

Performance of AO predictive control in the presence of non-stationary turbulence.

Maaïke van Kooten¹, Niek Doelman^{1,2}, and Matthew Kenworthy¹

¹Leiden Observatory, Niels Bohrweg 2, 2333 CA Leiden, The Netherlands

²TNO Technical Sciences, Stieltjesweg 1, 2628 CK Delft, The Netherlands.

ABSTRACT

Current high contrast imaging systems are limited at small angular separations, preventing the direct imaging of Earth-size exoplanets close to their host stars. One primary cause of this is the performance of the extreme adaptive optics (AO) system which is dominated by the servo-lag error at small angular separations. Prediction can be used to help reduce this servo-lag error. Most AO systems today do not use predictive controllers (integral controller), resulting in a phase correction that is proportional to the current measured wavefront and does not take into account the evolution of the wavefront phase fluctuations between measurement and correction. To improve upon this, we focus on the development of different predictive control strategies that estimate how the wavefront phase fluctuations evolve over a time horizon equal to the servo-lag. Moreover, the statistical properties of the turbulence phase are non-stationary (change in time) which needs to be included in the AO analysis for high contrast imaging systems. First, we show that the non-stationary properties of the atmosphere can be modeled using a von Karman covariance function which can incorporate variations in time of the Fried parameter, outer scale, and wind velocity. Using this new disturbance model of turbulence, the performance of three different predictors—based on a steady state, a recursive, and an adaptive linear-minimum-mean-square-estimator (LMMSE)—is tested under non-stationary conditions. We feed the prediction algorithms wavefront slopes from a 11-by-11 subaperture Shack-Hartmann wavefront sensor. A 97 actuator deformable mirror applies the predictor’s phase correction. We present the latest results of our simulations and compare our predictors with the common integrator. We show that under our varying wind conditions, the root-mean-square wavefront error can increase by a factor of two for both the integrator and our predictors. In our simulations, we do not observe an increase in performance using an exponential forgetting factor adaptive LMMSE.

Keywords: High Contrast Imaging, Adaptive Optics, Predictive control

1. INTRODUCTION

Since the first direct image of an exoplanet,¹ high contrast imaging (HCI) has been a very active field. Capable of providing insight into planet formation, characterizing an exoplanet’s atmosphere, and searching for signs of habitability, HCI spatially separates the planet’s photons from the ones emitted by the host star. With multiple dedicated instruments such as SPHERE² on the VLT, GPI³ on Gemini South, and SCExAO⁴ on Subaru many new techniques, better coronagraphs, and low noise detectors have been developed. These efforts have resulted in contrast levels of $10^{-4} - 10^{-5}$ at $0.5'' - 1''$ separation⁵ in the near infrared (see Figure 1).

However, to answer the fundamental questions of: is the earth unique? Is our solar system unique? And how did we form? we need to aim for contrast levels of $10^{-7} - 10^{-9}$ at separations of $0.05'' - 1''$ (in the visible) to find cold jupiters, neptunes, and rocky planets. Currently, the limiting factor is the correction achieved by the adaptive optics systems within the HCI systems. Specifically, the chromatic index and servo-lag error prevent us for observing at small separations as shown in Figure 2. We focus on minimizing the servo-lag error. One approach is to run our AO system faster, requiring high speed deformable mirrors, fast wavefront sensors, and bright guide stars. Another approach is to focus on understanding what happens to the induced wavefront phase fluctuations to correct for the changes that occur during the servo-lag. This is done through predictive control methods.

