

Real-time Shack-Hartmann wavefront sensor reference scaling compensates for beacon size

Michael D. Olikier^a, Michael J. Steinbock^b, Daniel E. Roskey^a, Odell R. Reynolds^b,
and Miles D. Buckman^b

^aLeidos, ^bAir Force Research Laboratory
Starfire Optical Range, 3550 Aberdeen Ave SE, Kirtland AFB, NM, USA

ABSTRACT

Past research in adaptive optics (AO) has demonstrated the link between apparent beacon extent and wavefront gradient estimation sensitivity of a classical Shack-Hartmann (SH) subaperture, particularly when using quad-cell detector regions. In most practical SH AO systems that utilize sodium beacon laser guide stars, the subaperture diameters are large enough to resolve beacon features. This is particularly relevant on larger telescope apertures where beacon elongation is present. Additionally, residual wavefront error can further broaden the subaperture point spread functions as the atmospheric seeing varies in time. This paper presents research conducted at the Starfire Optical Range over the past 8 years in implementing a robust solution on the site’s 3.5m telescope and AO system. Emphasis is given towards the practical aspects that must be considered beyond the pure theory, which has been presented in several prior works. A high signal-to-noise strategy has been implemented that estimates the aperture-averaged subaperture sensitivity (related to beacon size) in real-time by exploiting the null-space of the least-squares wavefront reconstructor. Careful consideration has gone into the implementation of this estimation method to avoid unintended effects, particularly at low light levels. Moreover, consideration of how to best utilize the estimated subaperture sensitivity within the AO pipeline will be presented.

Keywords: Wavefront sensing, adaptive optics, beacon elongation, wavefront reconstruction

1. INTRODUCTION

A Shack-Hartmann wavefront sensor (SHWFS) aims to measure the irradiance-weighted wavefront gradient over each of its subapertures. Subapertures consist of a lenslet that focuses pupil light to a focal plane array (FPA) at the lenslet’s focal distance. One system design trade must establish whether to design the subaperture with a quad-cell or a larger $N \times N$ pixel sensing region at the FPA. While the larger FPA sensing region can allow larger linear dynamic range for measuring gradients as well as potentially lower gradient estimation errors, it also requires lower pixel instantaneous fields of view (IFOV). In other words, the formed subaperture spots must span more than one pixel in an $N \times N$ sensing scheme so that the gradient estimation is linear as a spot transitions across the sensing region. On the other hand, a quad-cell sensing region can support larger IFOV’s, where the spot formed can be smaller than a pixel width. Figure 1 graphically depicts these design options.

In effect, the forced broadening of the spot for an $N \times N$ subaperture sensing array accumulates more sensor dark current and read noise than a quad-cell—simply due to the increased pixel count. This in turn requires the $N \times N$ sensing design to have higher beacon return fluxes to maintain equivalent signal-to-noise ratio (SNR). For applications or system designs that are signal starved, an active quad-cell sensing geometry can be attractive for this reason.

Unfortunately, the quad-cell sensing geometry comes with a complication avoided in the $N \times N$ geometry. As the apparent beacon size changes, so does the subaperture’s estimation of wavefront slope for a given input tilt. In the $N \times N$ geometry, since the diffraction-limited spot spans multiple pixels, enlarging the spot further has very minimal effect. Figure 2 compares this impact for both FPA geometries. This relationship between a subaperture’s input wavefront tilt and the measured centroid movement defines the subaperture’s sensitivity. More specifically, the sensitivity is the linear region’s slope of the “s-curves” about zero input tilt. Given this

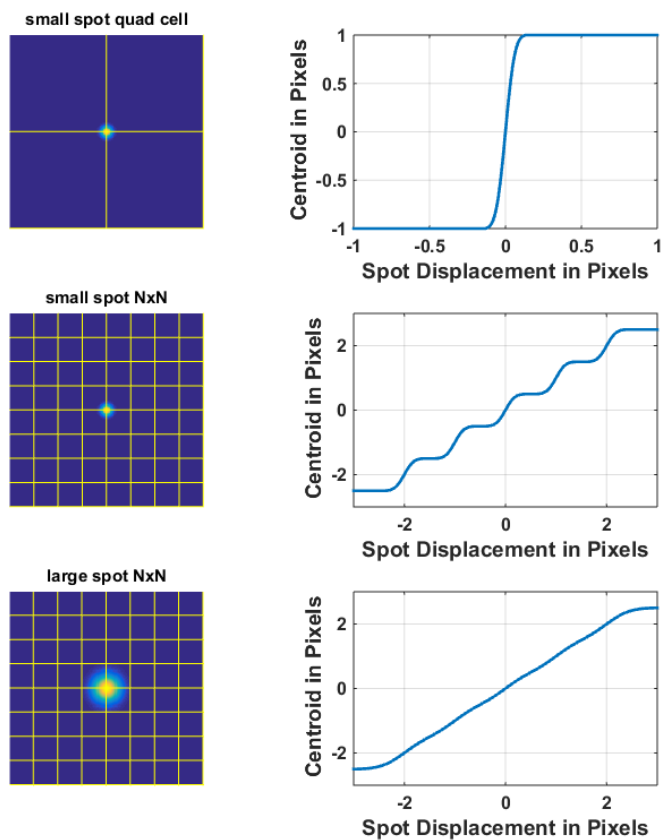


Figure 1. Three Shack-Hartmann WFS optical designs, with the spot relative focal-plane array on the left column and the corresponding centroid response “S” curves in the right column.

