

A Latency-Tolerant Architecture for Airborne Adaptive Optic Systems

Robert Burns¹, Eric Jumper², Stanislav Gordeyev³
University of Notre Dame, Notre Dame, IN, 46556

Abstract

A new data-driven method for latency compensation in adaptive optics systems is developed and evaluated in this paper. Conventional adaptive optic controllers typically assume a single timestep of latency in the feedback loop, which is an appropriate assumption for low-frequency applications such as atmospheric optics compensation. However, the controller frequencies needed for aero-optic applications can approach the tens of kilohertz. Cumulative latency in the feedback loop originating from digital communication links, sensors, processing, etc. can exceed one timestep and thus significantly reduce the performance of adaptive optic controllers, even though open-loop frequencies of individual components are sufficiently fast for the application. Our method uses proper orthogonal decomposition as a basis for wavefront model reduction and an artificial neural network to predict the evolution of the associated modal coefficients over a short temporal horizon equivalent to the amount of latency present in the feedback loop. The method is capable of significantly augmenting the performance of conventional adaptive optics controllers. This algorithm has been evaluated in closed-loop simulation with disturbance data gathered from the Airborne Aero-Optics Laboratory. Over a 5-step prediction window, the new controller could reduce OPD_{rms} in the worst cases by over 35%. Over a single timestep window, mitigations of greater than 55% are realized.

Nomenclature

α	= viewing angle (rad)
β	= modified elevation angle (rad) or integrator gain (-)
ϵ	= vector of wavefront prediction residuals
λ	= laser wavelength (m)
ρ	= air density (kg/m^3)
Σ	= diagonal matrix of singular values
Φ	= matrix of POD modes (-)
φ	= eigenmode vector (-)
A	= state-space matrix describing linearized wavefront dynamics
Az	= azimuth (rad)
a	= modal coefficients
C	= compensator transfer function (-)
D	= aperture diameter (m)
d	= single-aperture aero-optic disturbance (m)
E	= root-mean-square residual error
El	= elevation (rad)
f	= disturbance frequency (Hz)

¹ Graduate Student, Department of Aerospace and Mechanical Engineering, Hessert Laboratory for Aerospace Research, Notre Dame, IN 46556, AIAA Student Member.

² Professor, Department of Aerospace and Mechanical Engineering, Fitzpatrick Hall of Engineering, Notre Dame, IN 46556, AIAA Fellow.

³ Research Associate Professor, Department of Aerospace and Mechanical Engineering, Hessert Laboratory for Aerospace Research, Notre Dame, IN 46556, AIAA Associate Fellow.

f	=	vector-valued nonlinear prediction function
G	=	plant model transfer function (-)
g	=	function describing flow evolution
h	=	observer function of special flow properties
H	=	number of hidden neurons (-) or sensor transfer function (-)
I	=	on-target beam intensity (W/m^2)
I_0	=	unaberrated beam intensity (W/m^2)
K_{GD}	=	Gladstone Dale constant (m^3/kg)
M	=	Mach number (-) or embedding dimension (-)
N_d	=	number of timesteps of delay in the feedback loop (-)
n	=	index of refraction (-)
OPD	=	Optical Path Difference (m)
OPL	=	Optical Path Length (m)
S	=	Strehl ratio (-)
St	=	Strouhal number (-)
U	=	coefficient matrix
V	=	freestream velocity (m/s)
v	=	wavefront vector (m)
W	=	wavefront vector (m)
x	=	vector of coefficients (-)
y	=	single-aperture aero-optic residual output (m)
z	=	vector of flow properties

I. Introduction

A major obstacle in the implementation of airborne laser systems is the aero-optical problem¹. The aero-optical problem arises from changing density in the medium of a laser beam's propagation in the near-field of the optical aperture: as density changes, the index of refraction of the medium changes, the laser beam will be distorted, and the airborne laser system will suffer a significant degradation in performance. A very common and useful configuration for laser pointing mechanisms is the classical turret, as they give airborne lasers a very large field of regard and enable the airborne system to illuminate a target with minimal limitations to its flight path. The aero-optical problem manifests itself on turrets in compressible flow as shown² in Figure 1: in turbulent boundary layers, pressure and temperature fluctuations contribute to significant density fluctuations as the boundary layer thickens; flow separation over the top of the turret results in large separation regions that roll up into "horn" and secondary vortices; necklace vortices form at the base of the turrets; and at sufficiently high Mach numbers (as low as $M = 0.55$), shocks will begin to appear at the top of the turret. All of these effects have a substantial aberrating effect on laser beams.

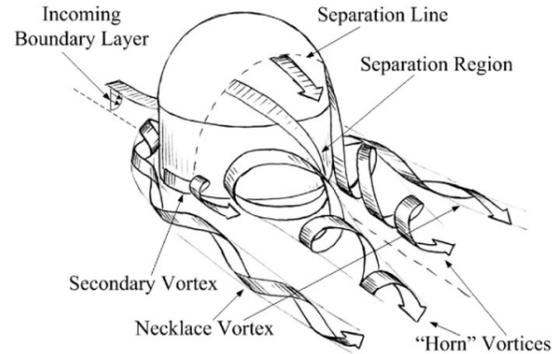


Figure 1: Optically-aberrating flow topology².

The aero-optics problem is depicted schematically in Figure 2. An initially planar wavefront is propagated outward from the interior of the turret in the direction of a target. The beam then encounters a separated, turbulent flow, which unevenly retards the progress of the beam by slowing its propagation through the air. In the near field, the previously planar wavefront has been imprinted with an aberration. In the far field, this aberrated wavefront will manifest itself as a laser beam of reduced intensity, scattered over a wider area. The index of refraction is related to the density of the medium through the Gladstone-Dale relationship³, $n(x, y, z) = 1 + K_{GD}(\lambda)\rho(x, y, z)$, where n is the index of refraction, λ is the wavelength, and K_{GD} is a constant. This relationship can be used to compute the exact aberration along a path by first computing the optical path length from some point in the aperture, $OPL(x, y, t) = \int_0^{z_1} n(x, y, z, t) dz$, and then subtracting the aperture-averaged mean to compute the optical path difference,

$$OPD(x, y, t) = OPL(x, y, t) - \langle OPL(x, y, t) \rangle \quad [1]$$

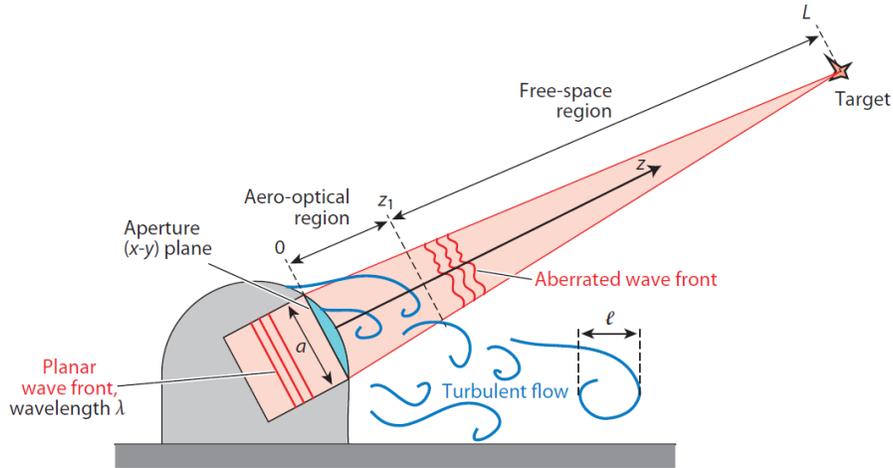


Figure 2: Schematic of the aero-optics problem¹.

