . When all of the actuator values are zero, the Fresnel plane cameras should see the ideal patterns and the delta irradiance vector should be zero. If our hypothesis is reasonable that the smoothed delta irradiance is dominated by refractive scintillation, we should be able to relate the actuator vector to the delta irradiance vector with a reconstructor matrix  $\mathbf{R}$ ,

$$\vec{\alpha} = \mathbf{R} \delta \vec{I}. \quad (5)$$

To find the reconstructor, we followed the familiar SVD pseudo-inverse process, commonly used with conventional AO to relate actuator displacements with WFS influence functions. We used 5000 randomly created Kolmogorov phase screens as our training set, recording the wavefront displacements at the DM actuator locations, but using the full-resolution phase screens to compute the Fresnel plane images. If the reconstructor model is correct, then Eq. (5) will be true for each pair of actuator and delta irradiance vectors. This allows us to write them together as a set, with the actuator and delta irradiance data forming large matrices related by the reconstructor matrix

$$\alpha = \mathbf{R} \delta \mathbf{I}. \quad (6)$$

Singular Value Decomposition allows us to find the closest least-square approximation to the pseudo-inverse of the delta irradiance matrix, the accuracy of which is determined by how many singular values (SVD-modes) we include.

$$\mathbf{R} \approx \alpha [\delta \mathbf{I}]^{-1}. \quad (7)$$

As with many uses of SVD, it is possible to over-fit the data, resulting in worse performance on similar data that was not in the training set. To determine the best number of modes, we computed 500 additional phase screens, and computed the residual variance as a function of included singular values (Figure 4). Since the variance increases beyond 1900 modes, we used 1900 modes to build the reconstructor.

### Open and closed-loop performance

The rms wavefront displacement in the test data was 794 nm, averaged over the illuminated actuators. The rms residual between the true value and the reconstruction was 75 nm, about 9.5% of the starting error. This one-shot open-loop correction is improved when iterated in closed loop, ultimately giving us an rms WFE of 40 nm, computed at the actuators. Since the actuator spacing at the entrance pupil was 10 cm, the DM fitting error was about 90 nm, considerably larger than the scintillation reconstruction residual. Note that photon or read noise was not included in this calculation, and the results are appropriate for a bright star limit.

Figure 6 shows five open-loop wavefront reconstruction examples. For each row, the left image is the phase screen wavefront displacements measured at the actuators. The circular aperture edges run from the center of the first column

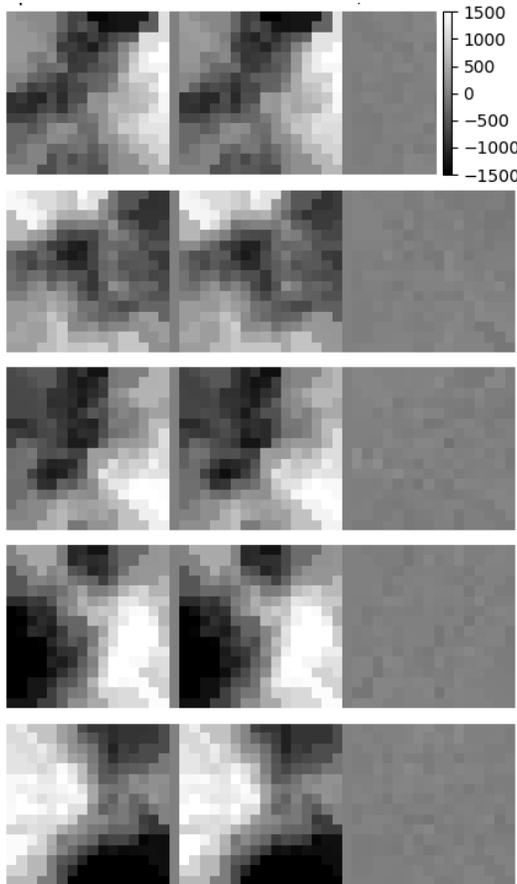


Figure 5. Five open-loop wavefront reconstruction examples. For each row, the left image is the phase screen wavefront displacements measured at the actuators. The center image is the reconstruction, and the right image is the residual minus any mean.

wavefront isn't known. Second, we can learn what types of residual wavefront errors are to be captured in the abbreviated full-resolution update by looking at the morphology of the computed images relative to the actual ones. Figure 9 shows this for the initial state, first, second, and fourth updates of the closed-loop simulation. Each box contains four rows, one for each Fresnel plane. In each row, the image on the left is what was seen by the camera. All four images were processed as described above, along with the moment terms, and the residual wavefront computed. This is an estimate of the incident wavefront minus the current DM setting, so it corresponds to a plane wave incident on a DM with only the computed actuator settings. Even though we can see smaller details in the images, we only have control over the actuators, so we set a test DM to the reconstructed actuator values and proceed to compute the corresponding images at the Fresnel planes, shown in the second image of each row. "Time step 1" is the initial look at the wavefront, before there are any corrections. The first plane image shows some interesting artifacts that are due to the cubic interpolation function. Presumably, this will not happen if we use a 2D spline-based interpolation or a mechanical model. By the second through fourth Fresnel planes, the irradiance appears artifact-free. The third image in each row is

to the last, with 15 actuator spacings between them. Unilluminated actuators are set to the value of the nearest illuminated actuator.

