



Universiteit
Leiden
The Netherlands

Predicting the future: Predictive control for astronomical adaptive optics

Kooten M.A.M. van

Citation

Predicting the future: Predictive control for astronomical adaptive optics. (2020, November 4).
Predicting the future: Predictive control for astronomical adaptive optics. Retrieved from
<https://hdl.handle.net/1887/138080>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/138080>

Note: To cite this publication please use the final published version (if applicable).

Cover Page



Universiteit Leiden



The handle <http://hdl.handle.net/1887/138080> holds various files of this Leiden University dissertation.

Author: Kooten, M.A.M. van

Title: Predicting the future: Predictive control for astronomical adaptive optics

Issue Date: 2020-11-04

Chapter 4

Implications for contrast as a result of the wind vector and non-stationary turbulence

“It’s the job that’s never started as takes
longest to finish”

J.R.R., *The Fellowship of the Ring*

Summary

In an adaptive optics (AO) system, finite time delays contribute a relatively large amount to the overall wavefront error. The temporal error, due to time delays, adds to the speckle halo which is visible in the science plane of a high-contrast imaging (HCI) instrument. The halo limits the achievable contrast near the inner working angle, preventing HCI from probing regions where we expect Earth-like planets. Our research focuses on understanding the non-stationary behavior of atmospheric turbulence and its effects on the temporal error and the contrast.

The wind vector greatly influences the turbulence-induced wavefront phase error. In the framework of the von Karman power spectral density (PSD) of wavefront phase, the wind speed changes the cut-off frequency as well as the gain of the temporal PSD. We aim to understand the behavior of the wind vector on sub-second timescales using data from the Thirty Meter Telescope site testing campaign at Mauna Kea. The nocturnal measurements we study are from 287 nights throughout the years of 2006-2008, taken at a height of 7 m with a maximum sampling rate of 60 Hz. We present an estimated PSD from the observed data to describe the wind speed located at Area E (northern plateau) on Mauna Kea. We compare the estimated PSD to the Kaimal PSD—a parameterised wind model for the lower atmospheric boundary layer. Studying the estimated PSD, we fit a fractionally integrated auto-regressive–moving-average model to the data. With this model we can create an artificial wind speed time series to be fed into phase screen simulations to generate realistic wind-varying phase screens. We examine how the contrast, achieved by

This chapter is published as: Maaïke van Kooten et al. (2018). “Implications for contrast as a result of the wind vector and non-stationary turbulence”. In: *Adaptive Optics Systems VI*. vol. 10703. International Society for Optics and Photonics, p. 107032C

the AO system, changes when subject to non-stationary turbulence induced by wind fluctuations for an integrator controller.

4.1 Introduction

In ground-based high-contrast imaging (HCI), current contrast levels — especially close into the host star — are limited by the performance of the adaptive optics (AO) system (Kasper, 2012). Specifically, by the finite time delay in the control loop, leading to the servo-lag error, that results in the so-called wind-driven halo seen by instruments such as SPHERE (Milli et al., 2017). Even after post-processing techniques such as angular differential imaging (ADI), the wind driven halo is still present as described by Milli et al. (2017). Efforts to improve control techniques to minimize the servo-lag error, and thereby the wind-driven halo, are on-going (Paschall and Anderson, 1993; Le Roux et al., 2004; Looze, 2009; Gray et al., 2014; Massioni et al., 2011; Correia et al., 2015; Doelman, Fraanje, et al., 2009; Hinnen et al., 2007).

In designing new controllers, the question arises as to how accurately input disturbances need to be modeled. Many AO simulation tools rely on performing a discrete Fourier transform (DFT) of a power spectral density (PSD) which gives the covariance function of the phase distortion due to atmospheric turbulence. For a temporally evolving atmosphere, a large phase screen (the length of which is the telescope diameter multiplied by the number of iterations to step through) is typically generated once and then shifted at a fixed rate simulating Taylor’s Frozen Flow Hypothesis (i.e, the wind blowing the turbulence across the telescope aperture). The principle can be used to generate a multi-layered atmosphere that is fitted to fixed C_n^2 and wind vector profiles. These methods provide a good basis for many AO systems and allow for relatively quick simulations that are representative of the atmosphere. However, depending on the temporal behavior of atmospheric parameters such as the outer scale, Fried parameter, and the wind vector, which we know evolve on timescales of minutes (or less; Ziad (2016)), these methods might not be adequate to simulate the performance for the next generation of large telescopes, especially for optimal performance or very high contrasts.