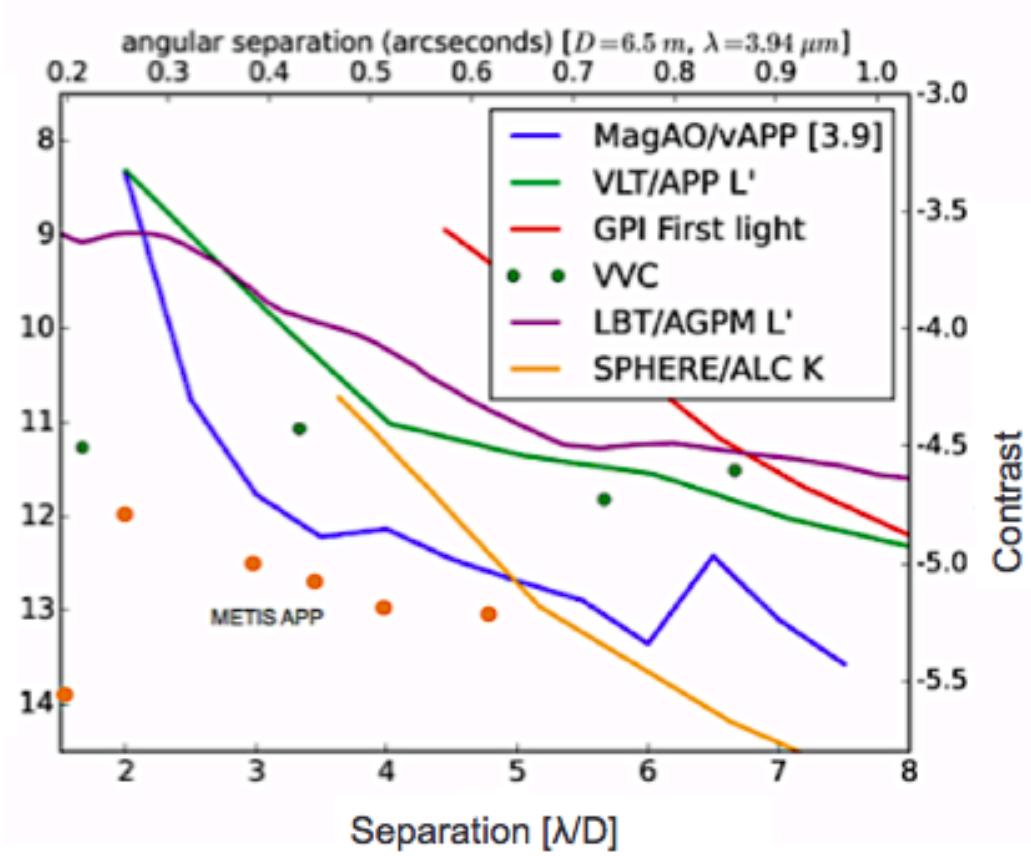


Figure 1. Current and future high contrast imaging capabilities.⁶ The ELTs HCI instrument METIS was added by authors.

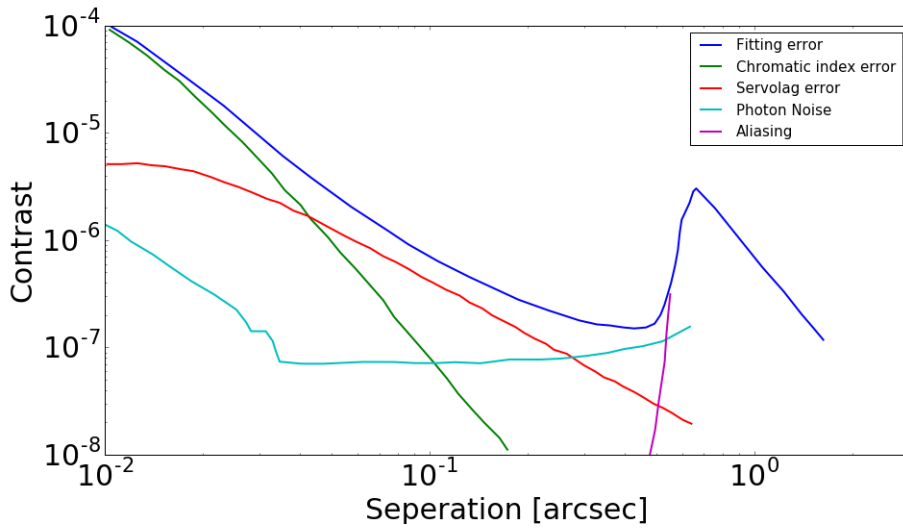


Figure 2. An example of an adaptive optics error budget showing that the servo-lag error is a major limitation of achievable contrast⁵ at small angular separations.

1.1 Predictive control in adaptive optics

As early as the 1970s,⁷ predictive control has been proposed to mitigate the effect of inevitable time delays in adaptive optics systems. To date, linear prediction has been very successful for tip tilt compensation, especially when the source of the vibrations are well characterized, using linear quadratic gaussian control (LQG; Kalman filter as the predictor).^{8–10} Work has been done on the LQG controller for full atmospheric disturbance.^{11–13} Different varieties of the Kalman filter predictor,^{14,15} which have allowed for filtering in real time, are now being used in laboratory testing¹⁶ as well in on sky testing.¹⁷ The data driven H_2 optimal control¹⁸ was also tested on-sky¹⁹ showing a reduction in the temporal error. Finally, adaptive predictors have been proposed²⁰ to track the time-varying behavior of atmospheric turbulence.

Similar to any control problem, the characteristics of the system (XAO system), disturbance, and the time scales all influence the choice of the controller. Due to the variations of the mean and variance of phase aberrations, atmospheric turbulence induced phase fluctuations are non-stationary. If we account for this behavior, we will be able to find a better performing predictive controller.