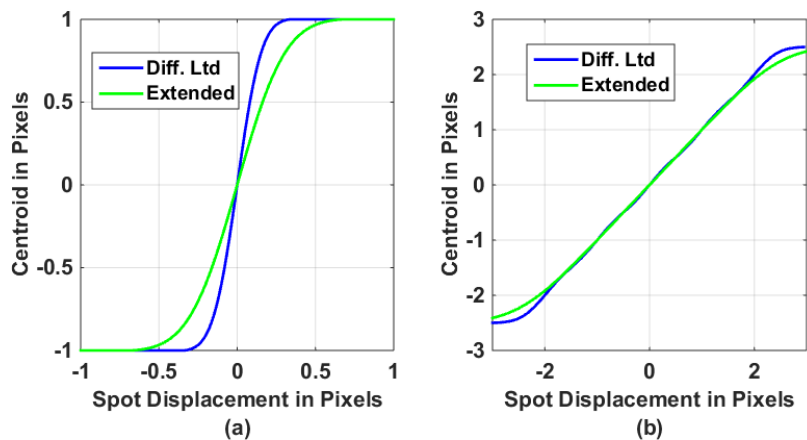


Figure 2. Comparison of a quad-cell (a) versus an $N \times N$ (b) sensing geometry Shack-Hartmann WFS in terms of their centroid response curves for two different spot sizes. In the case of the $N \times N$ geometry, the two spot sizes produce nearly the same response curve, but the same is not true for the quad-cell geometry.




Figure 3. A graphical depiction of “reference bleed-through.” An example residual phase error is sensed on a WFS (a). That same WFS was calibrated to have a hypothetical reference aberration (derived from \vec{s}_{ref}) (b). If the WFS then made a measurement with a beacon that led to 0.5 sensitivity, (c) shows the reconstructed measurement. Without sensitivity scaling, the full-sensitivity reference would be subtracted out of the measurement, leaving the right figure to be applied to the DM. Note the reference sine wave printed on (c) is inverted at half-magnitude because of the bleed-through.

definition, it is clear that the s-curve slope changes between the two spot size inputs to the quad-cell subaperture in Fig. 2.

Two distinct issues manifest for the quad-cell geometry that must be addressed:

1. The overall SHWFS gain changes based on the size of the beacon. This trickles down into the overall system gain. System operation will be suboptimal unless these changes are accounted for. Other factors such as sensing SNR also must be factored here.
2. When subtracting the lab-measured point-source gradient reference vector from the measurements, each vector set will have been measured with different effective WFS sensitivities. In extreme cases this leads to “reference bleed-through,” where the high sensitivity point-source reference gradients are subtracted from on-sky measurements with low sensitivity. Figure 3 depicts this graphically in a one-dimensional simplified case.

Note that the overall system (control loop) gain, referred to in the first point, cannot be used to compensate for the second point. This is because the system gain multiplier is applied in a step after the reference gradients are subtracted from the measured gradients.

This paper presents a real-time sensitivity estimator algorithm and guidance on incorporating an appropriate correction to the AO data processing pipeline, aimed at addressing the second issue above. Traditional system gain optimization schemes can then be implemented to compensate for the overall system gain. As mentioned, this work applies only to SHWFS’s that use quad-cell sensing geometries in each subaperture. Additionally, work is underway to evaluate this algorithm’s utility towards pyramid WFS’s, as the pyramid WFS behaves similarly to quad-cell SHWFS’s.

1.1 Prior Work

This paper is far from the first to examine this quad-cell SHWFS sensitivity issue. In 2000, Véran *et al.* proposed using the SHWFS’s reference gradient set (\vec{s}_{ref} in Sec. 2.1.1) to estimate sensitivity.¹ By assuming stationarity, one can average measured WFS gradients sufficiently long enough that the average should approach a sensitivity scaled version of \vec{s}_{ref} . This has challenges in capturing all of the local system aberrations as well as averaging

Send correspondence to: AFRL.RDSS.OrgMailbox@us.af.mil

sufficiently long enough to truly see the turbulence content average out. In 2003, van Dam modified this technique to use the slope discrepancy sub-space, as opposed to the entire modal space.² This modification avoids the issues of averaging out the atmospheric content, and initial success was presented for the Keck I observatory. However, later conversations with the Keck team indicated practical issues limited the overall capability, which have precluded ongoing use.³ In 2010, researchers from the Starfire Optical Range presented work underway to implement and transition a variant of the null-space sensitivity estimation algorithm up to the main telescopes.^{4,5} At this point, only simulation and laboratory experimentation had been conducted.

This paper pushes forward from these prior efforts, reporting on-sky operational lessons learned and practical algorithmic tweaks. A surprising amount of learning was required in this transition from laboratory to on-sky use. It is hoped that this paper will ease similar transitions for other observatories.