In this paper, we focus solely on the behavior of the wind vector and how it affects current adaptive optics systems. Within the context of HCI, we aim to understand how varying wind speed contributes to the wind-driven halo and therefore the contrast. The main objective of this paper is to answer the following question: to what extent do instantaneous wind fluctuations affect contrast? To answer this overarching question, we discuss the following questions:

1. How does wind vary during an observation?
2. How do wind fluctuations influence the contrast?

By answering these research questions, we can determine whether more efforts are needed to characterize wind at high temporal frequencies as well as use this information to dictate what control schemes could be used in place of an integrator to minimize the wind-driven halo.

4.1.1 Servo-lag error and wavefront sensor integration

To answer the question of how wind fluctuations affect the AO system, we can look at how wind speed affects the AO system by relying on the Frozen Flow hypothesis. In Hardy (1998), the servo-lag error is expressed as $\alpha_{servo}^2 = (\frac{T_{control}}{\tau_0})^{5/3}$, with $T_{control}$ being the closed-loop control time and τ_0 being the time constant (which depends on the wind speed). Note that a large amount of the control time is the wavefront sensor (WFS) integration time. Therefore, as the wind speed increases we can keep our servo-lag error constant by decreasing the WFS integration time. However, the WFS integration time also determines the signal-to-noise ratio of the WFS and therefore the measurement error. The optimal integration time must be found by minimizing the summed error term of the servo-lag error and the measurement error.

4.1.2 Outline

The first part of the paper focuses on understanding how the wind varies on sub-second timescales and how to model these fluctuations. In Section 4.2 we present the behavior of the observed wind vector. In Section 4.3 we estimate the wind speed PSD, compare it to PSDs from literature, and determine a corresponding stochastic model of the wind speed. The second part of the paper focuses on wind fluctuations and contrast. In Section 4.4, we look at the effects of varying wind on the contrast.

4.2 Observed Wind Vector Data

4.2.1 Overview of data

The data used in this work are from the Thirty Meter Telescope (TMT) site testing campaign at Mauna Kea, Hawaii taken between 2006 and 2008 by Schock et al. (2009). The data consist of raw sonic anemometer wind speed values taken at 7 m height, with a maximum sampling rate of 60 Hz. The wind speed is measured in three directions with a range of ± 60 m/s for the horizontal wind directions u_x , and u_y . A total of 287 nights (after the removal of flagged data, nights that were only a few minutes, as well as nights with points outside the measurement range) with at least one hour of data remaining for analysis are used in this work. See Table 4.1 for the exact nights used for this work.

4.2.2 Wind statistics

Nightly

We take the u_x and u_y wind speed measurements to determine the horizontal wind speed and wind direction. Looking at each night in the data set, we see two different trends for the wind speed: constant mean and variance (nearly stationary) and slowly changing mean wind speed (Figure 4.1). For nearly stationary nights, the variance of the wind speed is large. For other cases, the overall change in wind speed can be

Table 4.1: Data from the site testing group at TMT for Hawaii (Schock et al., 2009) having at least one hour of valid wind data.

Year	Month	Day	Total Nights
2006	July	18-27	10
	August	10-12	3
	October	23-31	7
	November	1,3-11,13,15-29	25
	December	1-13	13
2007	May	11-28,30-31	20
	June	1-8,10,12-30	28
	July	1-19,21-31	30
	August	1,3-13,16-31	28
	September	1-17,18-30	29
	October	1-2,4-6,8-14,17-20,24,26,28-31	21
	November	2-3,5-6,9-20,22,24,27	17
2008	March	22-24,29-31	6
	April	1-6,10-17,21-24,26-28	21
	May	1,3-5,7-21,23-31	27
	June	1	1

large for a night. The variance of the wind speed is, however, small for short time frames.

We also observe two different trends in wind direction: very constant wind direction or random fluctuations. These random fluctuations occur for slower wind speeds, though are still pronounced even when wind speeds less than 1 m/s are filtered out (therefore it is not a consequence of reduced certainty at low wind speeds). From Mahrt (2011), we expect these wind direction fluctuations at lower wind speeds during the night, although what causes them is not fully understood within the field of atmospheric boundary layer science. The large fluctuations seen at slow wind speeds only occur a few nights out of the data set and are not representative of behavior of the wind vector.