In this paper, we focus on understanding the input disturbance —atmospheric turbulence— of an AO system on short time scales. We look at how the disturbance model affects the performance of an AO system using basic prediction methods. The predictors are described in Section 2. The statistical behavior and the non-stationarities of turbulence are described in Section 3. We present, in Section 4, the results our AO simulation with our new turbulence model and we discuss the impact of the results on predictor design. Finally, we outline our ongoing and future work.

2. PREDICTION FILTER DESIGN

We implement a linear finite order predictor. We denote a single phase point within a phase screen as y_i . Assuming that the future predicted value of a given phase point, \hat{y}_i at the discrete time index $t + 1$, is a linear combination of previous phase values, $\vec{w}(t)$ – a $P \times 1$ column vector containing a collection of previous true phase values, we define the following, where P is the number of prior phase points:

$$\hat{y}_i(t + 1) = \vec{A}_i \cdot \vec{w}(t) \quad (1)$$

$$\vec{w}(t) = (y_0(t) \quad y_1(t) \quad y_2(t) \quad \dots \quad y_P(t))^T \quad (2)$$

The prediction vector, \vec{A}_i , is then determined by minimizing the linear-minimum-mean-square-error cost function:

$$\min_{\vec{A}_i} \langle ||y_i(t) - \vec{A}_i \vec{w}(t)||^2 \rangle \quad (3)$$

The strength of this method is that the algorithm immediately extracts any spatio-temporal information from the data provided.

We implement our linear-minimum-mean-square-estimator (LMMSE) in three different ways. A batch method LMMSE is performed by storing measurements and determining the filter once enough measurements have been taken (after Q time steps). During this period, the LMMSE cannot be used and is considered offline. Solving batch-wise, we first define

$$\vec{W} = (\vec{w}(t) \quad \vec{w}(t - 1) \quad \vec{w}(t - 2) \quad \dots \quad \vec{w}(t - Q)) \quad (4)$$

and

$$\vec{Y} = (y_i(t + 1) \quad y_i(t) \quad y_i(t - 1) \quad \dots \quad y_i(t - (Q + 1)))^T \quad (5)$$

By substituting \vec{W} and \vec{Y} for \vec{w} and \vec{y} in Equation 3, respectively, we find the prediction vector to be:

$$\vec{A}_i = \vec{W}^+ \vec{Y} \quad (6)$$




Figure 3. The phase points at $t=N$ to use to predict the point at $t=N+1$.

Here a pseudo inverse (denoted by $+$) is performed to find the prediction vector. We can also rewrite this equation such that the prediction vector is the product of the inverse of the auto-covariance matrix and cross-covariance vector.

$$\vec{A}_i = \mathbf{C}_{\vec{w}\vec{w}}^+ \vec{C}_{\vec{w}\hat{y}} \quad (7)$$

In the batch-wise approach, once the prediction vector is found it is fixed for the rest of the simulation. By making use of the matrix inversion lemma, we can form a recursive solution in which the covariance matrices are updated at each time step (Equations 8 and 9), removing the need for a pseudo inverse calculation (greatly reducing the computational load).

$$\mathbf{C}_{\vec{w}\vec{w}}^{-1}(t) = \mathbf{C}_{\vec{w}\vec{w}}^{-1}(t-1) - \frac{\mathbf{C}_{\vec{w}\vec{w}}^{-1}(t-1)\vec{w}(t-1)\vec{w}^T(t-1)\mathbf{C}_{\vec{w}\vec{w}}^{-1}(t-1)}{1 + \vec{w}^T(t-1)\mathbf{C}_{\vec{w}\vec{w}}^{-1}(t-1)\vec{w}(t-1)} \quad (8)$$

$$\vec{C}_{\vec{w}\hat{y}}(t) = \vec{C}_{\vec{w}\hat{y}}(t-1) + \hat{y}(t-1)\vec{w}^T(t-1) \quad (9)$$

The recursive solution can go online immediately with the initial covariance being set to diagonal matrices with large values (as done with recursive least squares methods). The solution is quick to converge (within tens of iterations). The memory of the recursive solution is not limited and therefore all previous measurements affect the solution. We also use a regularization parameter to stabilize the recursion. Finally, we introduce an exponential forgetting factor to the recursive method, creating a simple adaptive method as our third LMMSE.

The LMMSE is convenient in that it allows us to choose the amount of input data used for prediction. We assume the phase screen is simply blowing across the pupil (i.e., the phase fluctuations on the upper half of the pupil were measured –at some point in time– on the lower part half of the pupil). This is known as the Taylor’s Frozen Flow hypothesis. Hence, knowing how far (for example units of subapertures) the phase screen is shifted by, we limit the number of previous phase screens to one. We further limited it such that a given phase point only uses its neighboring phase points (within a radius of 5 grid points) in the prediction algorithm (see Figure 3). This allows for us to reduce the computation time as well as study how the LMMSE behaves for different amounts of input information. A similar thing can be done to build a temporal only predictor by assuming no spatial correlation and that the wavefront phase at each point can be represented by an independent time series. The overall performance might be optimized by providing multiple previous phase screens.