2. SENSITIVITY ESTIMATOR AND COMPENSATION

In designing an algorithm to estimate and compensate for sensitivity changes to a SHWFS, a number of desirable algorithm traits were identified:

1. The algorithm should be able to estimate the WFS sensitivity whether the AO or tracking system is in open or closed loop operation.
2. The algorithm should provide sensitivity estimates with high repeatability and low variability.
3. The algorithm should utilize a high SNR approach that can operate even when the WFS beacon is dim.
4. It is undesirable to introduce dithers due to the inherent impact on system performance.

These considerations, if met, form a practical solution over the wide operational range of most telescope systems.

The following subsection defines the proposed algorithmic approach to meet the aforementioned traits.

2.1 Algorithm Formulation

Starting with the desire for open and closed loop functionality, coupled with a non-dithered approach, it was noted that modes within the WFS reconstructor’s null-space might be exploited. Least-squares WFS reconstructors are generally formulated by pseudo-inverting the equation

$$\vec{s} = \mathbf{G}\vec{\phi}, \quad (1)$$

where \vec{s} is a vector of phase gradients, $\vec{\phi}$ is a vector of phases, and \mathbf{G} defines the relationship between the phases and gradients, called the “geometry matrix.”^{6,7} For AO, \mathbf{G} defines the relationship between applied deformable mirror (DM) commands to sensed wavefront gradients. It is noted that there are a number of ways \mathbf{G} can be defined or identified for a system, but these specifics are generally inconsequential to the discussion and math that follows. Pseudo-inverting Eq. 1,

$$\hat{\phi} = (\mathbf{G}^T \mathbf{G})^{-1} \mathbf{G}^T \vec{s} = \mathbf{R}\vec{s}, \quad (2)$$

the wavefront reconstructor \mathbf{R} is found. As seen in Eq. 2, \mathbf{R} allows calculation of wavefront estimates given wavefront slope measurements \vec{s} from a WFS.

Given \mathbf{G} and \mathbf{R} , a corresponding null-space projection operator \mathbf{P} can readily be calculated

$$\mathbf{P} = \mathbf{I} - \mathbf{G}\mathbf{R}, \quad (3)$$

where \mathbf{I} is the identity matrix. Remember that \mathbf{P} operates on wavefront slopes \vec{s} , isolating any slopes that are within the null-space of \mathbf{R} . In other words, \mathbf{P} ’s range includes any slope modes that would not be reconstructed into contributions of $\hat{\phi}$. Because these gradient components do not impact the resultant wavefront reconstructions, any gradient energy within this null-space should be preserved, even when the AO system is operating. Going further, the specific gradient modes within \mathbf{R} ’s null-space can be examined by performing a singular value decomposition (SVD) of \mathbf{P} .⁸ While the specific modal content of \mathbf{P} , and thus \mathbf{R} ’s null-space,

will vary system-to-system (and potentially day-to-day for influence function based reconstructors), there will generally be a subspace that consists of all possible gradient circulations for that system’s sensing geometry. David Fried termed this gradient subspace and its related rotational wavefront equivalent the “hidden phase.”⁹ This terminology describes the fact that rotational gradients cannot be translated to a continuous sheet and thus why gradient circulations are part of a reconstructor’s null-space.

This research proposes to exploit these rotational modes within a reconstructor’s null space. Further, atmospherically distorted wavefronts for paths with a Rytov number $\sigma_\chi^2 < 0.25$ do not contain significant rotational energy. For most astronomically-relevant applications (upward looking, night-time), this condition is generally met. So, we have a gradient modal sub-space that is not induced by the atmosphere or from the AO system running in closed-loop.

2.1.1 The Core Process

A set of measured wavefront gradients is given by

$$\vec{s} = S(\vec{s}_{recon} + \vec{s}_{null,atm} + \vec{s}_{null,ref} + \vec{s}_{null,noise}), \quad (4)$$

where S is the WFS sensitivity, \vec{s}_{recon} is any reconstructable gradient energy, $\vec{s}_{null,atm}$ is any atmospheric (or noise) energy in the null-space, $\vec{s}_{null,ref}$ is any null-space energy in the gradient reference \vec{s}_{ref} , and $\vec{s}_{null,noise}$ is any noise-induced gradient energy in the null-space*. Using \mathbf{P} , the null-space content of the gradient reference can be found

$$\vec{s}_{null,ref} = \mathbf{P}\vec{s}_{ref}. \quad (5)$$

To be clear, \vec{s}_{ref} is the gradient reference measurement, generally taken with a pristine input wavefront through the WFS optics, to calibrate out any local WFS optical aberrations. Note that a time-average of Eq. 4 can be expressed as

$$\langle \vec{s} \rangle = S(\vec{s}_{null,ref}), \quad (6)$$

where $\langle \rangle$ represents a time average. Equation 6 highlights the fact that time-averaged atmospheric and noise gradient components average to zero, while the system’s null-space energy content should be fixed for a given WFS sensitivity. Therefore, one should be able to calculate the WFS sensitivity through the use of time-averaged gradients.