The root-mean-square of the optical path difference is a very convenient metric for estimating the performance of aero-optic systems since it enables the estimation of the ratio of the Strehl ratio: peak on-target beam irradiance to peak beam irradiance of the corresponding unaberrated beam. For instantaneous aberrations following a Gaussian spatial distribution, this can be done through the Maréchal large aperture approximation^{4,5,6}, shown in Eq. 2. This relationship shows the heavy dependence of airborne laser system performance on wavefront aberration.

$$S(t) = \frac{I(t)}{I_0(t)} \approx e^{-\left(\frac{2\pi OPD_{rms}(t)}{\lambda}\right)^2} \quad [2]$$

Clearly, it is desirable to maximize the Strehl ratio by mitigating aero-optical aberrations. It is conventional to quantify reductions in aero-optical aberration using the following definition,

$$dB = 10 \log_{10} \frac{OPD_{rms,corrected}}{OPD_{rms,uncorrected}} \quad [3]$$

Since laser wavelength is application-dependent, it is useful to consider the impact of wavefront aberration mitigation in non-dimensional terms by rearranging the equation in terms of a corrected and uncorrected Strehl ratio. This relationship is plotted in Figure 3 for several amounts of mitigation. The reader will note that a -3dB (50%) correction would improve an uncorrected Strehl ratio of 0.2 to 0.67, for instance. It can be concluded that even seemingly-modest improvements to OPD_{rms} can lead to substantial increases in the Strehl ratio.

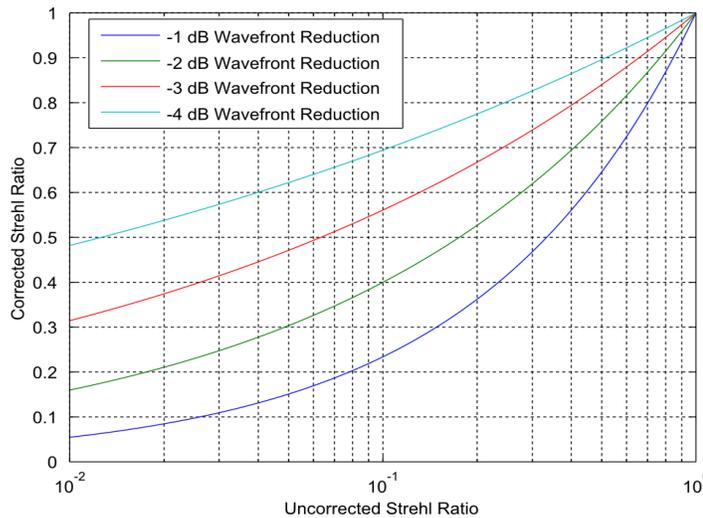


Figure 3: Effect of wavefront mitigation on Strehl Ratio.

A very large body of work has been performed to characterize aero-optical effects on the Airborne Aero-Optics Laboratory⁷ (AAOL), which is shown in Figure 4. AAOL consists of a pointing and a tracking aircraft. The pointing aircraft shoots a diverging beam at the tracking aircraft's turret, so that the aero-optics of the pointing aircraft manifest primarily as tip/tilt aberrations. The aero-optical aberrations of the tracking turret are imprinted on the laser beam near the aperture and then measured by a Shack-Hartmann wavefront sensor in the tracking aircraft. An important discovery in the AAOL program is that for some applications, the aero-optic disturbances in the turbulent boundary layer prior to separation may be tolerable; however, the strong aberrations present in the separated wake region aft of the turret are likely very deleterious to the performance of laser-based systems.



Figure 4: Picture of the Airborne Aero-Optics Laboratory (AAOL). The pointing aircraft is depicted on the left and the tracking aircraft on the right.⁷

An important discovery in the AAOL program is that aero-optical aberrations tend to collapse onto a curve that is a function of two quantities: a lookback angle, α , and a modified elevation angle, β . These two quantities can be expressed⁷ in terms of azimuth and elevation as shown in Figure 5, and expressed mathematically as follows,

$$\alpha = \cos^{-1} [\cos(Az) \cos(EI)] \quad [4]$$

$$\beta = \tan^{-1} \left[\frac{\tan(EI)}{\sin(Az)} \right] \quad [5]$$

Aero-optical aberrations must be mitigated using some technique for many applications involving either communications or directed energy. There are two primary ways of achieving this: flow control and adaptive optics. Flow control devices are generally very application-dependent. Transonic flow control devices consisting of porous screens have been developed by Gordeyev et al⁸ that relocate shocks away from the turret aperture, forcing the sharp optical discontinuity introduced by the shock to form upstream and subsequently dissipate into less optically-degrading structures. Smith and Gordeyev have studied devices⁹ that break up large, optically-active eddies in the boundary layer into smaller eddies that introduce less optical distortion. Vukasinovic et al¹⁰ used high frequency synthetic jets to reduce the optical aberrations present in the wake of a hemisphere. Later, Vukasinovic et al¹¹ used a hybrid of passive and active flow control: by using a forward-protruding plate to decouple the turret wake from the necklace vortices, the effectiveness of the oscillating jets upstream of the optical aperture to control the turret wake was enhanced. Vukasinovic et al¹² have also tested control jets in the transonic regime to control separation both upstream and downstream of the shock to reduce the sharp velocity and density gradients present in the shear layer. This type of flow control was studied in more detail by Gissen et al¹³ using Particle Image Velocimetry (PIV) and Schlieren imagery.

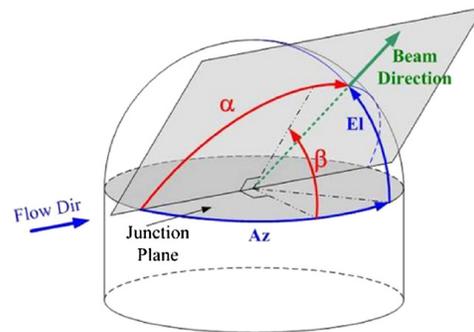


Figure 5: Definition of α and β in terms of azimuth and elevation⁷

Another means of aero-optic mitigation is to use adaptive optics. These systems imprint the conjugate of the currently-present aero-optic aberration onto a deformable mirror, and then send the pre-corrected laser beam out from the turret. When the corrected beam then encounters the aero-optic aberration, it emerges from the flow as nearly planar. These techniques are currently employed to great effect on ground-based telescopes to correct for atmospheric optical effects. However, these aberrations occur on the order of 1-10 Hz and thus the computational and system requirements are not extremely demanding for currently-available technology. However, aero-optical phenomena consisting of shocks, turbulent boundary layers, shear layers, etc. occur on the order of kilohertz on full-scale turrets at realistic Mach numbers. In particular, the dominant frequency content of separated flow on a turret in the wake region typically occurs near a Strouhal ratio of about 1:

$$St = \frac{fD}{V}, \quad [6]$$

where f is the frequency of the disturbance in hertz, D is the turret diameter, and V is the freestream velocity. For the Airborne Aero-Optics Laboratory, the disturbance frequencies observed in the wake region were concentrated near 1 kHz. Thus, a real-time adaptive optic system must be able to perform corrections at a minimum of 10kHz, assuming 10 discrete corrections per disturbance cycle. This is a demanding computational requirement given that wavefronts must be represented by matrices that are at minimum 10x10, and sometimes as large as 30x30. Despite this, existing technology is rapidly approaching a capability level where these real-time corrections are in fact feasible. The major obstacle will be cumulative latency in the system. Time spent on digital calculations, physical sensor response, analog-to-digital conversion of sensor signals, electronic amplifiers, and electromechanical response of deformable mirrors will all contribute to a significant degree of delay in a high-speed adaptive optic system.

