



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 3

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

“Go back?” he thought. “No good at all! Go sideways? Impossible! Go forward? Only thing to do!”

J.R.R. Tolkien, *The Hobbit*

Summary

The performance of current high-contrast imaging systems is limited at small angular separations, preventing the direct imaging of Earth-size exoplanets close to their host stars. One primary cause of this limitation is the performance of the extreme adaptive optics (AO) system, which is dominated by the servo-lag error at small angular separations. Prediction of the turbulence induced optical phase can be used to help reduce this servo-lag error. Most AO systems today make use of an integral control law, resulting in a phase correction that is proportional to the current measured wavefront. This approach does not take into account the evolution of the wavefront phase fluctuations between measurement and correction. To account for the changes during the servo-lag, we focus on the implementation of different predictive control strategies that estimate how the wavefront phase fluctuations evolve over a time horizon equal to the servo-lag. Furthermore, the statistical properties of the turbulence phase are non-stationary and might have a significant impact on the performance of AO systems for high-contrast imaging systems. We therefore include non-stationary turbulence in our analysis. First, we model the non-stationary properties of the atmosphere using a von Kármán covariance function which can incorporate variations in time of the Fried parameter, outer scale, and wind velocity. We then vary the Fried parameter using an ARIMA model and simulate increases in wind speed using both ramp and step-like jumps in wind speed. Using this new disturbance model of turbulence, the performance of three different predictors – based on a steady state, recursive, and adaptive linear-minimum-mean-square-estimator (LMMSE) – is de-

terminated under non-stationary conditions. We present the 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. From these tests, we conclude that a more in-depth study into non-stationary modelling is needed for atmospheric turbulence.

3.1 Introduction

Since the first direct images of an exoplanet by Marois et al. (2008), high-contrast imaging (HCI) has been a very active field. Capable of providing insight into planet formation, characterizing exoplanets' atmospheres, 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 SPHERE (Beuzit et al., 2008) on the VLT, GPI (Macintosh et al., 2014) on Gemini South, and SCExAO (Jovanovic et al., 2015) 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 $0.5''$ - $1''$ separation (see Kasper, 2012, for more details) in the near infrared as shown in Figure 3.1.

To answer the fundamental questions of: is the earth unique? is our solar system unique? and how did we form? we, however, need to aim for contrast levels of $10^{-7} - 10^{-9}$ at separations of $0.05'' - 1''$ (for visible wavelengths) to find cold jupiters, neptunes, and rocky exoplanets. Currently, the limiting factor is the correction achieved by the adaptive optics (AO) systems within the HCI systems. Specif-




Figure 3.1: Current and future high contrast imaging capabilities from Otten et al. (2017) for various instruments and coronagraphs.

ically, the chromatic index and servo-lag error prevent us from observing at small separations as shown in Figure 3.2. Note that the fitting error can be reduced by increasing the number of actuators for the deformable mirror (DM). We focus on minimizing the servo-lag error. One approach is to run our AO system faster, requiring high speed DMs, fast wavefront sensors, and bright guide stars. Another approach is to focus on the evolution of wavefront phase fluctuations during the delay and correct for it using predictive control methods.

3.1.1 Predictive control in adaptive optics

As early as the 1970s (see Hardy, 1998), predictive control has been proposed to mitigate the effect of inevitable time delays in AO 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; see Petit, Conan, Kulcsár, Raynaud, and Fusco, 2008; Doelman, Fraanje, and Breeje, 2011; Poyneer et al., 2016). Paschall and Anderson (1993), Le Roux et al. (2004), and Looze (2009) have all worked on the LQG controller for full atmospheric disturbance. Different varieties of the Kalman filter predictor have been implemented by Gray et al. (2014) and Massioni et al. (2011), and allow for filtering in real-time. For LGQ control, Petit, Conan, Kulcsár, and Raynaud (2009) have done laboratory testing while Correia et al. (2015) have performed on-sky tests. The data-driven H_2 optimal control implemented by Hinnen et al. (2007)




Figure 3.2: An example of an adaptive optics error budget showing that the servo-lag error is a major limitation of achievable contrast from Kasper (2012) for small angular separations. The fitting error can be reduced by increasing actuator density, and chromatic index error is reduced by sensing close to the science imaging wavelength.

was also tested on-sky and Doelman, Hinnen, et al. (2004) report a reduction in the temporal error. Finally, adaptive predictors have been proposed by Doelman, Fraanje, Houtzager, et al. (2009) to track the time-varying behavior of atmospheric turbulence.

