Multiple aperture approach to wavefront prediction for adaptive-optic applications

Matthew R. Kemnetz* and Stanislav Gordeyev †

University of Notre Dame, Notre Dame, Indiana, 46446, USA

A multiple aperture approach for the correction of aero-optical distortions is presented. Measurements were conducted in Notre Dame's Tri-Sonic facility at $M = 0.2$ at $40kHz$ with an $Ap/\delta \approx 6$. In this work we investigated the convective nature of the wavefront in order to verify previous work as well as test the limits of a simply convective assumption. We present two models for wavefront correction, a simply convective model and a linear convolution approach. We found that for a separation of $\Delta/\delta \approx 3-4$ a simply convective model of wavefront propagation was not sufficient to characterize the downstream wavefront from upstream data. POD analysis showed that although the large scale structure of the wavefront primarily convects, the small scale structure of the wavefront quickly loses correlation as the separation increases. The convecting-only assumption was then replaced with the time-delayed linear convolution, where information over several sequential time points in the upstream location were used to predict the downstream wavefront. If the first 40 POD modes were used in predicting the downstream wavefront, the reduction of the overall wavefront error was found to be -8 dB for the small separation of $\Delta/\delta = 1$ and nearly -5 dB for the largest separation of $\Delta/\delta = 4$.

Nomenclature

δ	Boundary layer thickness
δ^*	Displacement thickness
θ	Momentum thickness
λ	Wavelength
ρ	Fluid density
OPD	Optical Path Difference
OPL	Optical Path Length
POD	proper Orthogonal Decomposition
$SHWFS$	Shack-Hartmann Wavefront Sensor
St_δ	Strouhal number based on δ
TWT	Tri-Sonic Wind Tunnel
U_c	Convective velocity

I. Introduction

A major hurdle for airborne high energy laser systems is the so-called aero-optic problem. When an otherwise collimated, coherent beam of light encounters a turbulent flow field that includes density fluctuations (compressible turbulent boundary layer; for example), its optical wavefront becomes aberrated. As planar wavefronts propagate through these unsteady density distributions, they become distorted and these

*Graduate Student, Department of Aerospace and Mechanical Engineering, Hessert Laboratory for Aerospace Research, Notre Dame, IN, 46556, AIAA Student Member.

†Research Associate Professor, Department of Aerospace and Mechanical Engineering, Hessert Laboratory for Aerospace Research, Notre Dame, IN, 46556, AIAA Associate Fellow.

distortions can be accurately measured by various wavefront sensors.[1,2] In the far field, these aberrations manifest themselves as a retarded on target intensity.

One way to mitigate the effects of aero-optic distortions is to utilize an adaptive-optic system. These systems use a wavefront sensor in conjunction with a deformable mirror. The conjugate of the measured wavefront can be imprinted on the deformable mirror thereby correcting for aberrations caused by the high speed turbulent flow. Unfortunately, the aero-optical regime is characterized by high speed flow phenomena such as shocks, turbulent boundary layers, and shear layers occurring at characteristic frequencies on the order of 1kHz. Previous work by Burns et al.[3] has demonstrated that for these frequencies, a major limiting factor for operation is the cumulative latency inherent in the adaptive optic system. Time spent on digital calculations, physical sensor response, analog-to-digital conversion of sensor signals, electronic amplifiers, and electromechanical response of deformable mirrors all factor into the overall latency of the adaptive optic system.[3] This latency problem was approached by Burns et al. by developing a predictive controller. In this manner measurements were made over the aperture and the controller predicted forward in time. The control action is then performed based on the predicted flow characteristics. The latency tolerance is built into the prediction step.

Fundamentally, in the root of the latency problem lays the fact that for most AO systems the aero-optical distortions are commonly treated as unknown a priori, which requires an unrealistically-fast closed-loop adaptive-optics system. We propose that an alternative approach for a latency tolerant system would be to use multiple apertures and eventually, a physics-based model. A second aperture could be placed upstream of the lasing aperture and measurements would take place at this upstream location. By the time measurements finished, the flow structures would have convected over the main lasing aperture. Instead of measuring at one location and predicting forward in time, multiple apertures allow the use an upstream measurement as the prediction for the flow over the aperture.

Previous work has shown that naive convection using Taylors Hypothesis isnt quite adequate to predict the wavefront on the downstream aperture.[4] In that experiment, wavefronts were collected after the laser had passed through two statistically independent boundary layers on opposing sides of the wind tunnel. The second boundary layer has no effect on statistical quantities but for the type of direct comparison methods done in that previous work and this experiment here, the influence of the second boundary layer should be eliminated. The influence of the second boundary layer has been eliminated in this work through the use of wall heating.[1] More information about how the wall was heated can be found in Sec. [A].

In this work we will further investigate the limits of the Taylor Frozen Flow Hypothesis as it applies to adaptive-optic correction using multiple apertures. We use a POD approach to decompose wavefronts in modes and corresponding coefficients. By assuming the boundary layer is statistically the same over both apertures, we can postulate that the POD modes are the same over both apertures and the only difference is in POD coefficients. The problem is reduced to predicting the downstream POD coefficients from the time series of the upstream ones. In this work we propose a linear convolution approach for the prediction of the POD coefficients. We believe that this model can be adapted and used in a diverse set of flow regimes such as boundary layers, shear layers, wake regions, and cavities.

II. Experimental Procedure

Wavefront data used in this work were collected in the Hessert Tri-sonic Wind Tunnel (TWT) at the University of Notre Dame. The TWT is a continuous flow indraft wind tunnel with an inlet contraction ratio of 150:1. The test section used has a cross section of 10 cm 9.9 cm. The tunnel test section is constructed of Plexiglas. The test section was 143cm in length from the end of the inlet to the diffuser. Optical access windows are installed on the upper and lower walls of the test section from 110 cm to 122 cm. Freestream velocity was measured using a static pressure port just upstream of the optical window and was held constant at M = 0.20 for the duration of the measurements. The boundary layer thickness, δ, at the measurement location was found to be approximately 16.0 mm, with δ^* to be about 2.65 mm and $\theta = 2.00$ mm. Based on these values we have $Re_\theta = 44,030$. Detailed description of the tunnel is provided in Smith.[5] A schematic of the experimental facility can be seen in Figure [1].