There are a number of methods that have been designed to solve this latency problem. Burns et al demonstrated the feasibility of a system¹⁴ that relates pressure signals surrounding a beam aperture and optical wavefronts that would allow fast computation of wavefronts from real-time pressure signals, thus significantly reducing the latency problem. However, this was only demonstrated for transonic flows where pressure fluctuations are very strong and easy to distinguish from noise. Nightingale et al¹⁵ have developed a hybrid flow control/adaptive optics approach that uses a phased-locked loop to synchronize the flow and a controller algorithm, enabling high-speed correction of periodic flows. Faghihi et al¹⁶ have developed a state-space model and filter for the prediction of POD mode coefficients that could also be used in a predictive controller. Goorskey et al¹⁷ have proposed a promising controller based on the Dynamic Mode Decomposition of Schmid¹⁸ that is capable of compensating for one timestep of delay. Their algorithm was demonstrated to show a significant improvement (approximately 25% on average, and in many cases more) in controller performance.

As we will now show, even a small number (<5) of timesteps of latency in a system is capable of substantially reducing performance of a closed-loop controller. Consider the simplified feedback control model of Figure 6. We assume that each subaperture can be treated independently of all others. While this is certainly not a good assumption in reality, we believe it is sufficient to obtain a better understanding of the impact of feedback latency. In this simplified model, G represents the transfer function of the deformable mirror, H represents the transfer function of the sensor, C represents the transfer function of the controller, d is the aero-optical disturbance, y is the residual disturbance not corrected by the controller, and r is the reference input (usually 0).

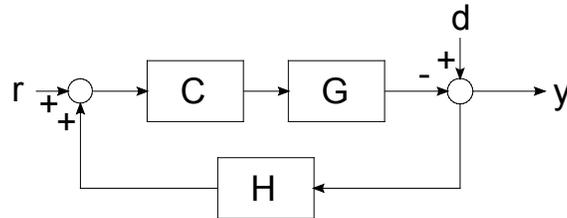


Figure 6: Simplified single-subaperture single-input single-output (SISO) adaptive optics model.

A common assumption in the analysis of adaptive optics systems is that the rise time of the deformable mirror and amplifiers is much smaller than the equivalent frequencies associated with disturbances. Additionally, a PI (proportional-integral) controller is often used as a starting point for AO control design. Finally, for this model we assume a sensor model that perfectly measures the residual disturbance, but with some number of timesteps of latency denoted by N_d . The transfer functions for these components in the z -plane¹⁹ are listed as follows:

$$G(z) = 1 \quad [7]$$

$$C(z) = \beta \frac{z}{z-1} \quad [8]$$

$$H(z) = \frac{1}{z^{N_d}} \quad [9]$$

We again stress that a real AO system is far more complex than this; but this should serve as a demonstration for the importance of feedback latency. We chose to model this in the discrete domain and then scale it to a notional controller operating at 25kHz to match AAOL using the Tustin transformation¹⁹, $z = e^{sT} \approx (1+sT/2)/(1-sT/2)$, where T is the sampling time of the system and s is the complex Laplace-domain variable. The sensitivity function of the AO controller in the z -plane is given as,

$$\frac{y}{d}(z) = \frac{z-1}{z-1 + \beta z^{1-N_d}}. \quad [10]$$

Next, we plot this this sensitivity function for several amounts of system latency as well as a gain value selected to achieve a 60 degree phase margin – a common performance target²⁰ in controls engineering. For our notional SISO system operating at 25 kHz, the sensitivity function for 1 to 5 timesteps of latency is shown in Figure 7. While a single timestep of latency certainly limits the disturbance rejection, additional latency dramatically reduces the disturbance rejection bandwidth. This latency may originate from the sensor and wavefront reconstruction as modeled here, or it may show up in communication links or even as “apparent” latency in the deformable mirror and amplifiers. In either case, it can be concluded that as the latency increases, the feedback gain must decrease. T

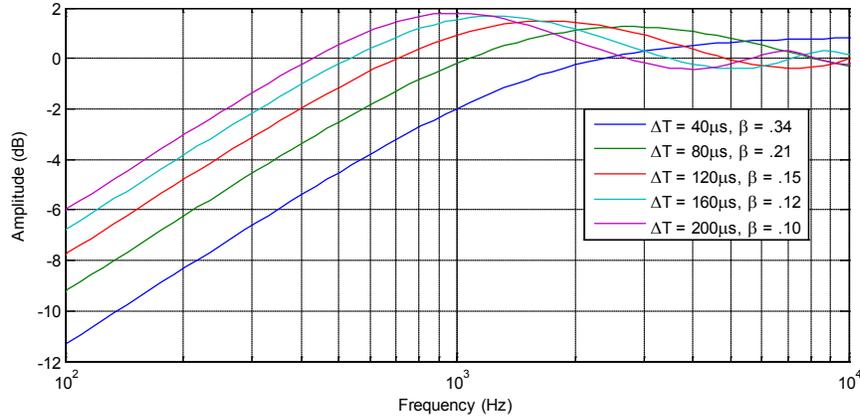


Figure 7: Effect of latency on a the disturbance rejection function of a properly tuned conventional controller operating at 25 kHz, holding phase margin = 60 deg.

As of this moment, there are no hard requirements on the amount of latency compensation that a system must be able to handle. One step of prediction may be sufficient, but it is not unlikely that the total latency in an adaptive optic system will exceed a single timestep for the reasons previously mentioned. In this paper, we evaluate the performance of a new type of adaptive optics controller. We implement intelligent algorithms that predict over a short temporal horizon the evolution of aero optic aberrations. We realize this through model reduction using Proper Orthogonal Decomposition (POD) coupled with a predictive Artificial Neural Network (ANN). We attempt to predict the flow’s evolution over multiple timesteps and show what performance gains are possible through simulation of a feedback controller based on disturbances from actual AAOL flight test data.

II. Theoretical Background

The objective of this paper is to develop a latency-tolerant architecture for airborne adaptive optic systems. As a first step, we propose the existence of a general prediction model. This predictor has the ability to precisely specify the state of the flow at some time in the future given only the “current” flow conditions. Therefore, it can be assumed that important attributes such as boundary conditions are in a sense built into this predictor function. Let the true state of the flow field at time step k be denoted as \mathbf{z}_k . The true state of the system evolves according to some non-linear vector-valued function \mathbf{g} , which is our general predictor,

$$\mathbf{z}_{k+1} = \mathbf{g}(\mathbf{z}_k). \quad [11]$$

In general, this function \mathbf{g} could be represented by a discretization of the Navier-Stokes equations. Practically, however, it is not possible to precisely know the true state of the entire flow field at any given time: this would require full knowledge of density, temperature, velocity, and pressure at each point in the flow. Given the chaotic behavior of the flow, even small measurement errors would quickly give rise to very large errors in the prediction.