Similar to any control problem, the characteristics of the (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 an improved 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 3.2. The statistical behavior and the non-stationarities of turbulence are described in Section 3.3. We present, in Section 3.4, the results of our AO simulations with our new turbulence model and we discuss the impact of the results on predictor design. Finally, we outline our on-going and future work in Section 3.5.

3.2 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:

$$\hat{y}_i(t + 1) = \vec{A}_i \cdot \vec{w}(t) \quad (3.1)$$

$$\vec{w}(t) = (y_0(t) \ y_1(t) \ y_2(t) \ \dots \ y_P(t))^T \quad (3.2)$$

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

$$\min_{\vec{A}_i} < \|\vec{y}_i(t) - \vec{A}_i \vec{w}(t)\|^2 > \quad (3.3)$$

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

$$\vec{W} = (\vec{w}(t) \quad \vec{w}(t-1) \quad \vec{w}(t-2) \quad \dots \quad \vec{w}(t-Q)) \quad (3.4)$$

and

$$\vec{Y} = (y_i(t+1) \quad y_i(t) \quad y_i(t-1) \quad \dots \quad y_i(t-(Q+1)))^T \quad (3.5)$$

By substituting $\vec{W}$ and $\vec{Y}$ for $\vec{w}$ and $\vec{y}$ in Equation 3.3, respectively, we find the prediction vector to be:

$$\vec{A}_i = \vec{W}^+ \vec{Y} \quad (3.6)$$

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 auto-covariance matrix and the cross-covariance vector.

$$\vec{A}_i = \mathbf{C}_{\vec{w}\vec{w}}^+ \vec{C}_{\vec{w}\hat{y}} \quad (3.7)$$

The batch prediction vector is found, it is fixed for the rest of the simulation. By making use of the matrix inversion lemma, we can also form a recursive solution in which the covariance matrices are updated at each time step (Equations 3.8 and 3.9), removing the need for a pseudo inverse calculation (greatly reducing the computational load).

$$\mathbf{C}_{\vec{w}\vec{w}}^{-1}(t) = \mathbf{C}_{\vec{w}\vec{w}}^{-1}(t-1) - \frac{\mathbf{C}_{\vec{w}\vec{w}}^{-1}(t-1)\vec{w}(t-1)\vec{w}^T(t-1)\mathbf{C}_{\vec{w}\vec{w}}^{-1}(t-1)}{1 + \vec{w}^T(t-1)\mathbf{C}_{\vec{w}\vec{w}}^{-1}(t-1)\vec{w}(t-1)} \quad (3.8)$$

$$\vec{C}_{\vec{w}\hat{y}}(t) = \vec{C}_{\vec{w}\hat{y}}(t-1) + \hat{y}(t-1)\vec{w}^T(t-1) \quad (3.9)$$

The recursive solution begins 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. To start, 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 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, we limit the number of previous phase screens to one. We further limit it such that prediction at a given phase point only uses neighboring phase points (within a radius of 5 grid points) in the prediction




Figure 3.3: The phase points at $t=N$ to use to predict the point at $t=N+1$.

algorithm (see Figure 3.3). This allows us to reduce the computation time as well as study how the LMMSE behaves for different amounts of input information.