3. NON STATIONARY TURBULENCE MODEL

The spatial behavior of atmospheric turbulence induced phase is often described by a power spectral density (PSD) function in which the Fried parameter, and outer scale can be used to describe the spatial phase variations. We consider a spatial and then derive the temporal PSD by applying Taylor’s Frozen Flow hypothesis which involves the wind speed and direction(see Section 3.3). Many simulations use this prescription to model the input disturbance into the AO control loop. In this situation, any time-shift corresponds to a pure shift of the phase screen due to wind (for a single layer). However, we know that this Frozen Flow approach ignores the time varying behavior of the atmosphere. Therefore, we aim to understand how the atmospheric parameters vary, as a result of the non-stationary behavior of turbulence, and what impact the variations have on the system.

We make use of the von Karman turbulence model by implementing Assemat’s²¹ infinite phase screen that allows for the new rows of atmosphere to be generated online. Their use of a von Karman covariance function to generate the phase allows for the wind speed and Fried parameter (we do not focus on the outer scale as little information is known on its fluctuations and it affects the temporal PSD the same way as the wind speed) to be changed online. To properly vary these parameters we need models that describe how the Fried parameter and wind speed change in time.

3.1 Fried parameter variations

The Fried parameter (or coherence length) is a measure of the strength of turbulence with large numbers corresponding to weak turbulence and small values corresponding to strong turbulence. It is by definition the path integral of the C_n^2 profile and therefore is directly related to changes in the temperature fluctuations in the atmosphere. There is no data on the Fried parameter at a frequency of 1 Hz or greater using available stereo-scidar measurements. Using wavefront sensor data, Doelman²² determined a fractional ARIMA model to describe the evolution of the Fried parameter at La Palma on the time interval of 100 seconds. We adopt this model, creating a time series by assuming the Fried parameter remains constant over these 100 seconds. We feed this series into our atmospheric phase generator.

3.2 Wind variations

Although much work has been done on the variation of wind speed over long time scales, including van der Hoven’s PSD²³ and time evolving PSD for applications in civil engineering,²⁴ we were unable to find an appropriate model useful for our proposes. The slope of the PSD changes significantly depending on the altitude and geometry of the location. Also, the behavior of wind varies from day time to night time. Wind data measurements (minute time scales) in flat agricultural areas (such as Cabauw in the Netherlands operated by Cesar Observatory) show very different variations (in amplitude and in mean wind speed) from data at Mauna Kea or Paranal with the same sampling. There are also added affects such as dome geometry and the telescope movement that contribute to the relative wind velocity seen by the AO system that are not described by these PSDs and data. Finally, the variation in wind direction also needs to be considered.

Due to the lack of data and wind models that accurately describe the wind at an observatory, we assume a worst case scenario – wind gusts combined with turbulent induced fluctuations. We use the Kaimal PSD to generate the small wind fluctuations caused by turbulence and assume the atmosphere is stationary for one second;²⁴ we then change the mean wind speed using a probability density function²⁵ and generate another time series of one second. In the end, we have a time series that resembles wind gusts with some fast fluctuations on top. Although this does not accurately represent the true wind fluctuations, we are subjecting the predictor to small variations and large variations in one time step as well as using wind speeds that are commonly seen at observatories. It also allows us to setup the pipeline for testing future wind speed models.

3.3 Effect on prediction and control

To understand how the wind, outer scale, and Fried parameter affect the disturbance model and the ability of the system to reject the disturbance, we look at the spatial and temporal von Karman PSDs. We define the spatial PSD, for a single point and a single layer, in Equation 10 (neglecting the inner scale) with L_0 as the outer scale. The Fried parameter is represented by r_0 .

$$\Phi(k) = 0.0229r_0^{-5/3}(k^2 + 4\pi^2L_0^2)^{-11/6} \tag{10}$$

For discussion purposes, we write the single-sided temporal PSD²⁶ (Equation 11), substituting in $C_n^2 dh \propto \frac{r_0^{-5/3}}{k^2}$ to give:

$$\Phi(f) \propto \left(\frac{v}{r_0}\right)^{5/3} \left(f^2 + \frac{v^2}{L_0^2}\right)^{-4/3} \quad (11)$$

Note, Equation 11 is slightly different from the temporal PSD used in our simulations by using Frozen Flow to evolve our phase screens.