The closed-loop simulation was simply to move the phase screen with the wind (horizontal motion of 2 m/s with 1 ms time steps). The plane was is passed through the phase screen, the entrance aperture, the DM its most recent actuator settings interpolated to a finer surface grid using a 2D cubic Delauney interpolation (`scipy.interpolate.griddata`). The field's scale was then changed according to the telescope magnification and propagated to each of the Fresnel planes where the images were computed, multiplied by the reconstructor and subtracted (unity gain) from the current actuator values. Interferometric images of the residual wavefront for the first four updates in closed loop are shown in Figure 6 at 700 nm. Note that the low spatial frequencies are immediately suppressed, which was our main goal. After two updates, the residuals are dominated by small-scale structure at or below the actuator spacing.

Figure 7 shows the instantaneous rms WFE during a 1 second simulation, along with an estimated terms in the error budget. The fitting error was largest at  $\sim 90$  nm, the scintillation algorithm residual, measured on the actuators, was  $\sim 40$  nm. The last term is the estimated residual after using an *ad hoc* GS-like single update, leaving us with an rms WFE of  $\sim 25$  nm.

Finally, a J-band science image is shown on a 4-decade log scale in Figure 8.

#### 4. THE FINAL UPDATE

The scintillation algorithm appears to work quite well, but it cannot alone give the performance advantages of the nlCWFS. To do that, we need to use the Fresnel plane images at the full diffraction limit of the telescope. To see how well we have succeeded, we can measure the residual wavefront error, or we can use the wavefront estimate to compute the Fresnel plane images and compare them to the camera images. This has at least two advantages: First, we can do the image comparison with actual data taken in the field, even when the true wavefront isn't known. Second, we can learn what types of residual wavefront errors are to be captured in the abbreviated full-resolution update by looking at the morphology of the computed images relative to the actual ones. Figure 9 shows this for the initial state, first, second, and fourth updates of the closed-loop simulation. Each box contains four rows, one for each Fresnel plane. In each row, the image on the left is what was seen by the camera. All four images were processed as described above, along with the moment terms, and the residual wavefront computed. This is an estimate of the incident wavefront minus the current DM setting, so it corresponds to a plane wave incident on a DM with only the computed actuator settings. Even though we can see smaller details in the images, we only have control over the actuators, so we set a test DM to the reconstructed actuator values and proceed to compute the corresponding images at the Fresnel planes, shown in the second image of each row. "Time step 1" is the initial look at the wavefront, before there are any corrections. The first plane image shows some interesting artifacts that are due to the cubic interpolation function. Presumably, this will not happen if we use a 2D spline-based interpolation or a mechanical model. By the second through fourth Fresnel planes, the irradiance appears artifact-free. The third image in each row is



Figure 6. Pupil field interferograms at 700 nm for the initial state through the first 4 updates.

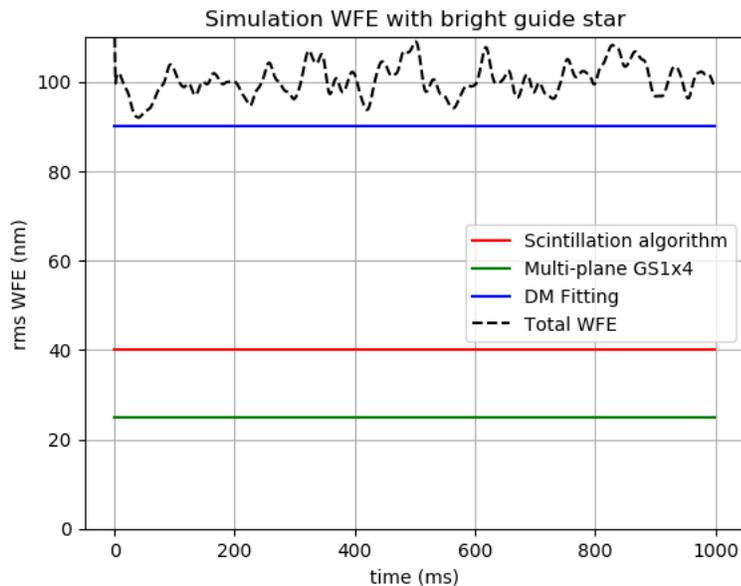


Figure 7. The instantaneous rms WFE during a 1 sec simulation. The blue line indicates the DM fitting error of 90 nm while the red line indicates the residual error from the refractive scintillation algorithm. The green line is an estimate of the WFE after a single update of the GS1xN algorithm.

the linear refractive scintillation algorithm.