3.3 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 (r_0), and outer scale (L_0) can be used to describe the spatial phase variations. We consider a spatial PSD and then derive the temporal PSD by applying Taylor's Frozen Flow hypothesis which involves knowing the wind speed, v , and direction (see Section 3.3.3). 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 Kármán turbulence model by implementing an infinite phase screen algorithm developed by Assémat et al. (2006) that allows for the new rows of atmosphere to be generated online. Their use of a von Kármán covariance function to generate the phase allows for v and r_0 to be changed at each time step of the simulation. In Figure 3.4 we show the temporal PSD for three different cases: the reference case, the case where v is increased/decreased by 15 m/s, and the case when L_0 is increased/decreased by 30 m. Wind speed has a large effect on the PSD shape — if the shape of the PSD were to change like this in real life, we would have non-stationary turbulence. We do not, however, focus on L_0 as little information is known on its fluctuations and its affect on the temporal PSD is much less than that of changing v . To properly vary these parameters we need models that describe how

the ν and r_0 change in time.

3.3.1 Fried parameter variations

The Fried parameter is a measure of the strength of turbulence, with large numbers corresponding to weak turbulence and small values corresponding to strong turbulence. By definition, it is the path integral of the C_n^2 profile, and is therefore directly related to changes in the temperature fluctuations in the atmosphere. Available stereo-scidar measurements are unable to provide r_0 sampled at frequencies greater than 1 Hz. Using wavefront sensor data, however, Doelman and Osborn (2016) determined a fractional ARIMA model to describe the evolution of r_0 at La Palma on the time interval of 100 seconds. We adopt this model, creating a time series by assuming the r_0 remains constant over these 100 seconds. We feed this series into our atmospheric phase generator. We also look at a constantly decreasing r_0 , representative of how seeing might improve over the course of a night.

3.3.2 Wind variations

Although much work has been done on the variation of ν over long time scales, including the van der Hoven (1957) PSD and time evolving PSD for applications in civil engineering by Deodatis et al. (2014), we were unable to find an appropriate model useful for our purposes. Moreover, the slope of the PSD changes significantly depending on the altitude and geometry of the location. Also, the behavior of wind




Figure 3.4: Temporal von Kármán PSD for three different cases: a reference PSD, a change in wind speed by 15 m/s, and finally a change in 30 metres for L_0 .

varies from day time to night time — many wind PSDs are for day time only. Observations of wind (minute time scales) in flat agricultural areas (at the CESAR observatory in the Netherlands) 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 effects such as dome geometry and the telescope movement that contribute to the relative v 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 (Pablo et al., 2004) to generate the small wind fluctuations caused by turbulence and assume the atmosphere is stationary for one second (see Deodatis et al., 2014). We then change the mean wind speed using a probability density function determined by MacMynowski et al. (2006) 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 v fluctuations, we are subjecting the predictor to small variations and large variations in one time step as well as using v that are commonly seen at observatories.

3.3.3 Effect on prediction and control

To understand how v , L_0 , and r_0 affect the disturbance model and the ability of the system to reject the disturbance, we look at the spatial and temporal von Kármán PSDs. We define the spatial PSD, for a single point and a single layer, in Equation 3.10 (neglecting the inner scale).

$$\Phi(k) = 0.0229 r_0^{-5/3} (k^2 + 4\pi^2 L_0^2)^{-11/6} \quad (3.10)$$

For discussion purposes, we write the single-sided temporal PSD (Harrington and Welsh, 1994) (Equation 3.11), substituting in $C_n^2 dh \propto r_0^{-5/3}$ to give:

$$\Phi(f) \propto \left(\frac{v}{r_0}\right)^{5/3} \left(f^2 + \frac{v^2}{L_0^2}\right)^{-4/3} \quad (3.11)$$

Note, Equation 3.11 is slightly different from the temporal PSD used in our simulations, although we too use Frozen Flow to evolve our phase screens.

Equations 3.10 and 3.11 show that when r_0 varies, both the spatial and temporal gains are changed. L_0 also affects the spatial PSD gain, as well as the cut-off frequency of the temporal PSD. Variations in v change the gain and cut-off frequency of the temporal PSD (i.e, Figure 3.4). An ideal predictor should be able to account for the variations in the cut-off frequency that are due to the changing disturbance's statistics. Time-invariant predictors are based on stationary disturbance statistics. Therefore, knowledge of the dynamical behavior of v and L_0 might be beneficial to achieve 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.