Equations 10 and 11 show that when the Fried parameter varies both the spatial and temporal gains are changed. The outer scale also affects the spatial PSD gain as well as the cut-off frequency of the temporal PSD. Variations in the wind velocity change the gain and cut-off frequency of the temporal PSD. A predictor ideally is able to account for the variations in the cut-off frequency that are due to the changing disturbance’s statistics (variations in gain are not an issue as the input signal to the predictor encodes any gain variations). Time-invariant predictors are based on stationary disturbance statistics. Therefore, knowledge the dynamical behavior of wind speed and outer scale might be beneficial to achieve to have the AO system best reject the turbulence phase fluctuations, as well as understand the penalties of using a stationarity based predictor under non-stationary conditions.

4. SIMULATIONS

We test three LMMSE predictors along with an integrator (zero-order predictor/no prediction) in an open-loop configuration. We simulate an 11-by-11 Shack-Hartmann wavefront sensor (SHWFS), adding measurement noise to the phase disturbance. We then feed the resulting phase screen (representing the atmosphere at time t) into our predictors and apply the predicted correction (predicted wavefront to time $t + 1$) to a deformable mirror with 97 actuators (DM97). We track the residuals to determine the performance of the predictors.

We do this for various input disturbances including: 1) stationary turbulence, 2) non-stationary turbulence with varying Fried parameter only, 3) non-stationary turbulence with varying wind speed only. We run our AO loop at 500 Hz for 4 seconds, we then average the root-mean-square error (rms) over 50 runs, regenerating the phase screens with the same PSD.

4.1 Results

We initially test the stationary turbulence model with the atmospheric parameters in Table 4.1. From Figure 4, we can see that the LMMSE predictor performs better than the integrator (no prediction) as expected, with the LMMSE having a rms of 0.41 radians and the integrator having an rms of 0.54. This validates the added value of prediction. The exact difference in performance is dependent on the wind speed, the spatial resolution, and the length of the delay. We see that as the wind’s speed reduces, the integrator and LMMSE become closer in performance. The batch and recursive LMMSE also have the same value when the batch predictor is turned on, which indicates our LMMSEs are behaving properly. Finally, as we move further away from the batch training period, the recursive solution becomes slightly better than the batch. As the recursive estimator continues to run it obtains a more accurate estimate of the covariance matrices.

Using the stationary turbulence, we looked at the prediction vector found by the LMMSE. When we limit the radius of points given to the algorithm to 2 grid points, the prediction vector easily finds the perfect predictor -giving the full weighting to the correct point upstream – as expected from the Frozen Flow principle. When using more grid points for prediction, the prediction algorithm has more trouble in finding the optimal solution due to the increase of redundancy in the regressor data. This is impart due to strong spatial correlation of phase between surrounding points. Therefore, the LMMSE is unable to distinguish a upstream phase point and finds an average solution. An increase of measurement noise is expected to alleviate this issue. Overall, this demonstrates the importance of selecting the regressor grid in relation to the amount of data required to converge.

We then vary the Fried parameter. Figure 5 shows that the Fried parameter varies only slightly (0.03% over 5 seconds) on such short time scales and the changes in the Fried parameter do not affect the solution for the predictors as expected. However, such small variations in the Fried parameter likely cannot be seen due to the injected measurement noise and the pixel scaling. Therefore, we also simulated large artificial ramps for the

Simulation name	Fried parameter	Wind speed	Outer Scale	Simulation Length
Stationary turbulence	13 cm	7 m/s	30 m	5 seconds
Non stationary turbulence- Fried parameter 1	13.001 -13.004 cm	7m/s	30 m	5 seconds
Non stationary turbulence- Fried parameter 2	25 -5 cm	7m/s	30 m	5 seconds (1.5 second ramp)
Non-stationary turbulence-wind speed	13 cm	4 - 7 m/s	30 m	5 seconds

Table 1. Summary of input parameters for different simulations.




Figure 4. Residuals (res) of the LMMSE predictors and the integrator compared to the full uncorrected wavefront error for stationary conditions with constant Fried parameter and wind speed.

Fried parameter value. In Figure 6, we show the results of changing the Fried parameter by 20 centimeters, from $D/r_0 \sim 5$ to $D/r_0 \sim 1$, over a few seconds. We see a dramatic decrease in the rms for the LMMSE predictors, the decrease for the LMMSE predictors is a factor of two larger than the integrator. We then study the behavior of the prediction vector in time. We plot the evolution of the prediction coefficient (one value in the prediction vector) for a few different phase points in Figure 7. We see that the prediction coefficients remain constant in time with slight fluctuations. Therefore, the recursive LMMSE solution is not affected by the change in Fried parameter. The decrease in rms followed by the increase in rms, Figure 6, is due to the overall decrease in variance of the phase screen as the coherence length increases, making the phase screens more uniform. The Fried parameter simulation results show that the regular fluctuations of the Fried parameter do no change the performance and on time scales of a few seconds, the Fried parameter varies slowly and the disturbance rejection remains unchanged.