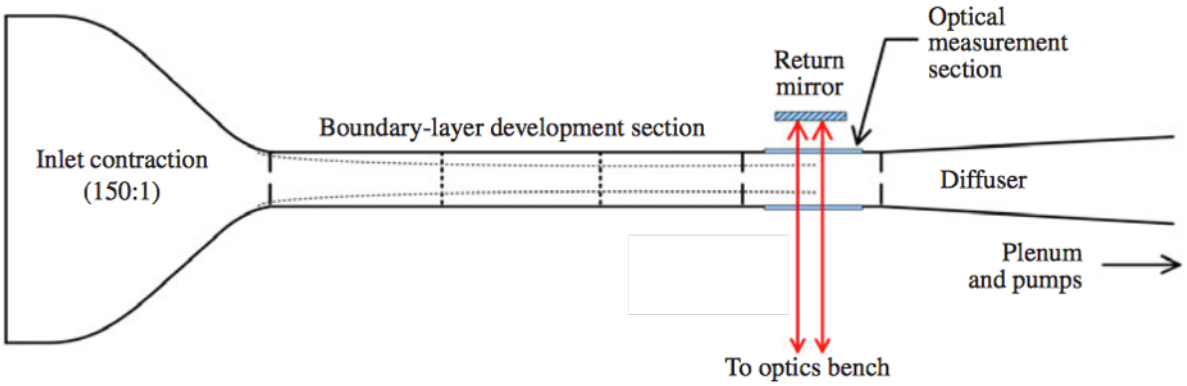

Figure 1. Schematic of the Notre Dame Tri-sonic Wind Tunnel Facility.

For this experiment the laser source was expanded to 1 inch through the collimator and then again to form a five inch beam that was sent through the test section. The beam was sent through the test section and reflected off a return mirror in the so-called double-path set-up. The returning beam was split off using a beam splitter and reimaged into a high speed wavefront sensor. Full 2-D wavefronts, resolved in both the streamwise (x) and the spanwise (z) directions, were collected using a high-speed Shack-Hartmann Wavefront Sensor (SHWFS). The collected wavefronts had a spatial resolution of 33 x 65. The sampling frequency was $43,300\,Hz$ and the wavefronts had a streamwise extent of 6.37δ and a spanwise extent of 3.60δ. A total of $45,000$ 2D wavefronts were collected.

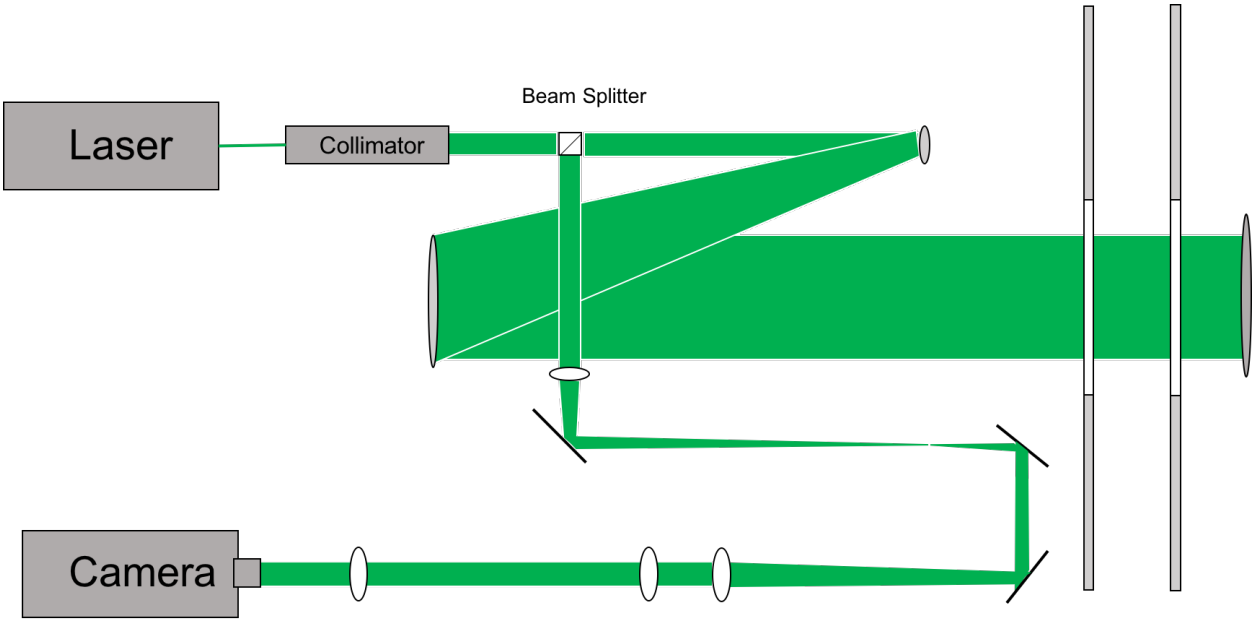

Figure 2. Schematic of the wavefront measurements using a 2-D Shack-Hartmann wavefront sensor.

A. Heated Wall

It has been shown[6] that the OPD_{rms} of a wavefront for a turbulent boundary layer with wall heating scales

$$OPD_{rms} \propto \left(M^2 + D\frac{\Delta T}{T_\infty} \right) \tag{1}$$

where $D \approx 2.2$ is an experimentally determined constant. For this experiment the freestream Mach number was $M = 0.2$ and the boundary layer was heated and held constant such that $\Delta T = 25°K$. This resulted in an OPD_{rms} contribution five time greater for the heated boundary layer of the test section than for the unheated boundary layer. In this manner the unheated boundary layer was optically suppressed.

B. Data Extraction

As mentioned in Sec. [II] one large, 5 inch, wavefront was collected. The data was then spatially sampled into four different streamwise separations. The sampled subapertures were chosen such that the separation from every point of the "Upstream" ($\Delta/\delta = 0$) data was 1δ, 2δ, 3δ, and 4δ respectively. After sampling, the data was Tip, Tilt, and Piston removed again to simulate data taken at varying streamwise locations with smaller apertures. A diagram of how the large aperture wavefront was sampled can be seen in Figure [3].