The proposed approach tracks the relative energy content measured in $\langle \vec{s}_{null,ref} \rangle$ to determine S , assuming the system $\vec{s}_{null,ref}$ is invariant. Using, $\vec{s}_{null,ref}$ from Eq. 5, it is then normalized,

$$\vec{\delta} = \frac{\vec{s}_{null,ref}}{\vec{s}_{null,ref} \cdot \vec{s}_{null,ref}}, \quad (7)$$

where \cdot represents an inner-product. Now, the sensitivity S of the on-sky SHWFS measurements, relative the gradient reference conditions can be simply calculated as

$$S = \langle \vec{s} \rangle \cdot \vec{\delta}. \quad (8)$$

Moreover, $\vec{\delta} \cdot \vec{s}_{recon} = 0$ because the null-space modal energy is orthogonal to the reconstructable modes.

At Starfire, an average gradient vector $\langle \vec{s} \rangle$ over four seconds is computed for the sensitivity algorithm. Because the null-space gradient component is not atmospherically or locally modified, it can reasonably be assumed that it is slowly varying.

With S computed, the reconstructed wavefront becomes

$$\hat{\phi} = \mathbf{R}(\vec{s} - S\vec{s}_{ref}), \quad (9)$$

*Instantaneous noise realizations (shot, read, and others) will couple energy into the null-space that is not representative of the signal of interest.¹⁰

which now appropriately scales down the reference gradients \vec{s}_{ref} to be on the same sensitivity level as the measured WFS gradients \vec{s} . S is recomputed every four seconds and fed back into the data pipeline, using the most recent $\langle \vec{s} \rangle$.

One might reasonably ask why not to apply S^{-1} to \vec{s} in Eq. 9, which seems mathematically equivalent to the approach proposed. In most cases $S < 1$, therefore applying S^{-1} to \vec{s} would mean applying a gain greater than one. Noting that \vec{s} contains a non-trivial noise contribution, Starfire prefers to scale down the reference \vec{s}_{ref} rather than applying a gain to a term that includes noise. However, it is noted that this approach leaves the system with a changing overall system (control law) gain, which should be compensated for with another approach that will generally account for this and the system SNR, among other things.

A non-trivial implementation detail is to ensure \mathbf{R} , used in Eq. 3 to calculate \mathbf{P} , is not modally-suppressed in any ways. Because of this, \mathbf{R} , used in Eq. 3, will likely be different than the actual reconstruction matrix used for the real-time wavefront reconstruction. Common examples of modal suppression include waffle suppression and tip/tilt removal.

The proposed algorithm makes two key assumptions:

1. The atmosphere does not impart wavefront energy to the modal space of $\mathbf{P}^{\dagger\ddagger}$.
2. Closed-loop operation does not modify wavefront energy in the modal space of \mathbf{P} .

Therefore, changes observed in the modal content of \mathbf{P} are attributed to changing WFS sensitivity of the static WFS aberrations that were recorded in the gradient reference vector. If modes are projected or left out of \mathbf{R} , they will inherently be present in \mathbf{P} and used to estimate sensitivity. In such cases, refer to the prior two assumptions to judge whether the modal content will invalidate the algorithm.

Further algorithm modifications and behaviors are noted in the following Section that followed on-sky testing.

3. ON-SKY TESTING

3.1 Initial Algorithm Characterization

The first desired trait, highlighted at the start of Sec. 2, states that the algorithm should operate seamlessly whether the system is in open or closed loop. Because the sensitivity is based on modal content that is within the null-space of the wavefront reconstructor \mathbf{R} , closing the loop does not reconstruct, nor does it impart DM commands within this sub-space. Figure 4 plots the open and closed loop power spectral densities (PSD's) for both the reconstructed and null-space modal sets. From the plot, it can be seen that closing the loop greatly reduces the PSD content in the reconstructed modes where the system has error rejection bandwidth, but little change is seen in the null-space modal content.

The second desired trait, highlighted at the start of Sec. 2, states that the algorithm should have a high repeatability with low variability. Figure 5 plots measurements of the system sensitivity over the span of two hours (including three instances closing the AO loop). A very slow drop in sensitivity was observed, which is attributed to slowly falling telescope elevation angle. Finer-scale variations are attributed to changing atmospheric conditions. Overall, we have found this metric to meet the desired repeatability and low variability quite well. Often excursions seen in the sensitivity estimations point operators to other system issues that may otherwise go unnoticed. To this end, it has been an unexpected utility. It is also noted that the changes in loop state did not result in discernable sensitivity estimate variations.

[†]Transient sensor noise induced energy does contribute to the null-space energy sensed; however, temporal averaging aims to mitigate these contributions.

[‡]Atmospherically induced branch-points do occur in paths with sufficiently high Rytov number.⁹ See Sec. 4 for discussion.