Long term statistics

Figure 4.2 shows a histogram (for the entire data set) of mean wind speed per night. Further analysis shows that the mean wind speed does not vary much on a month to month basis. The entire data set has a mean wind speed of 6.55 m/s. This mean value is in good agreement with the value of 5.70 m/s published by Schock et al. (2009), with differences in the inclusion/exclusion of data leading to our value.

The wind direction is more variable on monthly timescales. Although the dominant direction for higher wind speeds is from the east, on a monthly basis we see other wind directions that dominate at lower wind speeds (Figure 4.4).




Figure 4.1: Wind speed for the entire nights of 11-07-2007 (top) and 31-03-2008 (bottom).




Figure 4.2: Histogram of mean wind speed per night for the entire data set.




Figure 4.3: Histogram of the nightly wind speed variance over the entire data set.




Figure 4.4: Wind direction during the months of 12-2006 (top) and 07-2007 (bottom). The colors indicate the wind speed in m/s

We make no further attempt to distinguish any long-term behavior or seasonal behavior of the wind vector due to the relatively limited size of our data set on such timescales.

4.3 Estimation of Wind Speed Power Spectral Density

We estimate the PSD of the wind speed using the three years of data summarised in Table 4.1. First, we estimate the PSD for each night by splitting the data into a time series of 3000 points. The time series are detrended and the mean is subtracted. We then estimate the PSD using the Lomb-Scargle periodogram method, which allows for non-uniformly sampled data. All estimated PSDs are then averaged together for one given night. We then average the PSDs of all nights to generate our final PSD (Figure 4.5). We limit our PSD estimation to frequencies less than 10 Hz to avoid the peak at 20 Hz due to artifacts from the instrument itself.

In logarithmic space, the PSD shown in Figure 4.5 falls-off approximately linearly after the cut-off frequency of 0.03 Hz. There is also a slight inflection at higher frequencies (above 2 Hz). The shape is consistent with energy/turbulent production at low frequencies injecting power which peaks and then falls off in the inertial subrange, eventually dissipating at higher frequencies (Pablo et al., 2004). We fit a power law to the inertial subrange, finding a slope of -1.53.




Figure 4.5: Estimated wind speed PSD for the 287 nights of data.

4.3.1 Comparison to other wind speed PSD models

Models in the literature for the PSD of wind speed are based on the following form:

$$nS(n)/\mu_f^2 = \frac{a}{(1 + bf)^c} \quad (4.1)$$

with a , b , and c being real-valued coefficients, n being normalized frequency and $S(n)$ being the PSD.

It has been shown that for minute to second timescales the slope of the wind speed PSD is dependent on geographical features (Hiriart et al., 2001). Therefore we compare our normalized PSD to three variations, the Davenport, a modified Kaimal, and best-fit by Pablo et al. (2004) (we will refer to this model as A&A from now on). In Figure 4.6 we see that the overall shape and slope of the different distributions are similar to our estimated PSD. It is important to note that the Davenport, Kaimal, and A&A curves depend on the friction velocity, μ_f , due to the amount of shearing that occurs in the boundary layer (atmospheric boundary layer is defined as the first few hundred meters depending on location and time). This is not necessarily applicable to the wind measurements taken at Mauna Kea at an altitude of 4 km .




Figure 4.6: Estimated PSD compared to the Davenport, modified Kaimal, and A&A best-fit (Hiriart et al., 2001).

4.3.2 Stochastic Model for wind speed

One of the goals of this paper is to determine a dynamic model to allow for the modeling of realistic wind variations for a variety of simulation lengths. We therefore focus on matching a model to the observed PSD for the full frequency range between 0.002 Hz and 10 Hz. To match the behavior described above, as well as ensure the model does not have a large number of coefficients, we look toward fitting an autoregressive fractionally integrated moving average stochastic model (i.e., an ARFIMA) as done by Doelman and Osborn (2016) for the Fried parameter. Noticing that our PSD differs at low frequencies, we expand the model to capture the cut-off, we further use a low order autoregressive tempered fractionally integrated moving average stochastic model (an ARTFIMA) as described by Sabzikar et al. (2015).