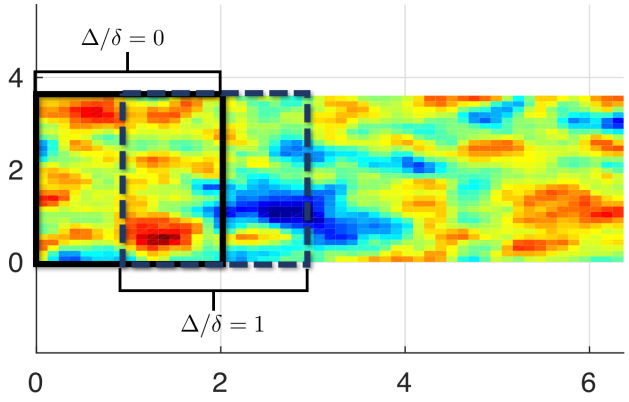

Figure 3. An example of the large aperture sampling to produce data for the $\Delta/\delta = 0$ and $\Delta/\delta = 1$ cases.

III. Results

A. Residual WF Error

To investigate different methods for the prediction of the downstream wavefronts define the global error, Σ.

$$\Sigma = \frac{\overline{\left\langle \left[W^{DN}(x,z,t) - \tilde{W}^{DN}(x,z,t) \right]^2 \right\rangle}}{W_{rms}^2} \tag{2}$$

Where $\tilde{W}^{DN}(x,z,t)$ is some estimate of the downstream instantaneous wavefront.

If we assume that the wavefronts simply convect, the downstream wavefront is just an earlier wavefront from the upstream aperture. A simple convective based estimate for $\tilde{W^{DN}}$ is

$$\tilde{W}^{DN}(x,z,t) = WF_{UP}(t-\tau) \tag{3}$$

where τ is the convective time delay, related to the convective speed, U_c, as $\tau = \Delta/U_c$. Using the global error equation defined in Eq. [2] we investigate the residual WF error plotted for varying streamwise separations and convective speeds. The residual WF error is plotted as a function of convective velocity in Figure [4].

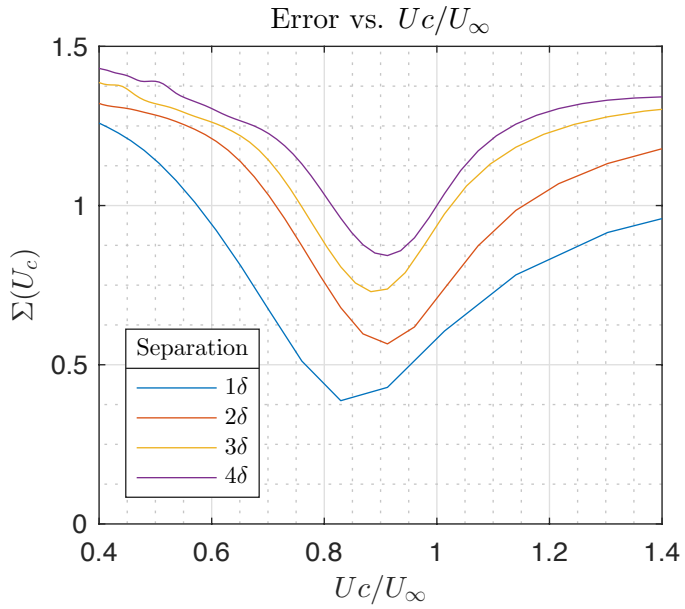

Figure 4. Plot of residual WF error vs. Convective velocity for various streamwise separations

The error measured using Eq. [2] quickly increases as a function of separation. Even for a modest separation of $\Delta/\delta = 3$ the error reaches a surprisingly high value of $\Sigma \approx 0.75$. For this reason we conducted POD analysis to better understand the large scale motion of the flow.

B. POD Analysis

The wavefronts were decomposed into global modes using an adaptation of the POD technique[7]

$$W(x,z,t) = \sum_n a_n(t)\phi_n(x,z). \tag{4}$$

POD modes were computed using the JPOD method.[8] The first 20 modes can be seen in Figure [5].

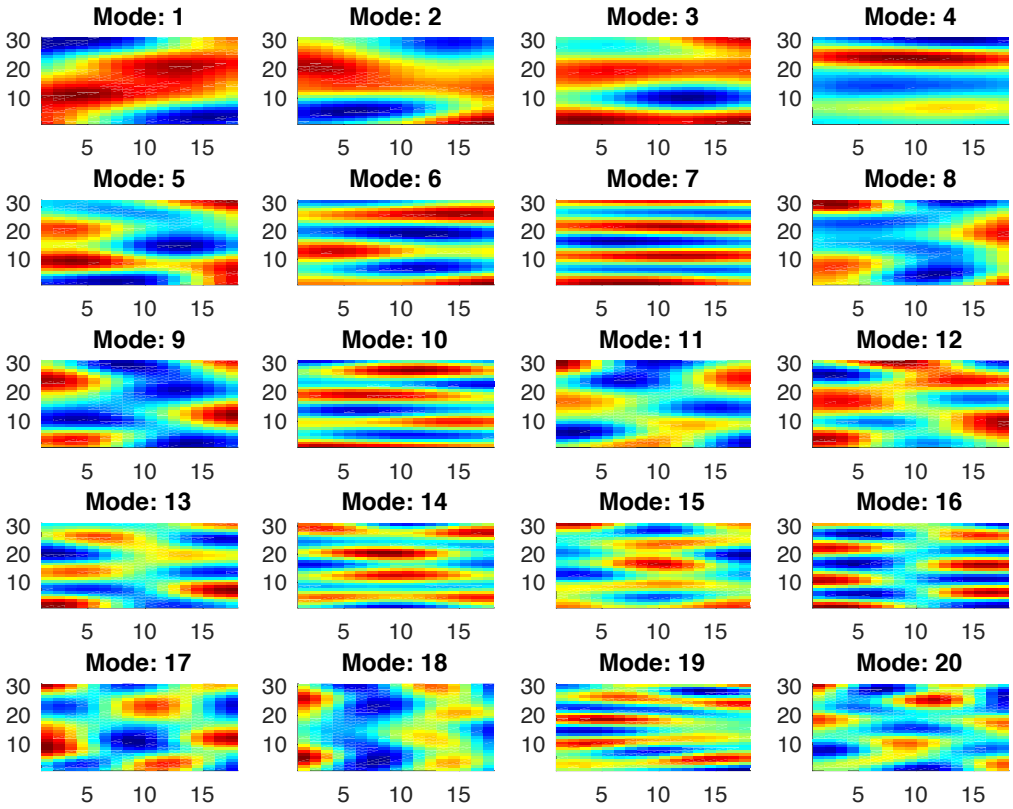

Figure 5. First 20 POD modes computed using the JPOD technique.[8]

POD modes, by the construction, are the same in both the upstream and the downstream locations; the only difference are in the temporal coefficients. If POD coefficients in the downstream location are somehow estimated using information from the upstream aperture, it will provide an estimate of the instantaneous wavefront evolution in the downstream location. These estimated wavefronts can be used as an input into an adaptive optics system to correct for boundary-layer aero-optical distortions.