Figure 5. Fried parameter varying according to an fractional ARIMA model.²²




Figure 6. Fried parameter behavior exaggerated as a ramp feature.

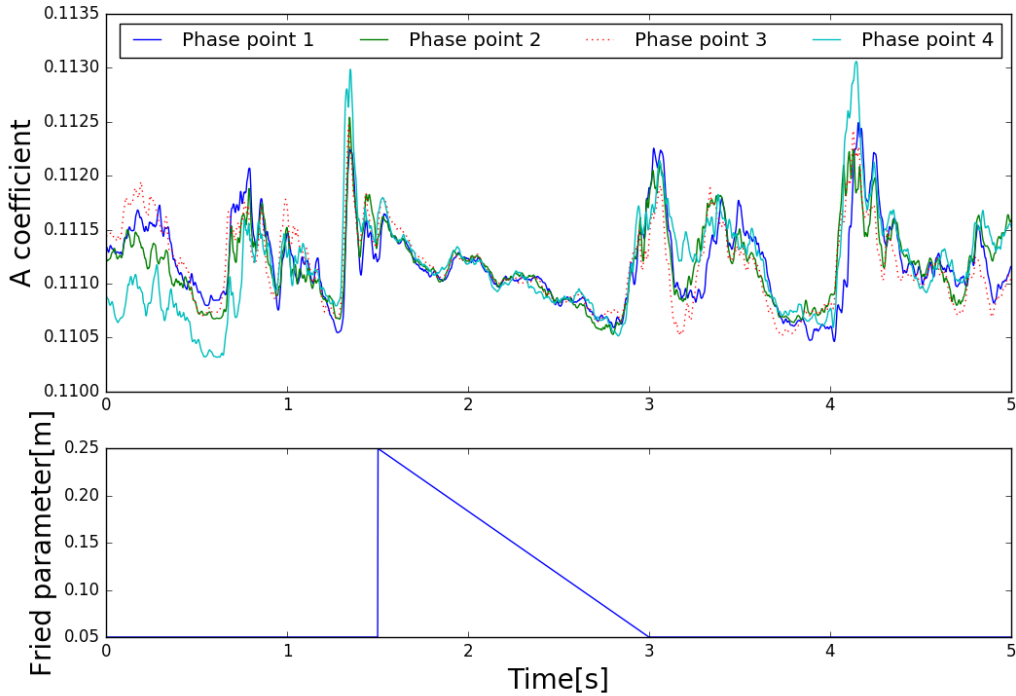


Figure 7. The evolution of a prediction coefficient for 4 different phase points as the Fried parameter varies.

We vary the wind speed according to section 3.2. In Figure 8, we see that wind jumps of 1m/s (realistic amplitude variations) can degrade the performance of not only the integrator but also the batch and recursive LMMSE. The batch method, as expected, does not update and therefore does not see the changing statistics. The recursive LMMSE, although it updates, cannot update fast enough to track the step in wind speed. In Figure 8, the rms wavefront error for all the predictors increases by a factor of two for the recursive LMMSE, at the second wind jump.

The tracking speed of the recursive approach is dependent on how long the prediction has been running; it has 'unlimited' memory and variations in the statistics are averaged into the estimator. Inserting an exponential forgetting factor in the covariance estimation of the recursive LMMSE, we form an adaptive LMMSE. We can see a very slight difference in the performance between the pure recursive and the forgetting factor recursive method (we tested forgetting factor values from 0.99-0.90 with no notable differences, values smaller than this cause numerical instabilities). However, we need more tests to quantify the added value of exponential forgetting in the recursive prediction algorithm.

From these tests, we show that any form of the LMMSE can provide better performance than a no prediction (integrator) approach but the LMMSE is not the optimal solution for varying atmospheric parameters. Due to the non-stationary nature of turbulence induced phase, a basic adaptive approach that is easy to implement has not shown a significant performance improvement. We need to understand specifically how the wind varies on subsecond timescales. Armed with a better understanding of the winds behavior (expected amplitude of jumps), we can begin to explore what type of predictor would be best. We can also begin to look at how information from telemetry and SCIDAR techniques could be combined with predictions methods. We can then decide on what adaptive approach should be best for HCI observations. The specific demands of the observing scheme could be optimized as well. If short exposures (a few seconds or less) are being taken, then say a minute between a set number of exposures could be used to recompute the predictor. If long term performance is needed, than the predictor that performs most reliably should be employed.




Figure 8. Performance of the predictors for varying wind parameters.