The PSD of an ARTFIMA(p, λ, f, q) is given by

$$\Phi(k) = \frac{|\Theta(e^{-ik})|^2}{|\Phi(e^{-ik})|^2} |1 - e^{-\lambda - ik}|^{-2d} \quad (4.2)$$

with $\Theta(z) = 1 + \theta_1 z + \dots + \theta_q z^q$ and $\Phi(z) = 1 - \phi_1 z - \dots - \phi_p z^p$ being the MA and AR parts of the model, respectively and k is the spatial frequency.

We find a lower order ARTFIMA(0, $\lambda, f, 0$), Equation 4.3, to be a good fit to our data (R-squared = 0.999) for frequencies higher than 0.05 Hz, as shown in Figure 4.7. Using this model, together with the histogram of wind speed variance per night (Figure 4.3), and the histogram of mean wind speed per night (Figure 4.2), we can simulate a variety of realistic wind conditions.

$$\Phi(k) = |1 - e^{0.02 - ik}|^{-2(0.91)} \quad (4.3)$$

Ongoing research is being conducted to further improve the fit and capture the correct behavior of the estimated PSD at the low frequencies, and not just the slope for higher frequencies.

4.3.3 Implementation

Using the stochastic model above along with the histogram of the nightly wind speed variances, we have presented a framework to allow for the modeling of wind speed fluctuations at the ground layer. By combining this model with turbulent phase generation environments such as the infinitely long and non-stationary phase screen method by Assémat et al. (2006) or the autoregressive method by Srinath et al. (2015), we can generate non-stationary phase screens. Furthermore, we can expand our non-stationary turbulence model to include Fried parameter fluctuations as found by Doelman and Osborn (2016) and test a variety of AO configurations.

Although we briefly investigate the behavior of wind direction fluctuations in Section 4.2, we do not attempt to model or determine the PSD of the wind direction fluctuations as the focus of the remainder of this paper is on contrast and wind speed in the framework of an AO system using an integral controller. We assume that the




Figure 4.7: Top panel shows the ARTFIMA fit to the estimated PSD. The relative PSD (the estimated PSD multiplied by the inverse of the fitted PSD) showing the error between the fitted ARTFIMA model and the estimated PSD from the data is shown in the bottom panel.

wind direction would not be a large factor in the behavior of the system due to the small nature of the wind fluctuations.

4.4 Implication of wind speed PSD on contrast

4.4.1 Simulation environment

To investigate the influence of variable wind speed on the contrast of an HCI system, we first simulate the residual wavefront error from an AO system and then propagate the phase error through a perfect coronagraph after which we can calculate the raw contrast (Jensen-Clem et al., 2018) from the coronagraphic point-spread-function (psf) using:

$$C_{raw} = \frac{\text{mean}[R(x, y)]}{\text{max}[PSF_{star}(x, y)]} \quad (4.4)$$

with $R(x,y)$ being the intensity in a region of interest. We focus on the behavior of a classical AO system assuming a single deformable mirror and a Shack-Hartmann WFS using a 4-quadrant detector centroiding algorithm. The controller is an integrator in which the current phase detected by the WFS is taken as the best representation of the wavefront to be corrected. Instead of performing an end-to-end simulation, as we only want residual wavefront errors for a variety of wind speeds, we use the analytical based AO toolbox, PAOLA, written in IDL (Jolissaint et al., 2006). The code uses a PSD-based approach in which the whole system is modeled together using an analytical expression for the residual phase PSD. The code allows for the user to define a deformable mirror (idealized, laboratory measured transfer function, or user defined). It also allows for the implementation of many pupil geometries. Finally, it allows for a ground layer natural guide star AO mode, which is used for this work.

Our WFS is sensing in the R-band and the natural guide star magnitude $m_r = 10$. The computational time is set to 7.5 ms; in this way to ensure the servo-lag error is the largest contributor to the residual wavefront error (overall latency including the WFS integration time). The deformable mirror was chosen such that the actuator grid samples the Fried parameter 4 times, guaranteeing a low fitting error. See Table 4.2 for list of input parameters in our simulation. Table 4.3 shows the residual root-mean-square (RMS) wavefront error where we can clearly see that the servo-lag and WFS noise are the largest contributions to the error budget.