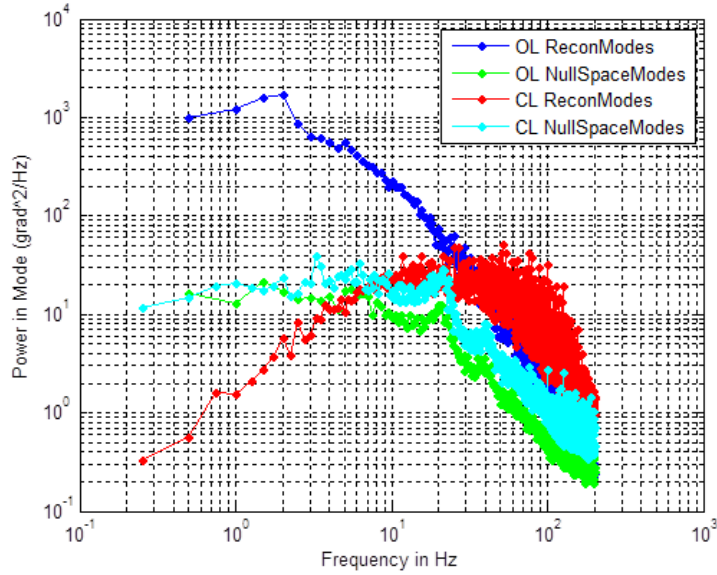


Figure 4. Power spectral densities (PSD's) for the null-space modal content and reconstructed modal content in both open and closed loop. Data sets for open and closed loop were taken sequentially to attempt to minimize atmospheric condition changes.

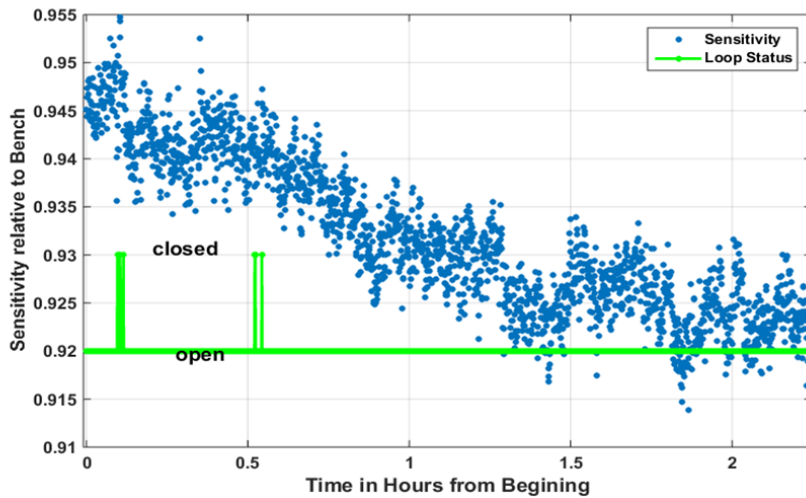


Figure 5. Estimated system sensitivity over the span of hours, in both open and closed loop. Low sensitivity variance is recorded, and long-term stability is seen. This data set was from observing a star with natural guide-star AO.

3.2 Low Signal Operation

Further on-sky testing at Starfire Optical Range revealed a number of unexpected findings that led to implementation enhancements that targeted operational robustness. Most notably, unexpected algorithmic behaviors were noted when WFS signal levels fell. As signal levels in the WFS drop, two things happen*: gradient smashing and pixel clipping. Gradient smashing is a process that sets a subaperture’s gradient to zero (after reference subtraction) if the sum of all four active pixels is less than a defined threshold. This is done to limit the impact of the lowest SNR subaperture gradient measurements to the overall wavefront estimate. Pixel clipping prevents negative pixel values after subtracting the WFS’s background frame by rounding any negative values up to zero. This is done to prevent near-zero denominators in estimating the centroid:

$$x_{cent} = \frac{(I_2 + I_4) - (I_1 + I_3)}{I_1 + I_2 + I_3 + I_4}, \quad y_{cent} = \frac{(I_3 + I_4) - (I_1 + I_2)}{I_1 + I_2 + I_3 + I_4}, \quad (10)$$

where I_1, \dots, I_4 are the pixel values within the quad-cell.

Unexpectedly, WFS sensitivity was calculated as falling once subaperture signal levels fell past a certain point—even on beacons that would otherwise provide higher sensitivities due to their size alone. At first, this behavior was thought to be an error in the algorithm. To diagnose, additional techniques were evaluated on the same data as the null-space algorithm, comparing the WFS-estimated tilt (looking at a star—not a laser beacon) to the tilt from a resolved high-speed tracking sensor (considered truth). Two distinct methods were devised for calculating sensitivities from the two signals:

$$S_1 = \frac{\sigma_{wfs}}{\sigma_{track}}, \quad (11)$$

where σ_{wfs} is the standard deviation of the WFS-derived tilt signal and σ_{track} is the standard deviation of the high-speed tracker tilt signal. This method takes the ratio of the WFS derived tilt standard deviation to the high-speed tracker standard deviation and is therefore simple to calculate. The second approach calculates the tracker-normalized correlation between the two track signals:

$$\Gamma_{wfs,track}(\Delta) = \frac{\sum_{i=1}^N \theta_{wfs}(i - \Delta)\theta_{track}(i)}{\sum_{i=1}^N \theta_{track}(i)^2}, \quad (12)$$

$$S_2 = \max(\Gamma_{wfs,track}(\Delta)), \quad (13)$$

where N is the number of data points (in time) for each sensor’s tilt series θ_{wfs} and θ_{track} and Δ is an integer frame number offset. For both approaches, the WFS and tracker should be clocked at the same sample rate, and a small time offset (much smaller than the overall data collect duration) between the two collects is tolerable. It should also be noted that the approaches for estimating sensitivity S_1 and S_2 are not generally useful in the same sense that the null-space approach is. For example, if running the AO system from a laser beacon, the WFS tilt measurements cannot be correlated to the tracker signal.