A more practical approach is to develop a model that relies solely on observations of relevant flow features while taking advantage of patterns that are recurrent in the flow. From the Gladstone-Dale relationship, it can be quickly seen that optical aberrations are a direct result of density fluctuations in the flow. While these density

fluctuations are typically a result of pressure or temperature changes, it is not necessary to precisely identify these variables in order to predict optical effects. Therefore, we shall make observations of the true state of the flow field using a wavefront sensor, which essentially gives information about the density field of the flow. Let this observation function be denoted \mathbf{h} ,

$$\mathbf{v}_k = \mathbf{h}(\mathbf{z}_k) + \boldsymbol{\varepsilon}_k, \quad [12]$$

where \mathbf{v}_k is the optical wavefront observed at time k and is a function strictly of the flow state plus some measurement error $\boldsymbol{\varepsilon}_k$. The function \mathbf{h} is surjective and not bijective – therefore, while it may be possible to make inferences about the state of the flow field \mathbf{z}_k , it is not possible to fully identify the flow field. This is acceptable as we are only interested in optical effects. Since it is not generally feasible to predict flow features based on underlying physical mechanics – i.e., to utilize Eq. 4 – we seek instead to approximate some other function that describes the evolution of optical wavefronts. This can be described as:

$$\mathbf{v}_{k+1} = \mathbf{f}(\mathbf{v}_k, \mathbf{v}_{k-1}, \dots, \mathbf{v}_{k-M+1}) + \boldsymbol{\varepsilon}_{k+1} \quad [13]$$

We seek to achieve our objective by approximating the function \mathbf{f} . The optimal choice of \mathbf{f} will minimize the mean-square error of all $\boldsymbol{\varepsilon}_k$, i.e. we wish to minimize $E = \sum_{k=1}^L \|\boldsymbol{\varepsilon}_k\|^2$, where L is the total number of observations available. This is known as a Non-Linear Autoregressive (NAR) problem, and can be solved in a number of ways as will be discussed further.

An additional challenge is that the number of elements in \mathbf{v}_k for useful and practical applications can be on the order of 10^3 . If one were to assume a sufficiently small time step that the temporal evolution of wavefronts could be described by a linear system, $\mathbf{v}_{k+1} = \mathbf{A}\mathbf{v}_k$, a full-rank matrix \mathbf{A} would have on the order of 10^6 elements and thus require far too much computational power for real-time application. A potentially more robust approach is to reduce the dimensionality of the problem using modal decomposition²¹. Additionally, Goorksey et al¹⁷ have used Dynamic Mode Decomposition¹⁸ to estimate a low-rank approximation of \mathbf{A} to improve real-time feasibility.

The problem of calculating POD modes, also known as Karhunen-Loève (K-L) modes, in the multi-dimensional case for discrete-time data is typically solved using the singular-value decomposition (SVD), but may also be solved by solving the eigenvalue problem on a dataset's autocorrelation matrix. In this case, let \mathbf{V} be a matrix of measurement snapshots of wavefronts organized by column vectors of samples ordered by increasing time as shown,

$$\mathbf{V} = [\mathbf{v}_1 \quad \mathbf{v}_2 \quad \dots \quad \mathbf{v}_N]^T. \quad [14]$$

Then \mathbf{V} may be decomposed using SVD as

$$\mathbf{V} = \mathbf{U}\boldsymbol{\Sigma}\boldsymbol{\Phi}^H, \quad [15]$$

and the spatial modes may be extracted from the columns of $\boldsymbol{\Phi}$,

$$\boldsymbol{\Phi} = [\boldsymbol{\varphi}_1 \quad \boldsymbol{\varphi}_2 \quad \dots \quad \boldsymbol{\varphi}_R], \quad [16]$$

where $R = \text{rank}(\mathbf{V})$. The temporal coefficients are then calculated from a projection of the spatial modes onto the original observations. Taking advantage of the fact that the pseudoinverse of an orthonormal matrix is its Hermitian transpose, this can be written as

$$\mathbf{x} = \boldsymbol{\Phi}^+ \mathbf{V} = \boldsymbol{\Phi}^H \mathbf{V}, \quad [17]$$

where \mathbf{x} is a matrix of temporal coefficients. This is an important fact in terms of computational efficiency since it means that the projection of wavefronts onto POD modes is a simple matrix multiplication. The POD modes are ranked by the importance of their contribution to the overall energy of the system. Quite often, the POD modes converge quickly to give a good low-dimensional model. Additionally, in the case of naturally occurring fluid flow, low order modes typically exhibit smooth behavior. If it is assumed that the POD modes do not change, then each wavefront can be decomposed as follows,

$$\mathbf{v}_k = \sum_{n=1}^N x_n(k) \boldsymbol{\Phi}_n, \quad [18]$$

where $\boldsymbol{\Phi}_n$ are the time-invariant POD modes, $x_n(k)$ are the coefficients at timestep k , and N is the desired truncation dimension. Additionally, we assume that M snapshots are sufficient to approximate the next wavefront in the sequence: we refer to M as the embedding dimension²². In this case, the NAR problem becomes a function of the modal coefficients, as shown,

$$\mathbf{v}_{k+1} = \boldsymbol{\Phi}_1^N \mathbf{g}(\mathbf{x}_k, \mathbf{x}_{k-1}, \dots, \mathbf{x}_{k-M+1}) + \boldsymbol{\varepsilon}_{k+1}, \quad [19]$$

where $\Phi_1^N \in \mathbb{R}^{L \times N}$ are the truncated set of POD modes and $\mathbf{g}: \mathbb{R}^{N \times M} \rightarrow \mathbb{R}^N$ is the nonlinear prediction function. Thus, it is necessary to use some method to estimate the function \mathbf{g} . Now that we have chosen a means of approximating a sequence of wavefronts it is necessary to develop a model that can predict the evolution of these modes. This is certainly challenging, as the behavior of the modal coefficients in a real flow will be, in general, non-linear. If the timestep is sufficiently small linear behavior can be assumed. If this assumption is made, methods such as DMD¹⁷ or linear state space modeling can be used.

A visualization of POD modes for the 114 degree viewing angle from AAOL data is shown in Figure 8. The reader may note that the modes tend to come in pairs that often capture the convective nature of turbulent flow, especially for the lower-order modes. The higher-order modes tend to contain higher spatial frequencies, and tend to become progressively less useful in capturing turbulent structures. It should therefore be expected that any predictive element should do a better job with the lower order modes than with higher order modes. Additionally, it is important to note that the POD truncation effectively imposes a limit on the performance of the controller. This limitation is shown in Figure 9.

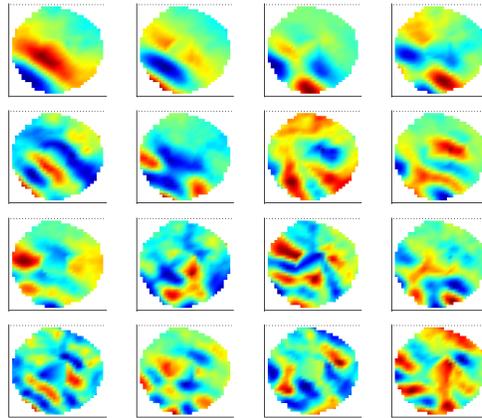


Figure 8: First sixteen POD modes from the 114.7 deg. viewing angle from AAOL flight tests, from most energetic (top-left) to least energetic (bottom right).