A plot of the normalized eigenvalues and the cumulative eigenvalues can be seen in Figure [6].

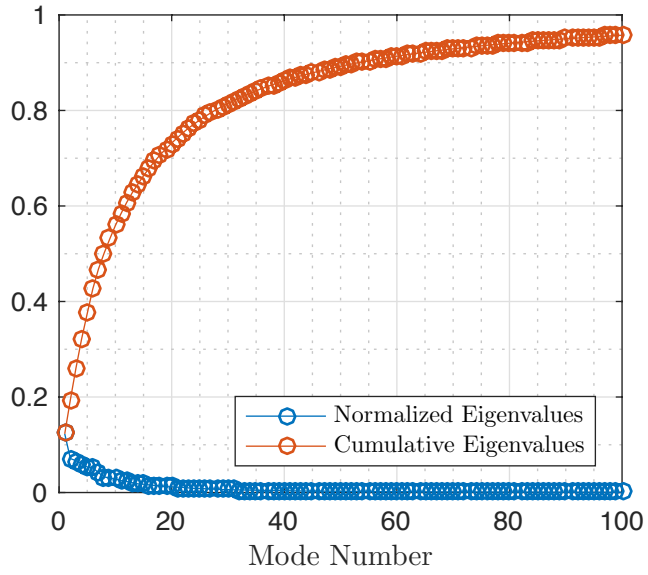

Figure 6. Normalized Eigenvalues vs. POD Mode number.

It is clear from Figure [6] that 80% of the energy of the wavefronts is contained in the first 30 modes. To better understand the dynamics of the flow and the importance of large scale structure we can now introduce a new description of error. We define the quantity relative error as

$$\epsilon(\tau, n) = \frac{\overline{\langle a_n^{UP}(t-\tau) - a_n^{DN}(t)\rangle^2}}{\lambda_n} \tag{5}$$

where $a_n(t)$ are the temporal coefficients defined in Eq. [4] and λ_n is the total energy. If one included only the first n modes in the wavefront prediction,

$$\tilde{WF}^{DN}(x, z, t-\tau) = \sum_{i=1}^{N} \tilde{a}_n^{DN}(t)\phi_n(x, z) = \sum_{i=1}^{N} a_n^{UP}(t-\tau)\phi_n(x, z) \tag{6}$$

the relative error, ϵ, and the global error due to the first N modes, $\Sigma(N)$ can be related,

$$\Sigma(N) = \frac{\sum_{n=1}^{N} \epsilon_n \lambda_n + \sum_{n=N+1}^{\infty} \lambda_n}{\sum_n \lambda_n} \tag{7}$$

This error shows the combined effect of not including higher-order POD modes into the estimation and the error in predicting the included low-order modes.

It can easily be shown that in the limit of $n \to N$, where N is the total number of POD modes, Eq. [7] is equivalent to Eq. [2]. We can now compare the error using Eqs. [5] and [7] and compare the results. The relative error was computed with the first 20 POD modes. The first 20 POD modes correspond to the large scale, low frequency content in a wavefront.

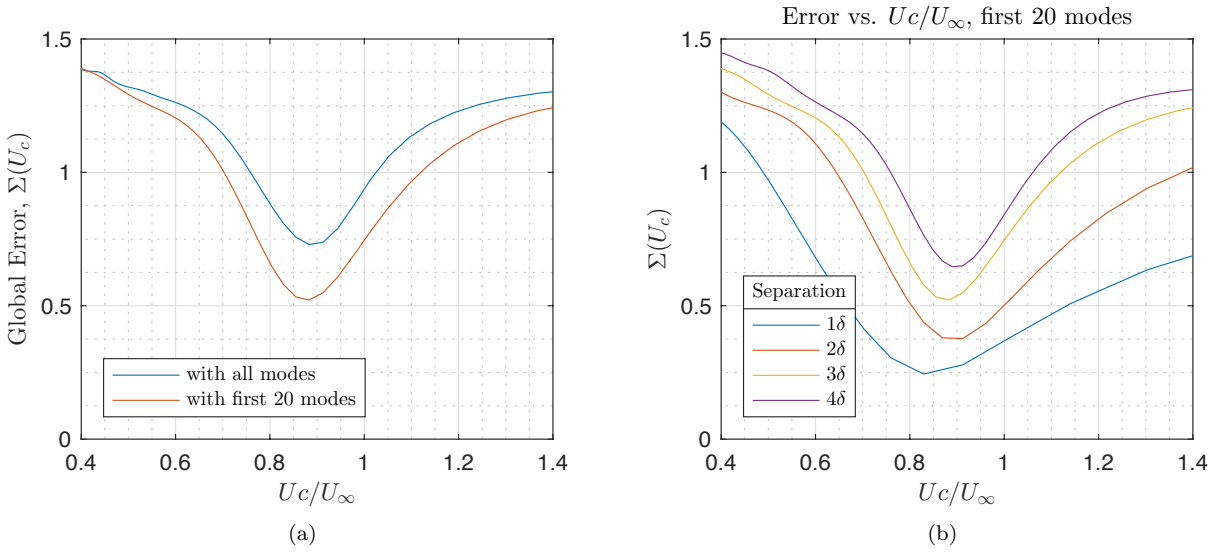

Figure 7. (a) Plot of residual WF error vs. Convective velocity for all POD modes and just the first 20 POD modes. (b) Global error using first 20 POD mode for varying separations.

There is a dramatic reduction in the residual wavefront error when just the first 20 POD modes are used. Despite the fact that the higher order modes (Mode # $>$ 20) account for just 20% of energy in the wavefront, they contribute significantly to the residual wavefront error. In Figure [7(b)] the global error is plotted for various streamwise separations using only the first 20 modes. Comparing with Figure [4] there is a dramatic reduction for all of the streamwise separtions. This implies that it is the lower order spatial modes that mostly convect while the higher order modes decay before reaching the downstream aperture. In the higher order modes there is simply no meaningful correlation between the upstream and downstream apertures.