We take the residual PSD found by PAOLA and create a phase screen which is propagated through a perfect coronagraph. We then propagate the wavefront to an image plane and determine the coronagraphic psf. We do this for many realizations and sum all the coronagraphic images to generate a long-term exposure image. We calculate the raw contrast according to Equation 4.4, the coronagraphic image is divided by the maximum of the summed non-coronagraphic images (determined by including the phase aberrations to take into account the Strehl ratio of the AO system). We average the contrast over an $8 \times 8 \lambda/D$ region of interest. We find for a perfectly flat wavefront (no aberrations) our coronagraph has a contrast of 10^{-24} , showing that our coronagraph is indeed perfect given our numerical accuracy.

Table 4.2: PAOLA input parameters for natural guide star ground layer adaptive optics.

Parameter	Value
Telescope Diameter	8 m
Fried parameter @ 500nm	0.144 m
Outer Scale	30 m
Wavelength band WFS NGS	R
Control system time lag	7.5 ms
WFS integration time	3-4 ms

To simulate a long-term exposure during a period with a given time series of wind fluctuations, we first take our wind fluctuations and generate a histogram to determine the probability of each of the wind speeds. We determine the optimized gain

and WFS integration time (an iterative analytic process done by PAOLA to minimize the measurement error and servo-lag error) for a given reference AO system with a specific reference wind speed. We then determine the residual PSD for each of the wind speeds in our histogram using the WFS integration time and gain determined for the reference wind speed. Finally, we take all the residual PSDs and generate many random phase screens and perform a weighted sum of images of which we can determine the raw contrast.

4.4.2 Simulations

We focus on the wind behavior of two nights, 11 July 2007 and 31 March 2008 (Figure 4.1). Taking the first 10 minutes of the time series, we find an average wind speed of 8.03 m/s and 12.25 m/s for the two nights, respectively. We use this information to calibrate our AO system, and thereby fix the controller gain and WFS integration time.

Table 4.3: Error budget for the mean wind speed of entire dataset 6.55 m/s

Error source	RMS of residual wavefront
Fitting	28.180 nm
Aliasing	2.676 nm
Servo-lag	68.454 nm
WFS noise	59.426 nm
Total	94.967 nm

For each of the nights we take two different 2.5 second segments of the wind speed fluctuations: one with a mean wind speed close to the beginning of the night, and one with a mean wind speed much larger. Examples are shown in Figure 4.8. We run the AO system for each wind speed in the time series with the reference AO




Figure 4.8: Wind speed during the nights of 2.5 seconds taken from 11-07-2007 (left) and 31-03-2008 (right).

setup and find the short-term exposure. We stack all the short term exposures to form a the long term exposure. This generates a science image that was taken with a fix AO gain while the wind speed was fluctuating as in Figure 4.8. We determine the contrast from long term exposure and compare it to the reference contrast found for the beginning of the night in when we assume the wind had not been fluctuating. Figures 4.9 and 4.10 show the results. We then repeat the simulations but set-up the AO system parameters to be optimal for the mean wind speed of the entire data set, 6.55m/s. Finally, we reference the high wind speed simulations to their mean wind speed values.

4.4.3 Simulation results

The results of the simulations are summarized in Figure 4.11. The dashed black line indicates the idealized dependence on $v^{2/3}$ from Guyon (2005) for the case where the servo-lag error consists only of the WFS integration time, and the controller has a unity gain. The black dots indicate the limit of our simulation by optimizing our AO system for both v_1 and v_2 (the reference wind speed and the changed wind speed, respectively). An analytical expression for this can be derived from Jolissaint et al. (2006) wavefront error expressions. Comparing the simulation limited values to the theoretical scaling, we see that the closer we optimize our AO system to the current mean wind speed (v_2/v_1 is close to 1), the closer the contrast is to the theoretical maximum (best) performance. Even when the system is optimized for a slightly higher wind speed than observed ($v_2 < v_1$), the contrast is not significantly impacted. However, when we are far away from our actual on-sky wind speed ($v_2/v_1 > 1.5$), we lose contrast dramatically; much faster than the idealized $2/3$ power scaling.