Figure 6(a) plots sensitivity calculations for an on-sky satellite pass that started fairly dim and got progressively brighter. All three sensitivity calculations show a similar drop in sensitivity at low light levels to the WFS. To this point, we began to investigate why—now trusting that the drop was a real effect.

An obvious effect as light levels drop, is higher relative noise levels (i.e. lower SNR), but the four-second averaging, mentioned in Sec. 2.1.1, should mostly eliminate noise impacts. Gradient smashing and pixel clipping also occur more as light levels fall, and these became the most suspected effects at play. A simulation of all three sensitivity estimation schemes was coded, and results can be seen in Fig. 6(b). The same general trends were found as in the on-sky data analysis[§].

[§]In Fig. 6, the bump seen at mid-light levels in the S_1 standard deviation approach is attributed to the impact of increasing sensor noise. As light levels continue to fall, the sensor noise is attenuated by gradient smashing and pixel clipping—bringing it roughly inline with the other approaches. At the lowest light levels, there is not enough full-aperture signal to correlate to the true track signal, leading the S_2 correlation approach to fall to nearly zero, while leaving S_1 slightly elevated.

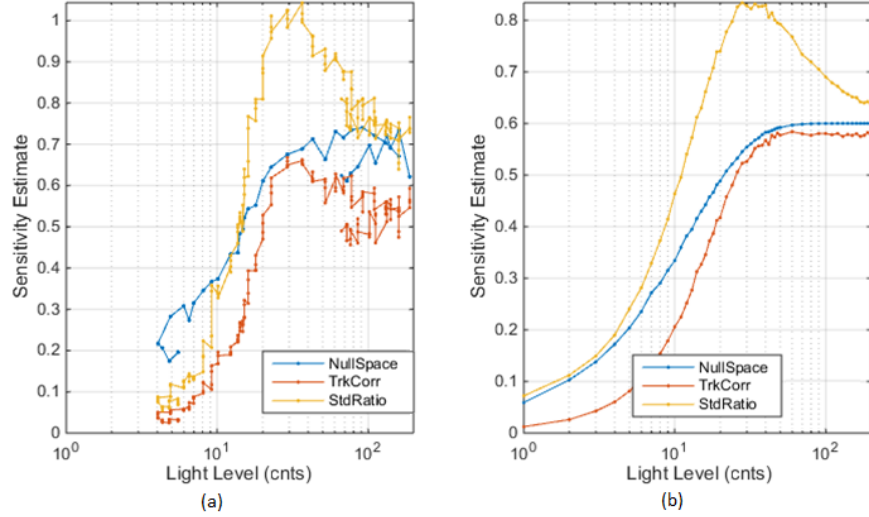


Figure 6. WFS sensitivity as a function of subaperture light levels is shown. The left subplot calculates this for an observed low-Earth orbit satellite pass that started dim and got brighter, while the right subplot is simulated data. Three methods for calculating the sensitivity are shown in each: Null-space refers to S , StdRatio refers to S_1 , and TrkCorr refers to S_2 .

It is true that both gradient smashing and pixel clipping lower the calculated WFS sensitivity, and the natural follow-up question is whether one should compensate for either or both of these effects through Eq. 9. To answer this, consider each individually. In assigning applicable subaperture gradients to zero, gradient smashing certainly lowers the average WFS energy content (in both reconstructable and null-space subsets), which, if included, would certainly lower the WFS sensitivity estimate from the algorithm. However, the sensitivity is meant to scale the reference gradients to the sensitivity of the remaining unsmashed gradient measurements. Therefore, it would not be desirable to base the sensitivity calculation on smashed gradient sets. Pixel clipping is less clear to analyze. By rounding negative pixel values to zero, gradient magnitudes are reduced[¶]. Again, if included, the algorithm will find this lowers the estimated WFS sensitivity. If the effected subaperture gradients were not smashed, these clipped gradients will be compared to the reference. To this end, it is recommended to include pixel clipping in estimating the sensitivity for this algorithm.

The proposed approach modifies how $\langle s \rangle$ is calculated, which then feeds Eq. 8. Instead of averaging at the gradient level, Starfire has adopted averaging at the WFS pixel level (i.e. I_1, I_2, I_3 , and I_4 for each subaperture). In this approach, clipping is included but gradient smashing is not (gradient smashing occurs after the pixel values are converted into gradient-space for a given frame). With this modification, Fig. 7 plots Starfire’s observed sensitivities over the span of one year. The same general trend is seen at low light levels as seen in Fig. 6, which drives an interesting and unexpected final conclusion: Light levels drive a WFS’s sensitivity!

3.3 Other Sensitivity Drivers

A number of other sensitivity drivers have been noted that were unexpected. These are described below.

3.3.1 Changing Sky Background Radiance

In a similar, yet subtly different manner to what has been discussed in Sec. 3.2, changing sky background light levels were observed to impact WFS sensitivity. In twilight hours, the whole-sky light level is changing on a minute-by-minute timeframe. During the daytime or near sunrise/sunset, the sky has a highly spatially variant background radiance. In these cases, the sensitivity changes due to scaling of the denominator of the centroid calculations in Eq. 10. This occurs because the WFS background frame (which is subtracted from every measurement frame) is only valid for the time and/or sky location that it was taken.