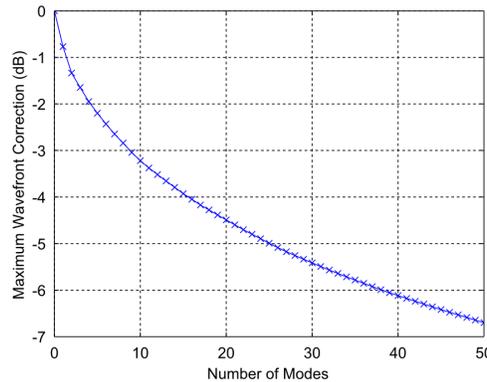


Figure 9: Maximum achievable correction for a given number of POD modes for the 114 deg. viewing angle.

In our case, we choose to use an Artificial Neural Network²³ to model the behavior of the POD coefficients, as they have been demonstrated to be very useful for non-linear dynamical systems²⁴. It has been shown that any multilayer network²⁵ can approximate, to an arbitrary degree of accuracy, any non-linear function with sufficiently many hidden neurons.

The network topology used in this paper consists of one input layer, one hidden layer, and one output layer as shown in Figure 10. This is the minimal topology needed to approximate an arbitrary non-linear system, and since it is desirable to keep the number of necessary computations to a minimum, this is deemed to be a good starting point for network optimization. The inputs to each neuron are essentially a weighted summation of every single network input with the addition of a bias term. These inputs are then passed to a sigmoid activation function (in this

case, a hyperbolic tangent). This sigmoid function gives the network its non-linear character. The outputs are then combined linearly to form the output layer. Thus, the training objectives are to find the input weights, output weights, input bias, and output bias that minimize mean-square error. This can be achieved using methods such as the Levenberg-Marquardt²⁶ algorithm or the iterative reverse propagation (iRPOP+)²⁷ algorithm, which are both adaptations of the classical Newtonian gradient descent algorithm that operate by backpropagating errors through the network.

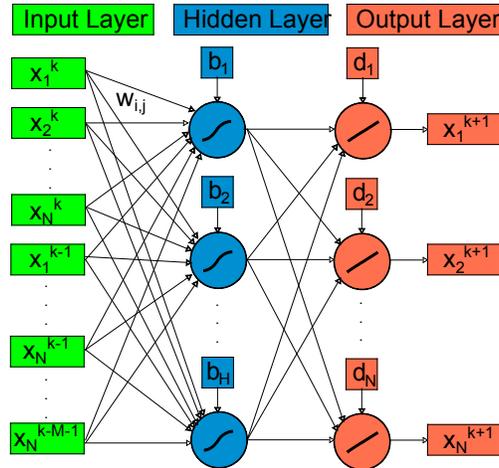


Figure 10: Predictive Artificial Neural Network topology. The network takes M input vectors prior to and including timestep k , to predict the system state at timestep $k+1$. H is the number of neurons in the hidden layer. N is the number of state variables in the vector x .

There are some additional considerations to be made when training the neural network. One must manage the quality of the regression on the training data set against the ability of the network to generalize to new data: in other words, a low mean-square error does not imply that the network will be able to accurately predict new patterns or even slight variations on previously-seen patterns. This is exactly analogous to the canonical example of fitting a high-order polynomial to a linear process corrupted with noise. The neural network can indeed detect some of the nonlinearities present in the underlying process and adapt to them, but if the network is too large then the overfitting problem will begin to manifest itself. However, this overfitting problem can show up in other ways. Even if the network is slightly too big, overfitting can be mitigated by either “early stopping” or regularization techniques. Early stopping consists of exiting the training loop before a local minimum on the residual error is reached during gradient descent: generally, this means that the weights in the network will be smaller in magnitude, forcing the network to tend toward linearity. Another method is to use explicit regularization, which imposes a penalty on the weights of the network. We use early stopping in this work, but will examine other regularization techniques in future work.

Another consideration is the local minimization problem: neural networks will typically not converge to a globally minimal mean-square error during training. One way to mitigate this is to train a collection, or an *ensemble*, of neural networks using similar training data and then average the output of each network to obtain a better result. The ensembling technique is widely^{28,29} used in practical applications of neural networks, and the authors have found it to be beneficial for the aero-optic prediction problem as well. This begs the question: should the training data set for each member of the ensemble be identical to all others? This has been the subjective of a fair amount of research in the statistical learning community. A number of methods have been developed to answer this question: two commonly-used approaches are adaptive boosting³⁰ (AdaBoost) and bootstrap aggregating³¹ (bagging). Since bagging has been used with a good deal of success in neural networks, we chose this approach for the analysis in this paper. Bagging simply consists of uniformly sampling from the training data set for the number of samples actually included in the set. In general, this leads to a number of unique samples in the training data set as well as a number of repeated samples.

Finally, the algorithm also uses a gain-tuning algorithm to obtain the optimal PI feedback gain for the training data. This is done using a simple gradient decent, which is run iteratively until convergence is achieved:

$$\beta_{n+1} = \beta_n - \frac{E'(\beta_n)}{E''(\beta_n)}. \quad [20]$$

III. System Architecture

The system architecture proposed is shown in Figure 11. One of the benefits to the predictive controller is that it is designed to augment a classical adaptive optics control loop in order to reduce the number of integration changes required to existing AO systems. The reader may observe that there are effectively two feedback loops: one optical feedback loop (the “outer” loop) as well as a feedback loop for the DM controller itself (the “inner” loop).

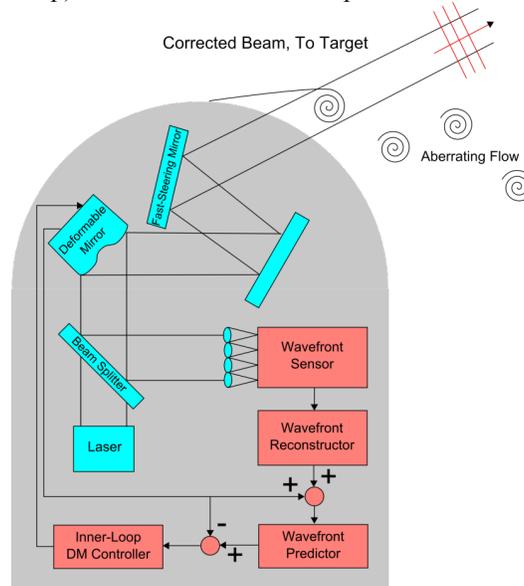


Figure 11: Augmented classical controller with predictive wavefront block.

In practical application, there are a few approaches that could be used to train the predictive controller. One approach is to “pre-train” the network and store a grid of POD mode sets and neural networks for various viewing angles and flight conditions. Another approach is to train the POD modes and neural network in real-time (but at a reduced update rate) and then update the network with the new information once those computations are complete. Such a system would operate at different rates: a high-speed wavefront control loop and a low-speed training loop that would synchronize with each other at the end of each cycle in the low-speed loop. There are potential issues with this, such as system stability during the transition region, but nonetheless this may be a viable option.

IV. Results

The algorithm was tested on actual flight data from AAOL on a range of viewing angles. The data was taken at 25kHz. In each case, the algorithm was trained on 1500 wavefront snapshots and then tested on the following 1500 snapshots. A simulation package was written in C and executed on the Notre Dame Center for Research Computing Cluster using all available AAOL flight test data.