A plot of the convective velocity for various mode numbers can be seen in Figure[8(a)]. The convective velocity for each mode number was computed by determining the minimum error between the time shifted upstream wavefront and the downstream wavefront.

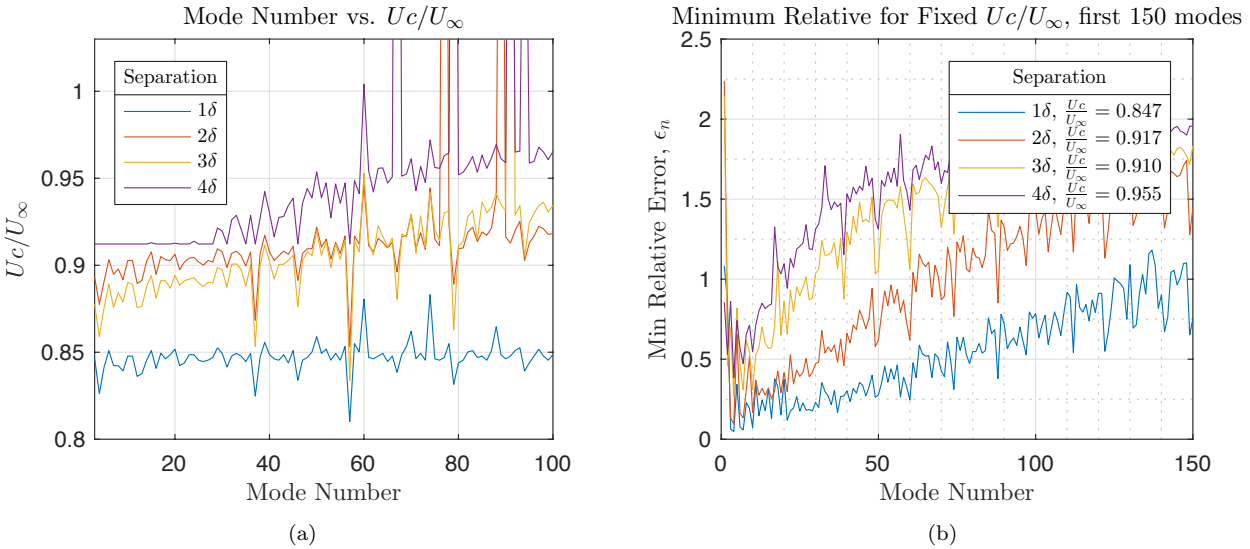

Figure 8. (a) Normalized convective velocity vs. Mode number. (b) Error vs. Mode Number.

In general, the computed convective velocities for each mode are constant. For the $\Delta/\delta = 3$ and 4 separations the convective velocity grows rather linearly. We believe this is mostly due to measurment error

caused by a drop in correlation between the upstream wavefront and the downstream wavefront for these larger separations.

Now that we know the convective velocity for each mode, we can look at how the error varies with mode number. We can fix the convective velocity and compute the relative error, ϵ_n through Eq. [5]. In Figure [8(b)] one can clearly see the trend of significantly reduced relative error for the lower order POD modes. This implies that it is the lower order spatial modes that mostly convect while the higher order modes decay before reaching the downstream aperture. The large error seen in the lowest order modes is likely an artifact of mechanical vibration. It is these lowest order modes that are most susceptible to contamination through vibration. In addition, for increasing separation, correlation in the higher order modes is lost. For $\Delta/\delta = 1$ correlation is maintained all the way through mode 150 while for $\Delta/\delta = 4$ correlation is largely lost by mode 30. For the larger separations the higher order modes corresponding to the fine grain structure of the flow have changed by the time they reach the downstream aperture. This provides further evidence that it is just the large scale structure of the flow that primarily convects.

C. Linear Convolution Approach

The coefficient estimation, given in Eq. [7] has several issues. Firstly, it relies on the upstream coefficient only at one fixed time delay. Secondly, the amplitude of the downstream coefficient is unchanged. In other words, it assumes a convecting-only wavefront, implying perfectly-correlated delayed wavefronts between the two apertures.

POD coefficients in the downstream location can be estimated from a linear time-delayed convolution with the upstream POD coefficients (for discrete time delays),

$$\tilde{a}_n^{DN}(t) = \sum_k \sum_{\tau \geq 0} A_{nk}(\tau) a_k^{UP}(t-\tau) \tag{8}$$

A-matrix provides different weights for different time delays and allows to use several sequential temporal coefficients in the upstream location to compute the modal coefficient in the downstream location.

In general, the $A_{nk}(\tau)$ matrix can be fairly large, which would require extra memory storage and might increase the time of computing the predicted coefficients. To make it more computational-efficient, it can be simplified by including only diagonal terms,

$$\tilde{a}_n^{DN}(t) = \sum_{\tau \geq 0} A_n(\tau) a_k^{UP}(t-\tau) \tag{9}$$

Simalarly to Eq. [5], the individual modal error of the estimation can be computed as,

$$\epsilon_n = \frac{\overline{\left[a_n^{DN}(t) - \tilde{a}_n^{DN}(t)\right]^2}}{\lambda_n} \tag{10}$$

Finally, the total time-averaged relative error of estimating the downstream wavefront using only first N POD modes can be computed using Eq. [7]. This error shows the combined effect of not including higher-order POD modes into the estimation and the error in predicting the low-order modes.

A-matrices in Eq. [8] were computed by Least Square Estimation technique for various separations between the upstream and the downstream apertures using the first 20,000 temporal points. The estimation was performed on the next 20,000 temporal points. The individual errors for each mode, computed using the full linear convolution, Eq. [8], is shown in Figure [9] for different separations. For the smallest separation of $\Delta/\delta = 1$, the relative errors are less than 20% for the first 100 modes. For the separation of $\Delta/\delta = 3$, the first 20 modes are predicted better than 20% of the relative error for the largest separation of $\Delta/\delta = 4$ the first 20 modes are predicted better than 30% of the relative error.