[¶]The denominator in Eq. 10 is increased and the numerators are reduced in magnitude.




Figure 7. A one-year look at recorded sensitivities plotted against subaperture light levels on Starfire’s telescope. Outlying features are generally caused by unconventional experimentation or telescope use.

An example case would be running the AO system on a natural star. The operator takes a background collect on the sky near the star, then uses the starlight as the AO beacon. The averaged background frame is subtracted from the WFS signal to (ideally) leave only the beacon signal. If this happens just after sunset (in twilight), the background collect will be relatively brighter than as time goes on. This means too much background will be subtracted from the sensed starlight, shrinking the centroid denominators. This is a rare case where Starfire has observed sensitivity estimates (correctly) greater than one. The opposite behavior will be true in twilight just before sunrise, as the sky brightness gradually increases in time. The same is true in twilight and daytime from a spatial sky aspect, in addition to the temporal aspect.

This algorithm compensates for the impact of changing (spatially or temporally) sky radiance levels, at least until a fresh background collect can be taken. It is notable that this effect has nothing to do with the beacon size.

3.3.2 Changing Turbulence or Elevation Angle

Also observed, is a dependence on the telescope elevation angle and temporally varying turbulence. This effect was noted in Sec. 3.1, but is worth repeating and further explanation. Telescope elevation angle ψ changes the volume of atmosphere that the observation path must look through, through a $\csc \psi$ relationship. This, as well as changing atmospheric conditions, effectively will change the optical system’s point-spread function (PSF) width, regardless of whether open or closed-loop AO is used. Changing the PSF width effectively changes the WFS’s sensitivity, also regardless of whether the AO beacon is inherently extended or not. Figure 8 shows this dependence for a single satellite pass that used laser guidestar AO. A distinct linear dependence of the sensitivity was found to the cosecant of the elevation angle. As noted previously, this impact is also found in natural guidestar AO data, with Fig. 5 showing this slower drift for a two-hour star observation.

3.3.3 Local System Optics

In transitioning this algorithm to a different AO system, erratic behavior for this algorithm was noted. After troubleshooting, it was determined that the system’s DM was introducing a non-trivial amount of energy into \mathbf{P} ’s modal space. Our WFS reference gradient set \vec{s}_{ref} , only encompassing the WFS optics alone, did not include this DM null-space energy. When transitioning to the sky or test bench source, the reported sensitivity increased dramatically. A similar behavior is seen, although to a lesser extent, at Starfire’s 3.5m system once the DM

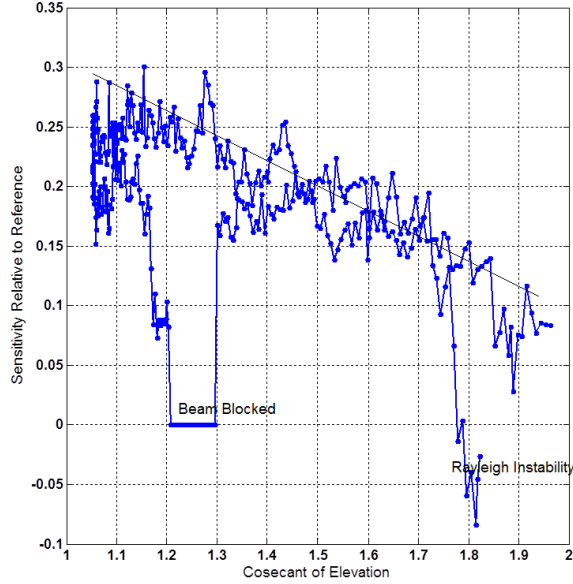


Figure 8. Sensitivity of Starfire’s laser guidestar AO system through a low-Earth orbit satellite pass. The independent axis is the cosecant of the telescope elevation angle, which is the linear scale factor of how much atmosphere the optical path looks through.

optics are included in the path, but the behavior was not investigated further until moving to this secondary system.

Ultimately, the sensitivity estimator’s reference gradient set \vec{s}_{ref} must be measured in a way that captures the primary null-space energy contributors within the local system. This new reference measurement is termed “the sensitivity baseline,” and the reference source is set up to capture (as much as possible) every optic up to the WFS. The sensitivity baseline is used only for the purposes of deriving $\vec{s}_{null,ref}$. A separate WFS reference is still taken in the traditional manner that captures (as much as possible) the WFS-specific optical train, which is used within the AO control system to remove any WFS-specific aberrations from the corrections.

At first, this required change was perceived as an annoyance, however, the DM’s additional energy in the null-space creates a more prominent $\vec{s}_{null,ref}$ which enhances the algorithm’s effective SNR in operation. Prior implementations induced static “racetrack”¹¹ in the gradients by intentionally rotating the SH sensor a very small amount (the smallest rotation achievable). This mode is in the null-space, and was thought to be required since it was assumed the system nominally did not contain any energy in the null-space. Once it was determined that sufficient energy in the null-space existed without the sensor rotation, the WFS rotation was removed. It is not clear whether all DM’s or optical systems will contain such null-space features, although any amount of sensor rotation will couple gradient energy to the null-space.