5. CONCLUSION

We show that the performance of a predictor to remove residual phase errors due to the inherent time delay in an AO system is dependent on how the wavefront disturbance behaves. We focus on understanding how the non-stationarities of the atmospheric turbulence, specifically wind and Fried parameter, might affect our predictors. We have described a method that can be used to model these non-stationarities and that the change in Fried parameter can be modelled following Doelman.²² We show that using this model, the variations in Fried parameter over short time scales do not change the performance of both the integrator and the LMMSE predictors. Inserting wind gusts (into the simulation), we see that the integrator and the LMMSE predictors are less efficient at rejecting the atmospheric turbulence corresponding to a factor of two increase in the rms wavefront error. To further improve on the prediction performance, we need to better understand how the wind speed evolves. We are currently focusing on developing a dynamical model to represent how wind might behave near an observatory.

ACKNOWLEDGMENTS

The authors would like to thank Leiden University, NOVA, METIS consortium, and TNO for funding this research.

REFERENCES

- [1] Marois, C., Macintosh, B., Barman, T., Zuckerman, B., Song, I., Patience, J., Lafreniere, D., and Doyon, R., “Direct imaging of multiple planets orbiting the star hr 8799,” **322**(5906), 1348–1352 (2008).
- [2] Beuzit, J.-L., Feldt, M., Dohlen, K., Mouillet, D., Puget, P., Wildi, F., Abe, L., Antichi, J., Baruffolo, A., Baudoz, P., Boccaletti, A., Carbillet, M., Charton, J., Claudi, R., Downing, M., Fabron, C., Feautrier, P., Fedrigo, E., Fusco, T., Gach, J.-L., Gratton, R., Henning, T., Hubin, N., Joos, F., Kasper, M., Langlois, M., Lenzen, R., Moutou, C., Pavlov, A., Petit, C., Pragt, J., Rabou, P., Rigal, F., Roelfsema, R., Rousset, G., Saisse, M., Schmid, H.-M., Stadler, E., Thalmann, C., Turatto, M., Udry, S., Vakili, F., and Waters, R., “Sphere: a planet finder instrument for the vlt,” *Proc.SPIE* **7014**, 7014 – 7014 – 12 (2008).
- [3] Macintosh, B., Graham, J. R., Ingraham, P., Konopacky, Q., Marois, C., Perrin, M., Poyneer, L., Bauman, B., Barman, T., Burrows, A. S., Cardwell, A., Chilcote, J., Rosa, R. J. D., Dillon, D., Doyon, R., Dunn, J., Erikson, D., Fitzgerald, M. P., Gavel, D., Goodsell, S., Hartung, M., Hibon, P., Kalas, P., Larkin, J., Maire, J., Marchis, F., Marley, M. S., McBride, J., Millar-Blanchaer, M., Morzinski, K., Norton, A., Oppenheimer, B. R., Palmer, D., Patience, J., Pueyo, L., Rantakyro, F., Sadakuni, N., Saddlemyer, L., Savransky, D., Serio, A., Soummer, R., Sivaramakrishnan, A., Song, I., Thomas, S., Wallace, J. K., Wiktorowicz, S., and Wolff, S., “First light of the gemini planet imager,” **111**(35), 12661–12666 (2014).