An optimization study was conducted using this simulation package to learn more about the predictive capabilities of this algorithm. Parameters varied were H , the number of hidden neurons; M , the embedding dimension; K , the number of prediction steps; and N , the truncation dimension (or number of POD modes used by the predictive controller). We found that as one may expect, as K increases the benefit from a higher value of N decreases: that is, the higher-order modes become harder to predict over a longer period of time. We found that H should be approximately 50% greater than the value of N . Finally, the optimal value of the embedding dimension, M , was found to be approximately 4. For a value of $M \leq 2$, the network was unable to make good predictions since it was underfitting the training data. For $M \geq 6$, the network suffered from overtraining – that is, it attempts to replicate nonphysical patterns and/or noise.

The mean disturbance rejection for five different amounts of latency is shown in Figure 12. As one may expect, the quality of the prediction and thus the mean disturbance rejection becomes worse as the latency increases. Nonetheless, these are very large improvements from the conventional controller.

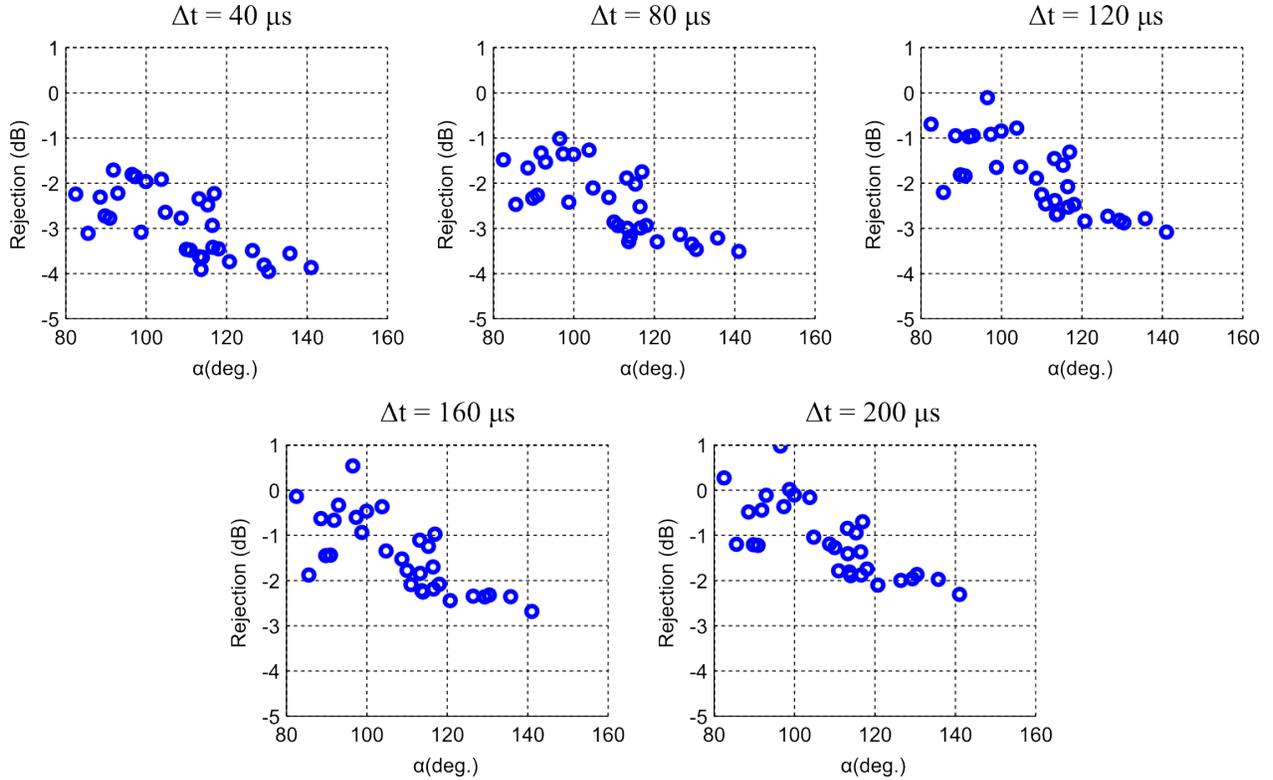


Figure 12: Mean disturbance rejection for the simulated controller using all available flight test data from AAOL.

Figure 13 shows the disturbance rejection function for three viewing angles and five different values of latency. The controller actually performs very poorly on the low-frequency end of the spectrum, which is exactly opposite of what one might expect from a PI controller. However, it seems to concentrate on suppressing disturbances near 1 kHz, which is where the dominant disturbances occur. This is favorable behavior, since it indicates that this neural network approach will automatically seek to suppress the dominant disturbances and ignore the less important frequencies. The reason this occurs is two-fold: clearly, POD and the neural network focuses on suppressing the dominant disturbances; however, a side-effect of only controlling 16 POD modes is that it effectively means that some of the disturbances are “filtered” out of the control loop. Effectively, this is a *modal* controller that simply focuses on suppressing the most optically-active modes.

This also suggests that future iterations of this control design could perhaps use split-frequency control or similar configuration. One may consider that a wavefront disturbance can be split into two parts for the modal controller we have analyzed in this work:

$$\mathbf{v}(t) = \mathbf{v}_{\text{controlled}}(t) + \mathbf{v}_{\text{uncontrolled}}(t) = \sum_{i=1}^N x_i(t)\Phi_i + \sum_{j=N+1}^{\text{Rank}(\mathbf{V})} x_j(t)\Phi_j. \quad [21]$$

That is to say: the uncontrolled part of the wavefront is orthogonal to the controlled part, since the controlled part of the wavefront is simply a superposition of the first N POD modes. Thus, it may be possible to use this property to develop a simple controller with a very low feedback gain for the “uncontrolled” part of the disturbance. This may yield a substantial improvement in controller performance on the low end of the frequency spectrum.