This is exactly what we had hoped to see. By explicitly isolating the linear portion of the scintillation due to large-scale refraction, we have reduced the wavefront residual WFE to a small enough value that it is at least significantly less nonlinear. The reason this works so well is that it is in an adaptive optics closed loop that removes most of the reason the open loop problem was nonlinear. With each iteration, we have changed the problem into an easier one to solve with less sophisticated algorithms.

So what to do about the final residual errors, the effect of which are fine speckle structure in each Fresnel plane. We have not yet explored these in detail, but there is clearly room for innovation. The only way to decide the best approach is to implement each algorithm or variant, and see how it performs in a Monte Carlo. We will only discuss one possible idea here that is based on the observation that the residual images quickly settle down to diffractive scintillation features that are not seen by the scintillation algorithm. In fact, as far as the refractive scintillation algorithm is concerned, the

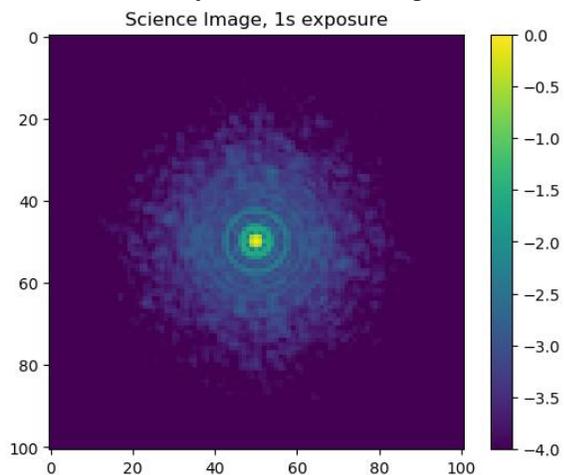


Figure 8. J-Band science image after 1 second using the scintillation algorithm. The color scale is 4 decade log.

the difference between the actual and computed images using the same positive-negative negative gray scale where negative image differences are shown in red. If we had the exactly right and were able to reproduce the turbulent wavefront accurately on our DM, the difference image would be zero. Instead, each plane shows a rich collection of positive and negative speckles. After a single update (“Time step 2”), the caustics are suppressed at all ranges in the actual images and the caustics in the computed images are replaced by mid-range focusing where caustics used to be, along with some strong isolated focusing features, larger than the  $r_0$ -sized features evident in the first plane where the scintillation is weak and presumably linear. The difference images continue to show a great deal of fine structure, but the caustics and the refractive scintillation regions are rapidly suppressed. After the second update (“Time step 3”), even these features are suppressed, and the computed images settle down to looking much like the ideal images, while the actual images continue to show small-scale scintillation features that are not seen by

the wavefront is nearly perfect. Since the computed update, which will be changing slowly like the low spatial frequency features blowing past the telescope, would be the canonical starting point for a Gerchberg-Saxton-like algorithm, we could argue that a reasonable first approximation to the pupil wavefront is zero. This has the advantage that the complex field can be precomputed once and used with the square root of the Fresnel plane images as soon as they are available. This allows the small-scale updates to be computed in parallel with the scintillation updates. Furthermore, since the remaining scintillation features are small as seen by each camera, it would appear that the scintillation AO servo has largely left wavefront features that focus onto one or the other planes where the pixel binning makes them relatively undetectable. If so, this suggests that we can likewise perform a mainly traditional GS update for each Fresnel plane alone, each giving an estimate of the residual wavefront that we can average and include in the next AO DM update. This further simplifies the computing by allowing the image from each camera to be processed independently and in parallel.