3.4 Simulations

We test the three different LMMSE predictors outlined in Section 3.2 along with a perfect integrator (zero-order predictor/no prediction) analogue in an open-loop configuration. To do this, we take the most recent open-loop measurement as the best estimate for the current phase aberrations; the gain is 1 as we do not have to worry about any response of the AO system. We simulate, in Python, 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 r_0 (two different cases),
3. and on-stationary turbulence with varying v .

The exact configurations of the input parameters for each case are summarized in Table 3.1. We run our AO loop at 500 Hz for 5 seconds, then average the root-mean-square (rms) error over 50 runs, regenerating the phase screens with the same PSD. We set $L_0 = 30$ m.

3.4.1 Results and Discussion

Stationary turbulence

We initially test the stationary turbulence model with the atmospheric parameters in Table 3.1 (Case 1). From Figure 3.5, 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 (rad) and the integrator having an rms of 0.54 rad. 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 v

Table 3.1: Summary of input parameters for different simulations.

Case	Simulation	r_0 [cm]	v [m/s]
1	Stationary turbulence	13	7
2	Non-stationary turbulence: r_0 1	13.001–13.004	7
3	Non-stationary turbulence: r_0 2	5–25	7
4	Non-stationary turbulence: v	13	4–7




Figure 3.5: Case 1: 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. The first second is the batch training period. The recursive predictor takes tens of iterations to converge which is not visible in the plot.

reduces, the integrator and LMMSE become closer in performance (not shown). 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 (not shown). 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. Due, in part, to strong spatial correlation of phase between surrounding points. Therefore, the LMMSE is unable to distinguish an 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.

Non-stationary turbulence

We then vary r_0 for cases 2 and 3 in Table 3.1. Figure 3.6 shows that as r_0 varies only slightly (0.03% over 5 seconds) on such short time scales the mean performance of the predictors does not change, as expected. However, such small variations in r_0 likely cannot be seen due to the injected measurement noise and the pixel scaling. Therefore, we also simulated large artificial ramps for r_0 as in case 3. In Figure 3.7, we show the results of changing r_0 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; the decrease 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 3.8. 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 r_0 . The decrease in rms followed by the increase in rms seen in Figure 3.7 is due to the overall decrease in variance of the phase screen as the coherence length increases, making the phase screens more uniform. Ultimately, the Fried parameter simulations show that the regular fluctuations of r_0 do not change the performance, and that, on time scales of a few seconds, when we vary r_0 slowly the disturbance rejection remains unchanged.

We vary v according to section 3.3.2 and present the results for our case 4. In Figure 3.9, we see that wind jumps of 1 m/s can degrade the performance of the integrator and the batch and recursive LMMSE. The batch method, as expected, does not update the predictor solution once set and therefore does not see the changing statistics. The recursive LMMSE is also unable to track the step in wind speed, even though the predictor is updated at each time step. In Figure 3.9, the rms wavefront error for all the predictors increases by a factor of two, at the second wind jump.

The tracking speed of the recursive approach is dependent on how long the pre-






Figure 3.6: Case 2: r_0 varying according to an fractional ARIMA model Doelman and Osborn (2016).

Figure 3.7: Case 3: r_0 exaggerated as a ramp feature.Figure 3.8: The evolution of a prediction coefficient for 4 different phase points as the r_0 varies.

diction 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.

3.5 Conclusions

We show that a predictor’s ability to remove residual phase errors due to the inherent time delay in an AO system depends on the statistical behavior of the wavefront disturbance. We focus on understanding how the non-stationary behavior atmospheric turbulence, specifically with respect to ν and r_0 , might affect our predictors. We have described a method that can be used to model these non-stationarities and that the change in r_0 can be modelled following Doelman and Osborn (2016). We show that using this model, the variations in r_0 over short time scales do not change the performance of the integrator and 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 resulting in a factor of two increase in the rms wavefront error. Determining a realistic model of wind speed fluctuations is a crucial next step in testing and developing predictive control algorithms as it can have a large effect on the AO performance.