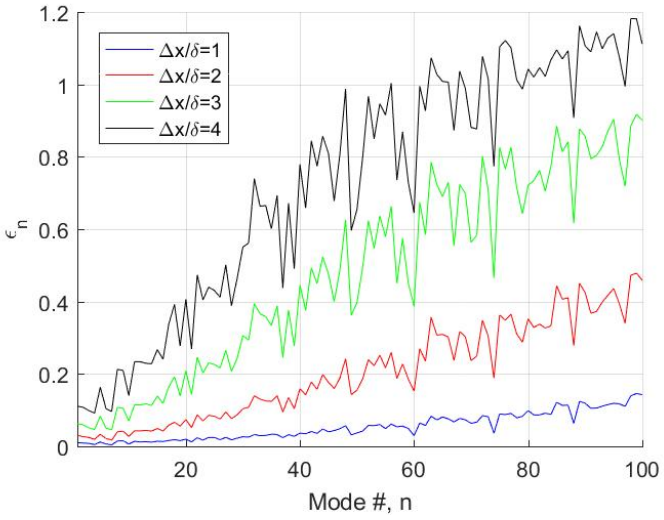

Figure 9. Relative errors computed using the full linear convolution, Eq. [8] for different separations.

To see the overall effect of using only a small number of low-order modes to estimate the downstream wavefronts, the total relative error is plotted in Figure [10] for all separations. An ideal correction using only the first N modes is also presented in Figure [10]. It shows that using only first 20 modes, it is possible to reduce the wavefront error by -5 dB or by a factor of 3.2 for small separations of $\Delta/\delta = 1$ and 2. Even for the largest separation $\Delta/\delta = 4$ the wavefront error can be reduced by -4 db. Inclusion of more modes will further improve the wavefront prediction, but in expense of additional computational cost, as the size of A-matrix in Eq. [8]is proportional to a square of the included POD modes. Also, for large separations of $\Delta/\delta = 3$ and 4 at some point the inclusion of more modes will not significantly improve the correction. For large separations, the correlations between the upstream and the downstream location gets weaker, and the correlation-based total error will approach a constant value. This value obviously is a function of the separation distance and is equal to approximately -7 dB for $\Delta/\delta = 3$and -5 dB for $\Delta/\delta = 4$, see Figure [10]. Thus, for large separations, there is no reason to include many modes, as there is no significant improvement in using more than 40 or so modes.

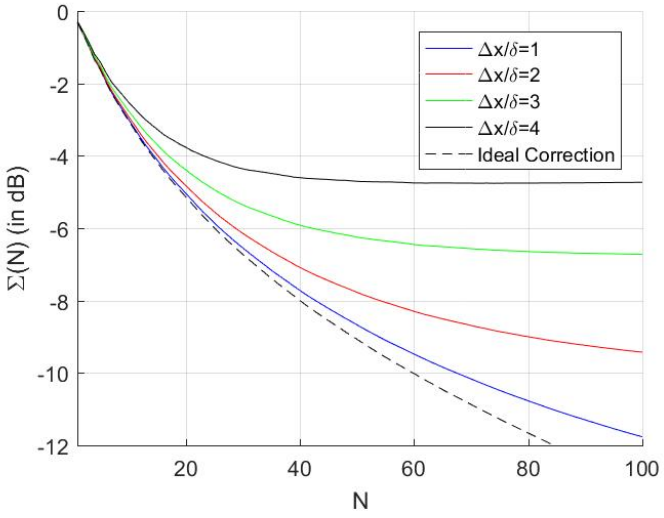

Figure 10. Total errors if only the first N modes are used to compensate the downstream wavefront, computed using the full linear convolution, Eq. [8], for different separations.

To further improve the computational effectiveness of the proposed prediction method, only diagonal

terms, 9, were kept to predict the downstream coefficients. New A-matrices were also computed by Least Square Estimation using 20,000 temporal points and the convolution was applied to the next 20,000 points. The relative modal errors in this case are shown in Figure [11(a)]. While obviously not as good as the inclusion of the possible correlation between the modes, the diagonal-only approach still provides less than 40% for the first 20 modes for separations up to $\Delta/\delta = 3$.

The related total error is presented in Figure [11(a)]. $\Delta/\delta = 1$ and 2, inclusion of the first 40 modes would still results in more than -5 db in the wavefront reduction. For $\Delta/\delta = 3$, an inclusion of 40 POD modes to predict the downstream wavefronts would give -3.9 dB, or by a factor of 2.5, reduction of OPDrms. For $\Delta/\delta = 4$, the corresponding reduction in OPDrms would be -2.9 dB, or by a factor of 2.0.

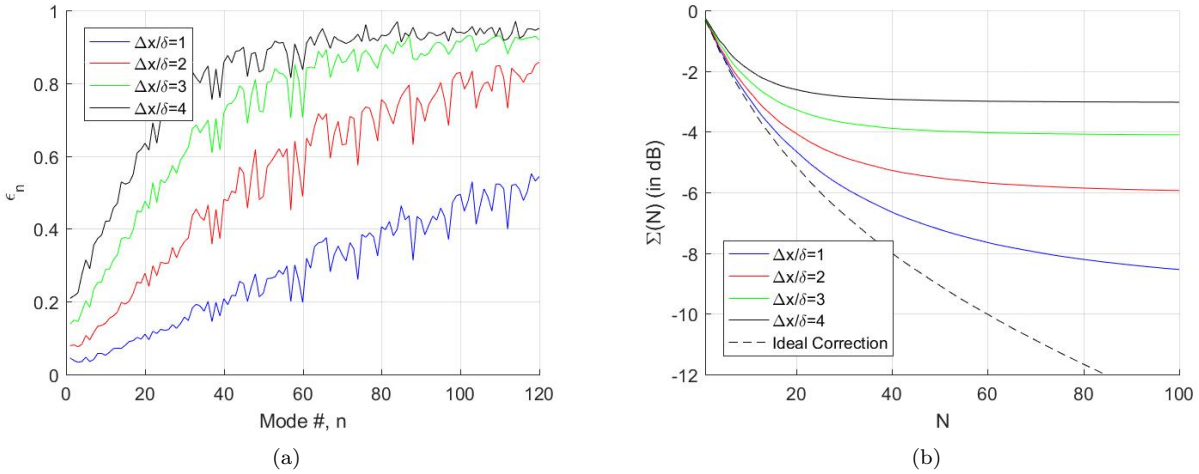

Figure 11. (a) Relative errors computed using the diagonal-only linear convolution, Eq. [9] for different separations. (b) Total errors if only the first N modes are used to compensate the downstream wavefront, computed using the diagonal only linear convolution, Eq. [9], for different separations.