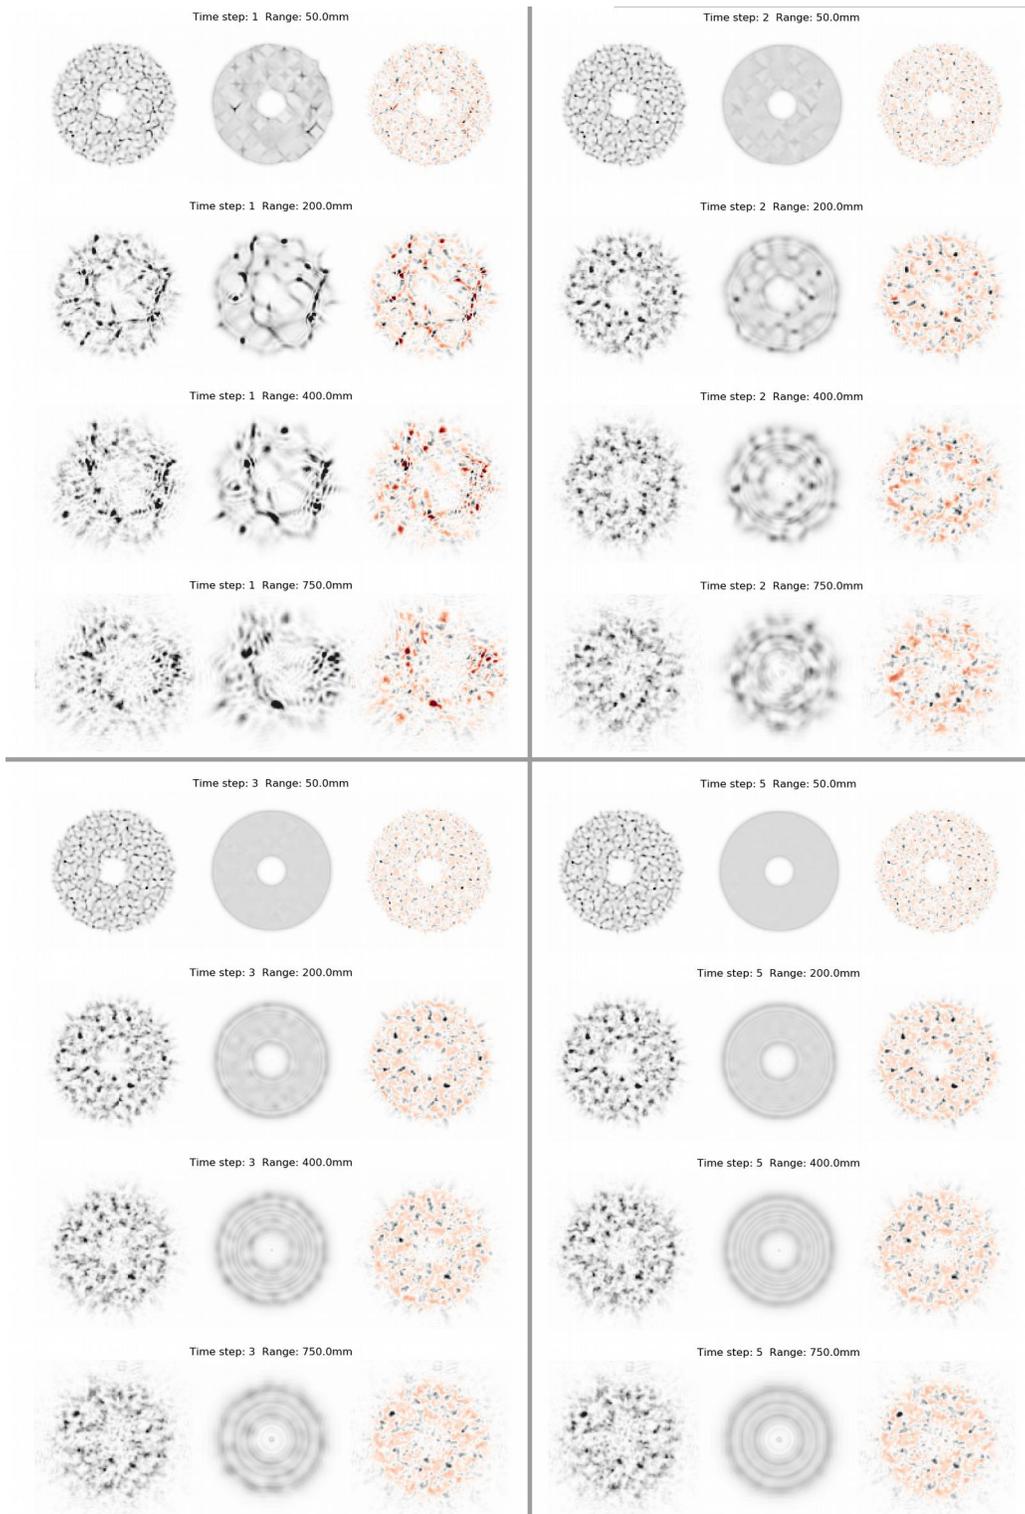


Figure 9. Fresnel plane image residuals for the initial state (top left), and after 1, 2, and 4 updates. Each row contains the actual Fresnel plane image from the camera; the image computed from the latest wavefront estimate using the scintillation algorithm; and the difference between the two.