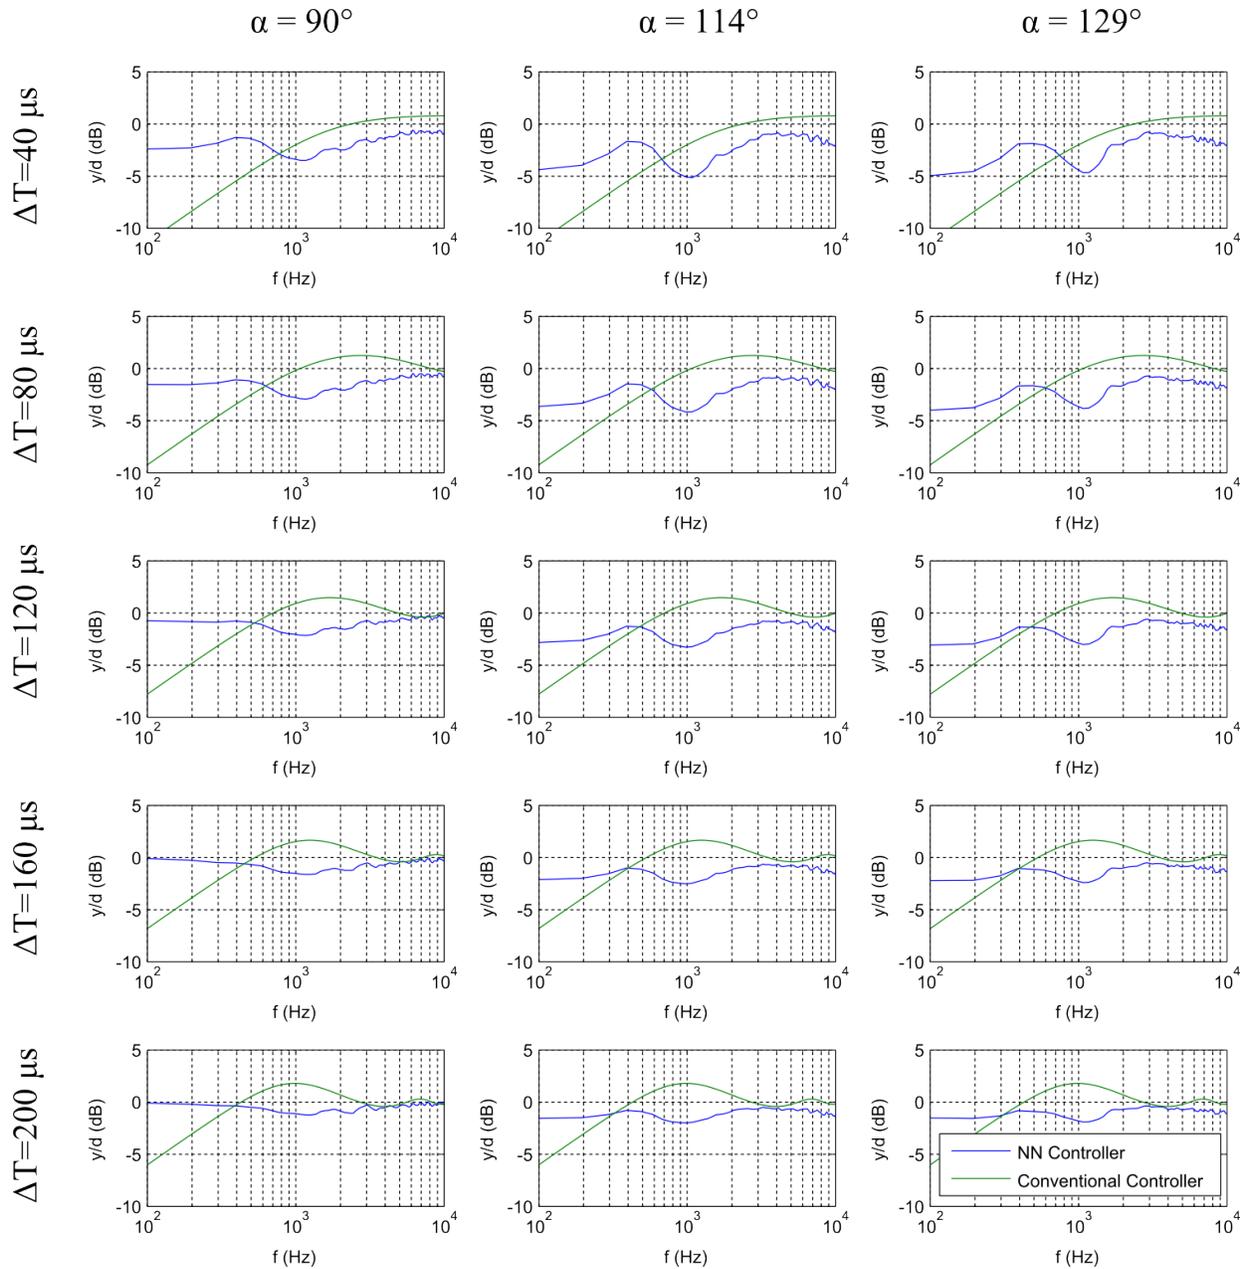


Figure 13: Error rejection function for a number of latency amounts for three different viewing angles.

A time series of the POD tracking performance of the algorithm for a viewing angle of 114 degrees is shown in Figure 14. The algorithm does a very good job of tracking the first few modes, with performance slowly tailing off as the constituent frequencies of the modes increase and the behavior of the modal coefficients becomes more complex. An important result from this observation is that as the desired prediction window increases, the number of modes that can be reliably tracked will decrease.

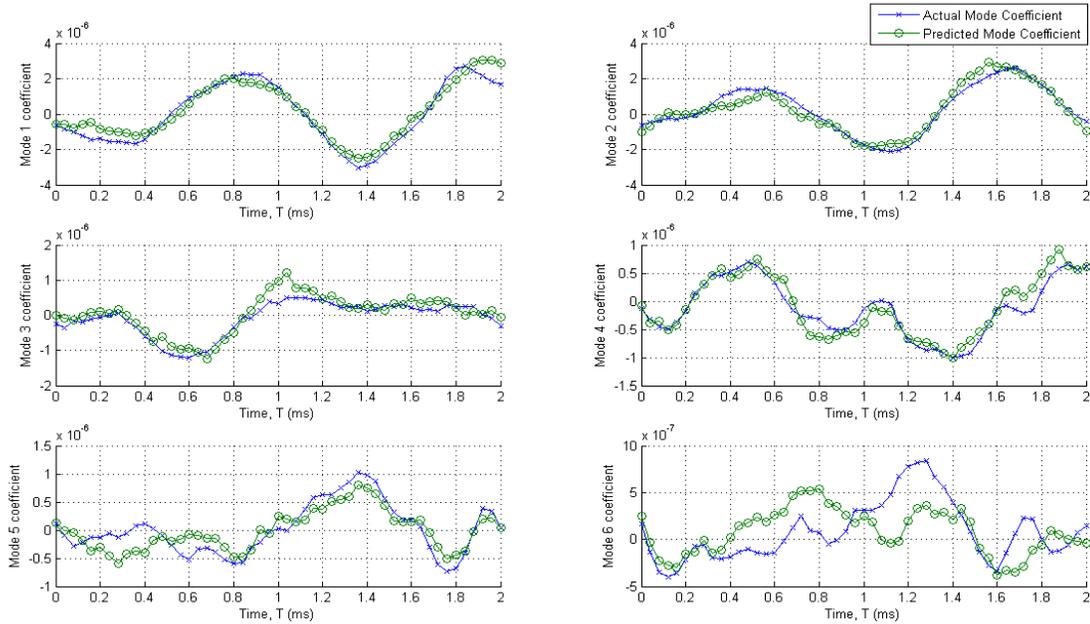


Figure 14: POD mode coefficient prediction for $\alpha = 114$ deg, 5 step prediction (200 μ s latency assumed), with $H = 24$, $M = 4$, $N = 16$.

V. Conclusions

The aero-optics problem is one of the primary obstacles in the implementation of airborne laser systems. We have shown simulations of a new type of adaptive optic controller that is capable of compensating for latency in a conventional adaptive optics controller, significantly increasing performance over multiple time steps of delay. At a single timestep (40 μ s) of latency for a 25 kHz controller, we have demonstrated the feasibility of better than 55% reduction in effective OPD_{rms} in the most optically-aberrating environments. At 5 timesteps (200 μ s) of latency, we have demonstrated the feasibility of approximately 35% reduction in OPD_{rms} . With this much latency, a conventional adaptive optics system would require such a low gain that mitigation would be hardly worthwhile. This degree of mitigation is very promising given that 200 μ s is a very long computational window. The results we have obtained so far show similar mitigation to the DMD-based controller of Goorskey et al¹⁷ for one timestep, but we attempt to predict over a larger temporal window. We would also like to point out that for larger turrets, the necessary sampling frequencies for effective control scale inversely with the turret diameter. For example, a notional turret three times as large as the AAOL turret would require only a ~ 8 kHz controller and the effective prediction window for the controller examined in this work would jump from 200 μ s to 600 μ s – an even softer computational requirement.