Diagonal A-matrices for selected modes are shown in Figure [12] for different separations. As mentioned before, A-matrices are essentially weights of different time-delayed upstream coefficients to compute the downstream coefficient. From Figure [12] it can be seen that A-matrix is essentially non-zero for a range of delays, with a peak at the time delay of Δ/UC. Also, the amplitude of A-matrix decreases for large-number POD modes and large separations, as the coefficients between the two apertures become less correlated. Thus, the predicted modal coefficients for large mode numbers will be suppressed, further improving the overall prediction of the downstream wavefront.

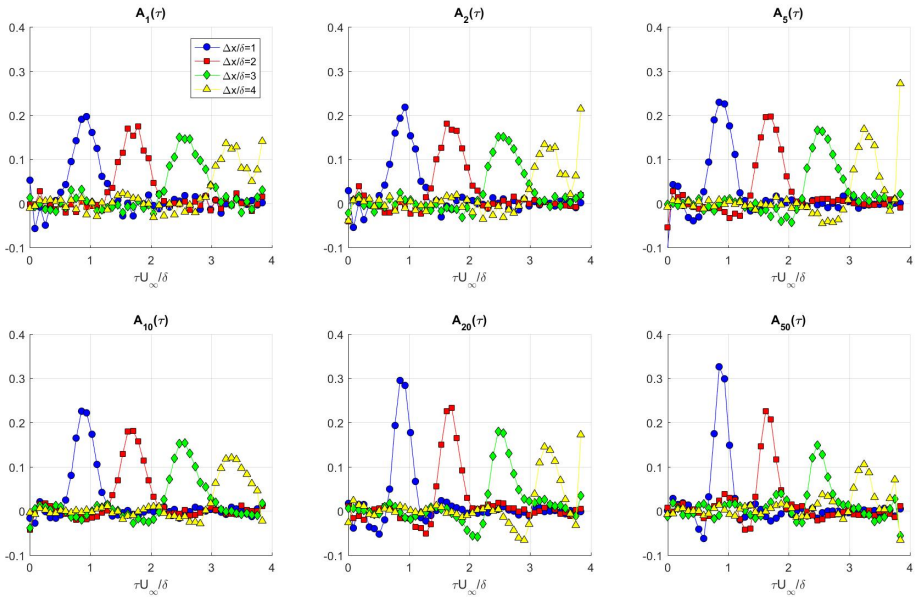

Figure 12. Diagonal A-matrices from Eq. [9] for various time delays and separation for selected modes.

D. Aero-Optical Jitter

In the typical analysis of wavefront data the three lowest order spatial modes are removed from the data. These are called the Tip, Tilt, and Piston modes. These modes are usually corrupted by mechanical vibration. In previous work,[9] we have developed a technique for stitching wavefront frames together reintroducing Tip, Tilt, and Piston back into the wavefront. Time is traded for space and one long wavefront $W(x, y)$ is generated from a collection of wavefront images $W(x, y, t)$. From this stitching analysis we were able to recover the Tip, Tilt, and Piston lost during wavefront processing. It is trivial to show the beam jitter is simply a function of the Tip and Tilt of the wavefront. In this manner the jitter caused by the aero-optics can be extracted from the stitching technique.

In the experiment conducted for this paper one large aperture of $Ap/\delta \approx 6$ was used. This large aperture data can be used to verify the stitching method as Tip and Tilt information from various smaller apertures can be directly extracted from it. The large aperture wavefront we re-apertured to form new datasets with varying aperture sizes. An example of this re-aperturing is shown in Figure[13]

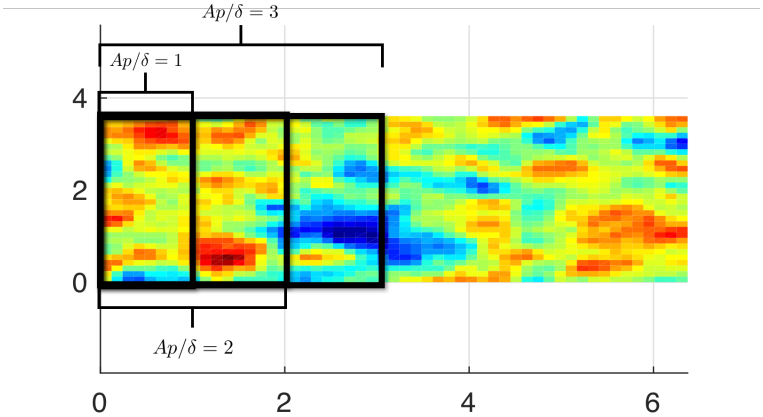

Figure 13. An example of the sampling done to create varying aperture sizes from the large $Ap/\delta \approx 6$ data.

overall unsteady Tip and Tilt (or jitter) was extracted from six smaller apertures $Ap/\delta = 1, 2, 3, 4, 5$, and 6. A plot of the aero-optical jitter, θ_{rms}, as a function of aperture size can be seen in Figure [14(a)]. The jitter was normalized as $\theta_{rms} \propto (\rho/\rho_{SL})(M^2 + D(\Delta T/T))$. Plotted alongside the experimental data is the predicted aero-optical jitter from the stitching method.[9]