The blue square points in Figure 4.11 are for simulations that have a constant (static) wind speed. The red triangles are the results from simulations with variable wind speed fluctuating about a mean speed (v_2) as shown in Figure 4.1. Overall, there is little difference between the static and fluctuating wind results. The results of both the static and varying wind simulations show contrast differences of no more than 5%, leading to the conclusion that the small fluctuations have no impact on the contrast within our simulation environment.

In order to achieve the best contrast possible for an integral controller, our AO system should be operating in the regime where the wind speed used for tuning the system is equal to the current on-sky wind speed, and therefore the wind speed needs to be tracked throughout the night, and the system re-tuned as needed. The improvement in contrast from tracking the wind speed is illustrated by the pentagons in Figure 4.11, which shows the improvement in contrast for specific 2.5 second period (test segment), on the night of 2007 as shown in Figure 4.1. We calibrated the AO system for different wind speeds. First, we tune the AO system to the overall mean wind from the three years of site testing data (6.55 m/s) and then run the AO system on the test segment, the resulting contrast is shown by point 1 in Figure 4.11. Then, the mean from the first ten minutes (point 2) of the night is used to tune the




Figure 4.9: Coronagraphic images for reference wind speed of 8.03 m/s (top) and random wind fluctuations around a mean wind speed of 16.43 m/s (bottom) corresponding to the time series on the top of Figure 4.1




Figure 4.10: Coronagraphic images for reference wind speed of 12.25 m/s (top) and random wind fluctuations around a mean wind speed of 15.80 m/s (bottom) corresponding to the time series on the bottom of Figure 4.1.




Figure 4.11: Contrast loss as a function of change in wind speed. The black dashed line indicates the theoretical limit. Various simulation runs are shown by the markers. The blue dashed line connects the outlined red points from 1 to 5 which indicate the improvement in performance when the reference wind velocity used to tune the AO system becomes closer to the mean wind speed of the 2.5 second time series from the night in 2007.

AO. Next, the AO was tuned to the 2 minute mean (point 3) before the test segment. Finally, the AO system is tuned to the mean of the wind speed 2 seconds before the test segment (point 4). Point 5 shows when the instantaneous wind speed is tracked. As the mean wind speed used to tune the AO system gets closer to the actual on-sky wind speed, the contrast for the controller is optimized.

4.5 Conclusions

Using data from the TMT site testing campaign, we find a PSD function to describe the wind speed fluctuations at Mauna Kea at 7 *m*. We model the fluctuations with an ARTFIMA(0,-0.02,0.90,0) model.

We study the effects of wind speed on the contrast. We find that sub-second wind speed fluctuations affect the contrast by less than 5% and would therefore not significantly reduce the contrast for current HCI. However, changes in wind speed on longer timescales do have a significant impact on the contrast, particularly when the wind becomes faster than what was expected given the AO controller gain and WFS integration time. We show that running an integral controller optimized for a different wind speed than observed at the time of correction leads to a much larger

loss in contrast than predicted by the fundamental limitations. To prevent this large loss in contrast, the AO system needs to track the wind speed at rates faster than 2 minutes for the given wind statistics on Mauna Kea.