Future work will focus on analyzing the stability properties of this controller, online learning, realistic dynamics models for the deformable mirror and amplifiers, computational efficiency, and improving the performance of the controller on the low end of the frequency spectrum. In particular, online learning is likely necessary for a robust real-time controller. The ability to adapt to slowly changing flow conditions is certainly feasible computationally with current hardware. However, updating the controller during these changes without affecting the stability of the controller could prove to be challenging, and the authors would like to develop a better understanding of this problem. The inclusion of realistic dynamics models for the deformable mirror and amplifiers will likely not prove to be fundamentally challenging, but are nonetheless needed to improve confidence in the method. Methods such as model predictive control (MPC) will likely be used in future work to better synergize the benefits of the POD/ANN predictive method with the feedback controller itself. Also, it may be desirable to use methods such as frequency splitting to improve controller performance on the low end, while leaving the higher-frequency components to the predictive controller.

References

- ¹ Wang, M., Mani, A., and Gordeyev, S., “Physics and Computation of Aero-Optics,” *Annual Review of Fluid Mechanics*, vol. 44, Jan. 2012, pp. 299–321.
- ² Gordeyev, S., and Jumper, E., “Fluid dynamics and aero-optics of turrets,” *Progress in Aerospace Sciences*, vol. 46, Nov. 2010, pp. 388–400.
- ³ Gladstone, J. H., and Dale, T. P., “Researches on the Refraction, Dispersion, and Sensitiveness of Liquids,” *Philosophical Transactions of the Royal Society of London*, vol. 153, 1863, pp. 317–343.
- ⁴ Mahajan, V. N., “Strehl ratio for primary aberrations: some analytical results for circular and annular pupils,” *JOSA*, vol. 72, 1982, pp. 1258–1266.
- ⁵ Ross, T. S., “Limitations and applicability of the Maréchal approximation,” *Applied optics*, vol. 48, 2009, pp. 1812–1818.
- ⁶ Porter, C., Gordeyev, S., and Jumper, E., “Large-aperture approximation for not-so-large apertures,” *Optical Engineering*, vol. 52, 2013, pp. 071417–071417.
- ⁷ Jumper, E. J., Zenk, M. A., Gordeyev, S., Cavalieri, D., and Whiteley, M. R., “Airborne aero-optics laboratory,” *Optical Engineering*, vol. 52, 2013, pp. 071408–071408.
- ⁸ Gordeyev, S., Burns, R., Jumper, E., Gogineni, S., Paul, M., and Wittich, D. J., “Aero-optical mitigation of shocks around turrets at transonic speeds using passive flow control,” *51st AIAA Aerospace Sciences Meeting*, 2013, pp. 2013–0717.
- ⁹ Smith, A. E., and Gordeyev, S., “Aero-Optical Mitigation of Turbulent Boundary Layers Using Large-Eddy Break-Up Devices,” *52nd Aerospace Sciences Meeting*, 2014, pp. 2014–0321.
- ¹⁰ Vukasinovic, B., Glezer, A., Gordeyev, S., Jumper, E., and Kibens, V., “Active control and optical diagnostics of the flow over a hemispherical turret,” *AIAA Paper*, vol. 598, 2008, p. 2008.
- ¹¹ Vukasinovic, B., Glezer, A., Gordeyev, S., Jumper, E., and Bower, W. W., “Flow control for aero-optics application,” *Experiments in Fluids*, vol. 54, Mar. 2013.
- ¹² Vukasinovic, B., Gissen, A., Glezer, A., and Gogineni, S., “Fluidic Control of Transonic Shock-Induced Separation,” American Institute of Aeronautics and Astronautics, 2013.
- ¹³ Gissen, A. N., Vukasinovic, B., Glezer, A., Gogineni, S., Paul, M. C., and Wittich, D. J., “Active Transonic Shock Control,” American Institute of Aeronautics and Astronautics, 2014.
- ¹⁴ Burns, R., Gordeyev, S., Jumper, E. J., Gogineni, S., Paul, M., and Wittich, D. J., “Estimation of aero-optical wavefronts using optical and non-optical measurements,” *52nd AIAA Aerospace Sciences Meeting*, 2014, pp. 2014–0319.
- ¹⁵ Nightingale, A. M., Goodwine, B., Lemmon, M., and Jumper, E. J., “Phase-Locked-Loop Adaptive-Optic Controller and Simulated Shear Layer Correction,” *AIAA Journal*, vol. 51, Nov. 2013, pp. 2714–2726.
- ¹⁶ Faghihi, A., Tesch, J., and Gibson, S., “Identified state-space prediction model for aero-optical wavefronts,” *Optical Engineering*, vol. 52, 2013, pp. 071419–071419.
- ¹⁷ Goorskey, D. J., Schmidt, J., and Whiteley, M. R., “Efficacy of predictive wavefront control for compensating aero-optical aberrations,” *Optical Engineering*, vol. 52, 2013, pp. 071418–071418.
- ¹⁸ Schmid, P. J., “Dynamic mode decomposition of numerical and experimental data,” *Journal of Fluid Mechanics*, vol. 656, Aug. 2010, pp. 5–28.
- ¹⁹ Ogata, K., *Discrete-Time Control Systems*, Englewood Cliffs, NJ: Prentice Hall, 1995.
- ²⁰ Ogata, K., *Modern Control Engineering*, Upper Saddle River, NJ: Prentice Hall, 2010.
- ²¹ Lumley, J. L., *Stochastic Tools in Turbulence*, Courier Dover Publications, 2007.
- ²² Kennel, M. B., Brown, R., and Abarbanel, H. D. I., “Determining embedding dimension for phase-space reconstruction using a geometrical construction.”
- ²³ Lippmann, R., “An Introduction to Computing with Neural Nets,” *IEEE ASSP Magazine*, vol. 4, Apr. 1987, pp. 4–22.
- ²⁴ Narendra, K. S., and Parthasarathy, K., “Identification and control of dynamical systems using neural networks,” *IEEE Transactions on Neural Networks*, vol. 1, 1990, pp. 4–27.
- ²⁵ Hornik, K., Stinchcombe, M., and White, H., “Multilayer Feedforward Networks Are Universal Approximators,” *Neural Networks*, vol. 2, 1989, pp. 359–366.
- ²⁶ Demuth, H., Beale, M., and Hagan, M., “Neural network toolbox™ 6,” *MATLAB® User’s guide*, 2008.
- ²⁷ Igel, C., and Hüsken, M., “Improving the Rprop learning algorithm,” *Proceedings of the second international ICSC symposium on neural computation (NC 2000)*, Citeseer, 2000, pp. 115–121.
- ²⁸ Hansen, L. K., and Salamon, P., “Neural Network Ensembles,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, Oct. 1990.

- ²⁹Zhou, Z.-H., Wu, J., and Tang, W., “Ensembling neural networks: many could be better than all,” *Artificial intelligence*, vol. 137, 2002, pp. 239–263.
- ³⁰Freund, Y., and Schapire, R. E., “A decision-theoretic generalization of on-line learning and an application to boosting,” *Computational learning theory*, Springer, 1995, pp. 23–37.
- ³¹Breiman, L., “Bagging Predictors,” *Machine Learning*, vol. 24, 1996, pp. 123–140.