We are currently focusing on developing a dynamical model to represent how wind might behave above an observatory. This will be used to further test our LMMSE predictor. With this model we will also study how contrast might be affected in simulation when comparing stationary and non-stationary atmospheric induced phase fluctuations.

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.
- Beuzit, Jean-Luc et al. (2008). “SPHERE: a planet finder instrument for the VLT”. In: *Proc.SPIE* 7014, pp. 7014 - 7014 –12.
- 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.
- Deodatis, G., B.R. Ellingwood, and D.M. Frangopol (2014). *Safety, Reliability, Risk and Life-Cycle Performance of Structures and Infrastructures*. CRC Press.




Figure 3.9: Performance of the predictors for varying wind parameters.

- Doelman, N., R. Fraanje, and R. den Breeje (2011). “Real-sky adaptive optics experiments on optimal control of tip-tilt modes”. In: *2nd Conference on Adaptive Optics for Extremely Large Telescopes* 1.1.
- Doelman, Niek J., Karel J. G. Hinnen, Freek J. G. Stoffelen, and Michel H.G. Verhaegen (2004). “Optimal control strategy to reduce the temporal wavefront error in AO systems”. In: *Proc.SPIE* 5490, pp. 5490 - 5490 –12.
- 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.
- Hardy, J. W. (1998). “Adaptive Optics for Astronomical Telescopes”. In: *Oxford University Press Jul 1998*.
- Harrington, Patrick M. and Byron M. Welsh (1994). “Frequency-domain analysis of an adaptive optical system’s temporal response”. In: *Optical Engineering* 33.7, pp. 2336–2342.
- 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.
- Jovanovic, N. et al. (Sept. 2015). “The Subaru Coronagraphic Extreme Adaptive Optics System: Enabling High-Contrast Imaging on Solar-System Scales”. In: *Publications of the Astronomical Society of the Pacific* 127, p. 890.
- Kasper, Markus (2012). “Adaptive optics for high contrast imaging”. In: *Proc.SPIE* 8447, pp. 8447 - 8447 –10.
- Kooten, Maaïke van, Niek Doelman, and Matthew Kenworthy (2017). “Performance of AO predictive control in the presence of non-stationary turbulence”. In: *Instituto de Astrofísica de Canarias*.
- 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.
- Macintosh, Bruce et al. (2014). “First light of the Gemini Planet Imager”. In: 111.35, pp. 12661–12666.
- MacMynowski, Douglas G., Carl Blaurock, George Z. Angeli, and Konstantinos Vogiatis (2006). “Modeling wind-buffeting of the Thirty Meter Telescope”. In: *Proc.SPIE* 6271, pp. 6271 - 6271 –12.
- Marois, C., B. Macintosh, T. Barman, B. Zuckerman, I. Song, J. Patience, D. Lafreniere, and R. Doyon (2008). “Direct Imaging of Multiple Planets Orbiting the Star HR 8799”. In: 322.5906, pp. 1348–1352.

- 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.
- Otten, G. P. P. L. et al. (Jan. 2017). “On-sky Performance Analysis of the Vector Apodizing Phase Plate Coronagraph on MagAO/Clio2”. In: *The Astrophysical Journal* 834, p. 175.
- 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.
- Petit, C., J.-M. Conan, C. Kulcsár, and H.-F. Raynaud (May 2009). “Linear quadratic Gaussian control for adaptive optics and multiconjugate adaptive optics: experimental and numerical analysis”. In: *Journal of the Optical Society of America A* 26, p. 1307.
- Petit, C., J.-M. Conan, C. Kulcsár, H.-F. Raynaud, and T. Fusco (2008). “First laboratory validation of vibration filtering with LQG control law for Adaptive Optics”. In: *Optics Express* 16, p. 87.
- Poyneer, L. A. et al. (Jan. 2016). “Performance of the Gemini Planet Imager’s adaptive optics system”. In: *OSA* 55, p. 323.
- van der Hoven, I. (Apr. 1957). “Power Spectrum of Horizontal Wind Speed in the Frequency Range from 0.0007 to 900 Cycles Per Hour.” In: *Journal of Atmospheric Sciences* 14, pp. 160–164.