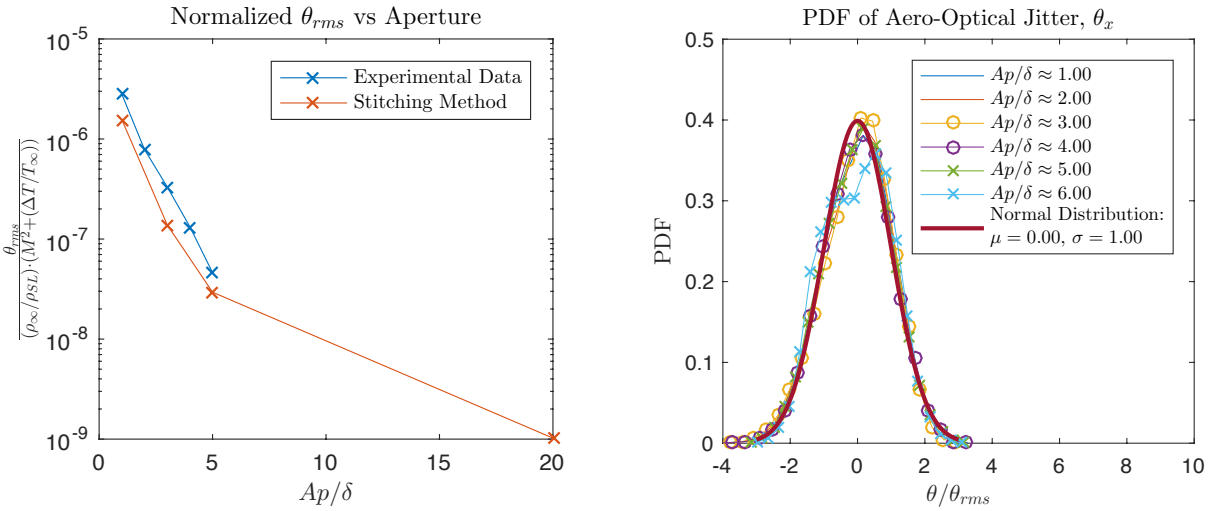

Figure 14. (a) Normalized θ_{rms} for both the experimental data the stitching method (b) PDF of the aero-optical jitter for the experimental data.

As one can see in Figure [14(a)], the stitching method did an good job in predicting the aero-optical jitter for various apertures thus verifying our stitching method.

Plotted in Figure [14(b)] arethe PDFs of the jitter for various aperture sizes. After normalizing by the root-mean-square of the jitter, the data collapses on a normal distribution of zero mean and a standard deviation of 1. This is very useful as, combined with Figure [14(a)], we can compute the jitter statistics for a given aperture size. This could help in the selection of a fast steering mirror for adaptive optic applications.

IV. Conclusion

In this work, measurements of aero-optical distortions due to a turbulent boundary layer are presented and discussed. Measurements were conducted in Notre Dame's Tri-Sonic Wind Tunnel at $M = 0.2$ with a wall temperature heated to $25°C$ above T_∞.

In this work we investigated the convective nature of the wavefront in order to verify previous work as well as test the limits of a simply convective assumption. Wall heating was implemented to suppress the influence of the second boundary layer. We found that for a separation of $\Delta/\delta \approx 3 - 4$ a simply convective model of wavefront propagation was not sufficient to characterize the downstream wavefront from upstream data. POD analysis showed that although the large scale structure of the wavefront primarily convects, the small scale structure of the wavefront quickly loses correlation as the separation increases. The uncorrelated small scale structures are enough to significantly increase the measured error for separations beyond $\Delta/\delta = 3$.

The convecting-only assumption was then replaced with the time-delayed linear convolution, where information over several sequential time points in the upstream location were used to predict the downstream wavefront. The weighting matrices were computed using Least Square Estimation technique. As this approach uses additional correlation information at various time delays, it clearly outperforms the simple convection-only approach. If the first 40 POD modes were used in predicting the downstream wavefront, the reduction of the overall wavefront error was found to be -8 dB for the small separation of $\Delta/\delta = 1$ and nearly -5 dB for the largest separation of $\Delta/\delta = 4$.

A simplified, more computationally-effective, version of the time-delayed convolution method with diagonal terms only was also tested. While not as quite effective as the full version of the convolution approach, it still showed -7dB reduction of the wavefront error for the small separation of $\Delta/\delta = 1$ and about -3dB

reduction for the separation of $\Delta/\delta = 4$, if the first 40 modes are used. So, the linear convolution-based method was found to be fairly effective in predicting the wavefronts several boundary layer thicknesses downstream using the information from the upstream location. As the linear convolution is essentially a matrix multiplication of temporal coefficients from the nearest past, it automatically solves the latency problem and it can straightforwardly adapted to be used with current adaptive optics systems.

We were also able to extract the aero-optical jitter from our data for various aperture sizes and compare with predictions made by the stitching method. We found that the stitching method, proposed earlier, did a good job in predicting the actual aero-optical jitter, thus verifying the stitching method. Finally, we computed the PDF of the aero-optical jitter and found that after proper normalization all of the data collapsed into the normal distribution. This implies that for a given aperture size the statistics of the aero-optical jitter can be determined. This data could be useful in the selection of a fast steering mirror for adaptive-optic applications.

References

[1]Gordeyev, S., Cress, J. A., Smith, A., and Jumper, E. J., "Aero-optical measurements in a subsonic, turbulent boundary layer with non-adiabatic walls," *Physics of Fluids*, Vol. 27, No. 4, 2015.

[2]Gordeyev, S., Smith, A. E., Cress, J. A., and Jumper, E. J., "Experimental studies of aero-optical properties of subsonic turbulent boundary layers," *Journal of Fluid Mechanics*, Vol. 740, 2014, pp. 214–253.

[3]Burns, W. R., Jumper, E. J., and Gordeyev, S., "A Latency-Tolerant Architecture for Airborne Adaptive Optic Systems," *53th AIAA Aerospace Sciences Meeting*, 2015, AIAA Paper 2015-0679.

[4]Kemnetz, M. R., Burns, W. R., and Gordeyev, S., "Multiple Aperture Approach for the Study of Large-Scale Boundary-Layer Structures," *47th AIAA Plasmadynamics and Lasers Conference*, 2016, AIAA Paper 2016-3530.

[5]Smith, A. E., *Evaluation of passive boundary layer flow control techniques for aero-optic mitigation*, Ph.D. thesis, University of Notre Dame, 2015, Doctoral Dissertation.

[6]Cress, J. A., Gordeyev, S., and Jumper, E. J., "Aero-Optical Measurements in a Heated, Subsonic, Turbulent Boundary Layer," *48th Aerospace Science Meeting and Exhibit*, 2010, AIAA Paper 2010-0434.

[7]Holmes, P., *Turbulence, coherent structures, dynamical systems and symmetry*, Cambridge University Press, Cambridge, UK; New York, 2012, ID: 756045732.

[8]Gordeyev, S., Lucca, N. D., Jumper, E. J., Hird, K., Juliano, T. J., Gregory, J. W., Thordahl, J., and Wittich, D. J., "Comparison of Unsteady Pressure Fields on Turrets with Different Surface Features using Pressure Sensitive Paint," *Experiments in Fluids*, Vol. 55, 2014, pp. 1661.

[9]Kemnetz, M. R. and Gordeyev, S., "Optical investigation of large-scale boundary-layer structures," *54th AIAA Aerospace Sciences Meeting*, 2016, AIAA Paper 2016-1460.