The algorithm we are calling GS1xN operates on one camera at a time, with N cameras being combined in the end. The complex field at each Fresnel plane is precomputed and its complex phasors at every pixel (  $\exp i \phi_{nm}$  ) extracted and kept in memory. When the next image is available, we use the square root of the preprocessed image as an estimate of the amplitude (units and normalization don't matter here), and the estimated Fresnel plane field

$$\psi_{nm} = \sqrt{I_{nm}} \exp(i \phi_{nm}) \quad (8)$$

Propagated back to the pupil plane using the usual FFT-based paraxial propagation technique. Once back in the pupil plane, the phase is computed and unwrapped and averaged with the estimates from the other cameras. We implemented this method in the closed loop simulation and did achieve a reduction in the noise-free rms WFE, estimated to be approximately 25 nm. For the simulated case however, there was no noticeable improvement in the result. However, this will not be the situation for all use cases.

The GS1xN algorithm is simple and very fast, but it may be that we need to start the wavefront retrieval starting from the wavefront estimate computed by the scintillation algorithm. In that case, we cannot compute the phasors in the Fresnel planes until the scintillation algorithm is complete, although computing square roots of the images can be done in parallel. If we do some form of GS1xN with a better initial estimate of the pupil field, we can still process each camera's wavefront estimate in parallel. However, if we use some form of GS round-robin, we will have to process each camera in sequence, which will take the greatest amount of time. Only a Monte Carlo "bake-off" will be able to decide the best approach, but using a simple and highly parallelizable algorithm like GS1xN may allow a second iteration in the tight time budget.

## 5. CONCLUSIONS

We have shown that the images used to drive the high-performance nonlinear curvature wavefront sensor (nlCWFS) can be made into a practical real-time system by smoothing the images to extract and use the strong refractive scintillation signal that causes the system to be so nonlinear to begin with. Correcting this signal with a linear reconstructor model rapidly eliminates the caustics and regional modulation characteristic of the large-scale aberrations that are traditionally difficult to correct. The remaining information that is unused by the scintillation algorithm is clearly displayed by the real-computed difference images, providing a useful diagnostic. The remaining wavefront error can be estimated using some form of a single-update Gerchberg-Saxton algorithm, with several unexplored approaches being suggested, each working on a much simpler problem than the one of accurately determining a large amplitude wavefront that is highly nonlinear.

Larger optical bandwidth Fresnel plane images will be adversely affected by atmospheric dispersion without an atmospheric dispersion corrector (ADC). Even with an ADC, the speckles will show a chaotic array of dispersion patterns, depending on the local prism effect where a given speckle originates. This phenomenon should be minimized very rapidly after beginning the closed-loop correction. It would also appear that the scintillation algorithm should be relatively unaffected when using an extended source as a beacon, so long as it is not significantly larger than the seeing disk.

## ACKNOWLEDGMENTS

This work was supported by AFRL contract FA9451-12-C-0004.

## REFERENCES

- [1] O. Guyon, "High Sensitivity Wavefront Sensing with a non-linear Curvature Wavefront Sensor," *PASP*, **122**:49-62, 2010. DOI: [10.1086/6496](https://doi.org/10.1086/6496)
- [2] Narayan, Ramesh. "The Physics of Pulsar Scintillation." *Philosophical Transactions: Physical Sciences and Engineering* 341, no. 1660 (1992): 151-65. <http://www.jstor.org/stable/53917>.
- [3] R. W. Gerchberg and W. O. Saxton. "A practical algorithm for determination of the phase from image and diffraction planes.," *Optik*, 35:237-246, December 1972.
- [4] M. Mateen, "Development and verification of the non-linear curvature wavefront sensor," Univ of Arizona Dissertation, 2015.
- [5] J. L. Codona, S. M. Flatté, D. B. Creamer, R. G. Frehlich, and F. S. Henyey, "Solution of the fourth moment of waves propagating in random media." *Radio Sci.*, 21:929-948, 1986.

- [6] Press, W. H.; Teukolsky, S. A.; Vetterling, W. T.; Flannery, B. P. (2007), "[Section 2.6](#)", *Numerical Recipes: The Art of Scientific Computing* (3rd ed.), New York: Cambridge University Press, [ISBN 978-0-521-88068-8](#)
- [7] Tyson, R. (2015). *Principles of Adaptive Optics*, Fourth Edition. Boca Raton: CRC Press.