- [4] Jovanovic, N., Martinache, F., Guyon, O., C.Clergeon, Singh, G., Kudo, T., Garrel, V., Newman, K., Doughty, D., Lozi, J., Males, J., Minowa, Y., Hayano, Y., Takato, N., Morino, J., Kuhn, J., Serabyn, E., Norris, B., Tuthill, P., Schworer, G., Stewart, P., Close, L., Huby, E., Perrinand S. Lacour, G., Gauchet, L., Vievard, S., Murakami, N., Oshiyama, F., Baba, N., Matsuoand J. Nishikawaand M. Tamura, T., Lai, O., Marchis, F., Duchene, G., Kotani, T., and Woillez, J., “The subaru coronagraphic extreme adaptive optics system: Enabling high-contrast imaging on solar-system scales,” *Publications of the Astronomical Society of the Pacific* **127**, 890 (Sept. 2015).
- [5] Kasper, M., “Adaptive optics for high contrast imaging,” *Proc.SPIE* **8447**, 8447 – 8447 – 10 (2012).
- [6] Otten, G. P. P. L., Snik, F., Kenworthy, M. A., Keller, C. U., Males, J. R., Morzinski, K. M., Close, L. M., Codona, J. L., Hinz, P. M., Hornburg, K. J., Brickson, L. L., and Escuti, M. J., “On-sky Performance Analysis of the Vector Apodizing Phase Plate Coronagraph on MagAO/Clio2,” *The Astrophysical Journal* **834**, 175 (Jan. 2017).
- [7] Hardy, J. W., [*Adaptive Optics for Astronomical Telescopes*] (July 1998).
- [8] Petit, C., Conan, J.-M., Kulcsár, C., Raynaud, H.-F., and Fusco, T., “First laboratory validation of vibration filtering with LQG control law for Adaptive Optics,” *Optics Express* **16**, 87 (2008).
- [9] Doelman, N., Fraanje, R., and den Breeje, R., “Real-sky adaptive optics experiments on optimal control of tip-tilt modes,” *2nd Conference on Adaptive Optics for Extremely Large Telescopes* **1**(1) (2011).
- [10] Poyneer, L. A., Palmer, D. W., Macintosh, B., Savransky, D., Sadakuni, N., Thomas, S., Véran, J.-P., Follette, K. B., Greenbaum, A. Z., Mark Ammons, S., Bailey, V. P., Bauman, B., Cardwell, A., Dillon, D., Gavel, D., Hartung, M., Hibon, P., Perrin, M. D., Rantakyrö, F. T., Sivaramakrishnan, A., and Wang, J. J., “Performance of the Gemini Planet Imager’s adaptive optics system,” *OSA* **55**, 323 (Jan. 2016).
- [11] Paschall, R. N. and Anderson, D. J., “Linear quadratic Gaussian control of a deformable mirror adaptive optics system with time-delayed measurements,” *OSA* **32**, 6347–6358 (Nov. 1993).
- [12] Le Roux, B., Conan, J.-M., Kulcsár, C., Raynaud, H.-F., Mugnier, L. M., and Fusco, T., “Optimal control law for classical and multiconjugate adaptive optics,” *Journal of the Optical Society of America A* **21**, 1261–1276 (July 2004).
- [13] Looze, D. P., “Linear-quadratic-gaussian control for adaptive optics systems using a hybrid model,” *J. Opt. Soc. Am. A* **26**, 1–9 (Jan 2009).
- [14] Gray, M., Petit, C., Rodionov, S., Bocquet, M., Bertino, L., Ferrari, M., and Fusco, T., “Local ensemble transform Kalman filter, a fast non-stationary control law for adaptive optics on ELTs: theoretical aspects and first simulation results,” *Optics Express* **22**, 20894 (Aug. 2014).
- [15] Massioni, P., Kulcsár, C., Raynaud, H.-F., and Conan, J.-M., “Fast computation of an optimal controller for large-scale adaptive optics,” *Journal of the Optical Society of America A* **28**, 2298 (Nov. 2011).
- [16] Petit, C., Conan, J.-M., Kulcsár, C., and Raynaud, H.-F., “Linear quadratic Gaussian control for adaptive optics and multiconjugate adaptive optics: experimental and numerical analysis,” *Journal of the Optical Society of America A* **26**, 1307 (May 2009).
- [17] Correia, C. M., Jackson, K., Véran, J.-P., Andersen, D., Lardièrre, O., and Bradley, C., “Spatio-angular minimum-variance tomographic controller for multi-object adaptive-optics systems,” *OSA* **54**, 5281 (June 2015).
- [18] Hinnen, K., Verhaegen, M., and Doelman, N., “Exploiting the spatiotemporal correlation in adaptive optics using data-driven h2-optimal control,” *J. Opt. Soc. Am. A* **24**, 1714–1725 (Jun 2007).
- [19] Doelman, N. J., Hinnen, K. J. G., Stoffelen, F. J. G., and Verhaegen, M. H., “Optimal control strategy to reduce the temporal wavefront error in ao systems,” *Proc.SPIE* **5490**, 5490 – 5490 – 12 (2004).
- [20] Doelman, N., Fraanje, R., Houtzager, I., and Verhaegen, M., “Adaptive and real-time optimal control for adaptive optics systems,” *European Journal of Control* **15**(3), 480 – 488 (2009).
- [21] Assémat, F., Wilson, R. W., and Gendron, E., “Method for simulating infinitely long and non stationary phase screens with optimized memory storage,” *Opt. Express* **14**, 988–999 (Feb 2006).
- [22] Doelman, N. and Osborn, J., “Modelling and prediction of non-stationary optical turbulence behaviour,” *Proc.SPIE* **9909**, 9909 – 9909 – 8 (2016).
- [23] van der Hoven, I., “Power Spectrum of Horizontal Wind Speed in the Frequency Range from 0.0007 to 900 Cycles Per Hour.,” *Journal of Atmospheric Sciences* **14**, 160–164 (Apr. 1957).

- [24] Deodatis, G., Ellingwood, B., and Frangopol, D., [*Safety, Reliability, Risk and Life-Cycle Performance of Structures and Infrastructures*], CRC Press (2014).
- [25] MacMynowski, D. G., Blaurock, C., Angeli, G. Z., and Vogiatzis, K., “Modeling wind-buffeting of the thirty meter telescope,” *Proc.SPIE* **6271**, 6271 – 6271 – 12 (2006).
- [26] Harrington, P. M. and Welsh, B. M., “Frequency-domain analysis of an adaptive optical system’s temporal response,” *Optical Engineering* **33**(7), 2336–2342 (1994).