References

- Assémat, François, Richard W. Wilson, and Eric Gendron (Feb. 2006). “Method for simulating infinitely long and non stationary phase screens with optimized memory storage”. In: *Opt. Express* 14.3, pp. 988–999.
- Correia, C. M., K. Jackson, J.-P. Véran, D. Andersen, O. Lardière, and C. Bradley (June 2015). “Spatio-angular minimum-variance tomographic controller for multi-object adaptive-optics systems”. In: *OSA* 54, p. 5281.
- Doelman, Niek, Rufus Fraanje, Ivo Houtzager, and Michel Verhaegen (2009). “Adaptive and Real-time Optimal Control for Adaptive Optics Systems”. In: *European Journal of Control* 15.3, pp. 480–488.
- Doelman, Niek and James Osborn (2016). “Modelling and prediction of non-stationary optical turbulence behaviour”. In: *Proc.SPIE* 9909, pp. 9909 - 9909 –8.
- Gray, M., C. Petit, S. Rodionov, M. Bocquet, L. Bertino, M. Ferrari, and T. Fusco (Aug. 2014). “Local ensemble transform Kalman filter, a fast non-stationary control law for adaptive optics on ELTs: theoretical aspects and first simulation results”. In: *Optics Express* 22, p. 20894.
- Guyon, Olivier (Aug. 2005). “Limits of Adaptive Optics for High-Contrast Imaging”. In: *The Astrophysical Journal* 629.1, pp. 592–614.
- Hardy, J. W. (1998). “Adaptive Optics for Astronomical Telescopes”. In: *Oxford University Press Jul 1998*.
- Hinnen, Karel, Michel Verhaegen, and Niek Doelman (June 2007). “Exploiting the spatiotemporal correlation in adaptive optics using data-driven H2-optimal control”. In: *J. Opt. Soc. Am. A* 24.6, pp. 1714–1725.
- Hiriart, D., J.L. Ochoa, and García, B. (Oct. 2001). “Wind Power Spectrum measured at the San Pedro Martir Sierra”. In: *American Astronomical Society* 37, pp. 213–220.
- Jensen-Clem, R. et al. (Jan. 2018). “A New Standard for Assessing the Performance of High Contrast Imaging Systems”. In: *Astronomical Journal* 155, p. 19.
- Jollissaint, Laurent, Jean-Pierre Véran, and Rodolphe Conan (Feb. 2006). “Analytical modeling of adaptive optics: foundations of the phase spatial power spectrum approach”. In: *J. Opt. Soc. Am. A* 23.2, pp. 382–394.
- Kasper, Markus (2012). “Adaptive optics for high contrast imaging”. In: *Proc.SPIE* 8447, pp. 8447 - 8447 –10.
- Kooten, Maaike van, Niek Doelman, and Matthew Kenworthy (2018). “Implications for contrast as a result of the wind vector and non-stationary turbulence”. In: *Adaptive Optics Systems VI*. Vol. 10703. International Society for Optics and Photonics, p. 107032C.

- Le Roux, B., R. Ragazzoni, C. Arcidiacono, J.-M. Conan, C. Kulcsar, and H.-F. Raynaud (Oct. 2004). “Kalman-filter-based optimal control law for star-oriented and layer-oriented multiconjugate adaptive optics”. In: *Proc.SPIE* 5490, pp. 1336–1346.
- Looze, Douglas P. (Jan. 2009). “Linear-quadratic-Gaussian control for adaptive optics systems using a hybrid model”. In: *J. Opt. Soc. Am. A* 26.1, pp. 1–9.
- Mahrt, Larry (2011). “Surface Wind Direction Variability”. In: *Journal of Applied Meteorology and Climatology* 50.1, pp. 144–152.
- Massioni, P., C. Kulcsár, H.-F. Raynaud, and J.-M. Conan (Nov. 2011). “Fast computation of an optimal controller for large-scale adaptive optics”. In: *Journal of the Optical Society of America A* 28, p. 2298.
- Milli, J. et al. (Sept. 2017). “Performance of the extreme-AO instrument VLT/SPHERE and dependence on the atmospheric conditions”. In: *AO4ELT5*.
- Pablo, Juan, Pérez Beaupuits, Angel Otárola, Fredrik T. Rantakyö, Roberto C. Rivera, Simon J. E. Radford, and Larsaake Nyman (2004). “ANALYSIS OF WIND DATA GATHERED AT CHAJNANTOR”. In: *ALMA Memo* 497.
- Paschall, R. N. and D. J. Anderson (Nov. 1993). “Linear quadratic Gaussian control of a deformable mirror adaptive optics system with time-delayed measurements”. In: *OSA* 32, pp. 6347–6358.
- Sabzikar, Farzad, Mark M. Meerschaert, and Jinghua Chen (2015). “Tempered fractional calculus”. In: *Journal of Computational Physics* 293, pp. 14–28.
- Schock, M. et al. (2009). “Thirty Meter Telescope Site Testing I: Overview”. In: *Publications of the Astronomical Society of the Pacific* 121.878, p. 384.
- Srinath, S., L. A. Poyneer, A. R. Rudy, and S. M. Ammons (Dec. 2015). “Computationally efficient autoregressive method for generating phase screens with frozen flow and turbulence in optical simulations”. In: *Optics Express* 23, p. 33335.
- Ziad, Aziz (2016). “Review of the outer scale of the atmospheric turbulence”. In: *Proc.SPIE* 9909, pp. 9909 - 9909 –12.