3.3.4 Modal Reconstruction Considerations

As mentioned in Sec. 2.1.1, it is essential to calculate \mathbf{P} using a full modal basis set in \mathbf{R}^{\parallel} . In the case of AO systems utilizing modal reconstruction, where one has the option of only reconstructing a subset of M modes for an K channel DM, the following relationship generally holds: $M \leq K$. For the purposes of this sensitivity estimation algorithm, a version of \mathbf{R} should be calculated such that $M = K$, even if a different \mathbf{R} is used in the real system. The reasoning behind this links back to the two algorithmic assumptions listed at the end of Sec. 2.1.1. If modes are left out of \mathbf{R} that is used to calculate \mathbf{P} , and if those modes have non-trivial energy in the observed atmosphere, the time-variant atmospheric energy in these left out modes will (incorrectly) drive the algorithm’s estimate of the WFS sensitivity. Generally speaking, only the rotational gradient sub-space

¹¹This does not imply that \mathbf{R} must be full-rank—it generally will not be.

(sometimes termed the slope discrepancy space¹¹) is desired in P . Using this definition, one can alternatively construct P , rather than Eq. 3.

4. CONCLUSIONS

A real-time algorithm to estimate a quad-cell Shack-Hartmann wavefront sensor's sensitivity is presented. The primary use for this algorithm is in preventing "reference bleed-through," where the in-lab gradient reference set is recorded at a different WFS sensitivity than on-sky measurements. The algorithm has been in use at the Starfire Optical Range for over one year. A separate overall system (control-loop) gain control algorithm that takes into account system SNR and other operational conditions is recommended to complement this algorithm.

The proposed algorithm has high-SNR, is repeatable, shows low-variability, works whether the system is in open or closed loop, and does not require dithering or other system modifications. The following operating conditions were found to impact sensitivity:

1. AO beacon extent relative a point source
2. Turbulence conditions and telescope elevation angle, which impact point-spread function width
3. Dim AO beacons
4. Changing sky radiance levels, whether spatially across the sky or temporally in twilight hours
5. Local system optical aberrations, if unaccounted for

Research related to this algorithm continues at Starfire Optical Range in two basic research areas. The first is investigating applicability towards the pyramid wavefront sensor, which is fundamentally quite similar to a quad-cell Shack-Hartmann wavefront sensor. Secondly, evaluation of this algorithm towards applications in moderate to high Rytov turbulence regimes is important for understanding daytime observation capabilities. In these stronger turbulence regimes, the atmosphere excites energy in the reconstructor null-space modes, which violates a core assumption of this approach. The long temporal averaging, currently used to mitigate other noise sources, may average out such atmospheric null-space energy, in a similar sense to Véran's approach.¹ Characterization of the extent of this impact will inform operational utility and potential alternative solutions for use in daytime applications.

ACKNOWLEDGMENTS

This research was funded over several years by the Air Force Research Laboratory. The authors would like to thank the technical critique, algorithm implementation, test support, and contract administration from the entire Starfire Optical Range team—both Government and contractor.

The views expressed in this document are those of the authors and do not reflect the official policy or position of the United States Air Force, the Department of Defense, or the United States Government.

REFERENCES

1. J.-P. Véran and G. Herriot, "Centroid gain compensation in Shack-Hartmann adaptive optics systems with natural or laser guide star," *J. Opt. Soc. Am. A* **17**, pp. 1430–1439, Aug 2000.
2. M. A. van Dam, "Measuring the centroid gain of a Shack-Hartmann quad-cell wavefront sensor by using slope discrepancy," *J. Opt. Soc. Am. A* **22**, pp. 1509–1514, Aug 2005.
3. M. van Dam, "Keck's WFS sensitivity compensator." Email correspondence, Oct 2017.
4. M. Olikier and D. Roskey, "New techniques for the live update of gain tables in NGS and LGS WFS operation," *Proc. SPIE* **7736**, p. 15, 2010.
5. K. Shtyrkova, M. D. Olikier, K. P. Vitayadom, D. W. Oesch, D. J. Sanchez, P. R. Kelly, C. M. Tewksbury-Christle, and J. C. Smith, "Experimental demonstration of real time gradient gain correction for sodium beacon laser guide star," Tech. Rep. ADA531341, Air Force Research Laboratory, June 2010.

6. R. Tyson, *Principles of Adaptive Optics*, Series in Optics and Optoelectronics, Taylor & Francis, 2010.
7. M. J. Steinbock, "Implementation of branch-point-tolerant wavefront reconstructor for strong turbulence compensation," Master's thesis, Air Force Institute of Technology, June 2012.
8. G. Golub and C. Van Loan, *Matrix Computations*, Johns Hopkins Studies in the Mathematical Sciences, Johns Hopkins University Press, 1996.
9. D. L. Fried, "Branch point problem in adaptive optics," *J. Opt. Soc. Am. A* **15**, pp. 2759–2768, Oct 1998.
10. D. L. Fried, "Adaptive optics wave function reconstruction and phase unwrapping when branch points are present," *Opt. Commun.* **200**(16), pp. 43 – 72, 2001.
11. T. J. Brennan, "Anatomy of the slope discrepancy structure function: characterization of turbulence," *Proc SPIE* **5087**, p. 12, 2003.