

# Direct inverse neural network control for nonlinear time-varying Adaptive Optics

Alberto Dall'Ora

Master of Science Thesis



# Direct inverse neural network control for nonlinear time-varying Adaptive Optics

MASTER OF SCIENCE THESIS

For the degree of Master of Science in Systems and Control at Delft  
University of Technology

Alberto Dall'Ora

May 11, 2023

Faculty of Mechanical, Maritime and Materials Engineering (3mE) · Delft University of  
Technology



Copyright © Delft Center for Systems and Control (DCSC)  
All rights reserved.



---

# Abstract

Micromachined membrane deformable mirrors (MMDMs) are commonly utilized in Adaptive Optics (AO) systems due to their relatively good performance and cost-effectiveness. However, these deformable mirrors often exhibit nonlinearity at high control magnitudes and a response that is dependent on external factors such as temperature and humidity. In order to overcome these nonidealities, AO controllers typically implement a linear proportional-integral closed-loop control. Nevertheless, if the model is inaccurate, multiple wavefront (WF) measurements are required, which slow down operations. To address these issues, this thesis proposes a novel approach based on the Direct Inverse Control (DIC) framework, which involves modeling and controlling the AO system using shallow neural networks. Specifically, the *specialized learning* DIC framework is employed. This approach consists of first identifying a forward model of the plant using a neural network, then placing the controller network in series with the plant one, and finally training the controller to make the overall system resemble an identity transfer function. Since the analyzed system is underdetermined, the controller loss function is augmented with a Lagrangian multiplier. This additional term also enables the regularization of the inversion process, which helps to reduce the risk of saturating actuators. The results of this study show that the proposed approach provides better modeling accuracy than benchmarks, especially in the working ranges where nonlinearities are present. As a result, it enables faster control convergence than the state-of-the-art method when generating large-phase wavefronts. Moreover, when operating online, the DIC-based method demonstrates better stability and similar tracking abilities to Recursive Least Squares. Overall, the proposed approach provides a promising solution to the challenges associated with using MMDMs in AO systems.



---

# Table of Contents

<b>Acknowledgements</b>	<b>xiii</b>
<b>1 Introduction</b>	<b>1</b>
1-1 Outline . . . . .	2
<b>2 Theoretical background</b>	<b>3</b>
2-1 Adaptive Optics . . . . .	3
2-1-1 Image formation . . . . .	3
2-1-2 Wavefront aberration . . . . .	5
2-1-3 Wavefront detection . . . . .	6
2-1-4 Wavefront correction . . . . .	8
2-2 AO identification and control . . . . .	10
2-2-1 Least Squares method . . . . .	11
2-2-2 Challenges and limitations . . . . .	14
2-2-3 Extensions . . . . .	17
2-3 Artificial Neural Networks . . . . .	20
2-3-1 Working principle . . . . .	21
2-3-2 Training . . . . .	21
2-3-3 Direct Inverse Control . . . . .	22
2-3-4 ANN in Adaptive Optics . . . . .	25
<b>3 Method proposal</b>	<b>27</b>
3-1 Adaptive Direct Inverse Control for Adaptive Optics . . . . .	27
3-1-1 Forward model identification . . . . .	27
3-1-2 Forward model inversion . . . . .	28
3-1-3 Online update of forward and inverse models . . . . .	31

---

<b>4</b>	<b>Results</b>	<b>33</b>
4-1	Experimental conditions . . . . .	33
4-1-1	Investigated test cases . . . . .	34
4-1-2	Benchmarks . . . . .	35
4-2	Wavefront prediction . . . . .	37
4-2-1	Nonlinear conditions . . . . .	37
4-2-2	Time-varying conditions . . . . .	38
4-2-3	Lack of persistent excitation . . . . .	39
4-3	Model inversion and ideal reconstruction efficiency . . . . .	39
4-4	Wavefront correction . . . . .	41
4-4-1	Nonlinear conditions . . . . .	41
4-4-2	Time-varying conditions . . . . .	45
4-5	Discussion . . . . .	49
<b>5</b>	<b>Conclusion</b>	<b>51</b>
	<b>Bibliography</b>	<b>57</b>

---

# List of Figures

2-1	Representation of the Law of Refraction [1] . . . . .	4
2-2	Spherical wavefront generated by a point source in a homogeneous medium [1] .	5
2-3	Example of object imaging through a thin lens [1] . . . . .	5
2-4	Example of an unaberrated and aberrated wavefronts when passing through a lens [1] . . . . .	6
2-5	Application of the phase conjugation principle through a deformable mirror [1] . .	7
2-6	On the left, a schematic of a Shack-Hartmann sensor [1]. WF is the incoming wavefront, MLA is the microlens array, and SP is the projected spot pattern. On the right, is an image of the centroids pattern from the SHWFS used for this project. 7	7
2-7	Graphic representation of the first 15 Zernike modes. This basis is commonly used in optics due to various interesting properties. Zernike polynomials are orthogonal on the unit disk, have simple formulations in the cartesian domain, and resemble typical optical deformations like coma and astigmatism. . . . .	9
2-8	Schematic of a Membrane deformable mirror (DM). The shape of the reflective layer is changed through electrostatic force, applied by electrodes beneath the surface. [2] . . . . .	10
2-9	Schematic of biased operation of a Membrane DM. Bi-directional operations are enabled by a slightly concave reference shape. [2] . . . . .	10
2-10	Representation of the similarity matrix $T$ and reconstruction efficiency $diag(T)$ , for different number of truncated modes $\delta$ . The gray scale represents the absolute value of the matrix, element-wise. All Zernike modes between 4 (top-left) and 39 (bottom-right) are represented. The histograms on the right depict the elements on the diagonal of the similarity matrix and range from Zernike mode 4 to 39. Truncating the lowest singular values decreases the reconstruction efficiency of highest-order modes, and enhances the cross-talking with other, off-diagonal, elements. . . . .	13
2-11	Maximum value of the control matrix $S$ for different number of truncated singular modes $\delta$ . . . . .	14
2-12	Experimental AO setup used for this thesis project. AO Lab, Delft Center for Systems and Control (DCSC), TU Delft. The deformable mirror is an OKO 96-channel MMDM with 25.4 mm aperture. The SHWFS consists of a Thorlabs $1024 \times 1024$ CMOS Camera and a 10 mm $\times$ 10 mm lens array. The laser beam has 640 nm wavelength, a diameter of 2.9 mm, and 1.2 mW power. . . . .	15

- 2-13 Modelling diagram of the DM implemented in our setup. The mirror presents a quadratic relationship between the applied voltages contained in  $\bar{\mathbf{u}}$  and the wavefront  $s$ . The OKO SDK applies the square root operator, together with scaling and shifting, to create a linear relationship between the control commands  $\mathbf{u}$  passed to the mirror and the wavefront  $s$ . . . . . 16
- 2-14 Nonlinear behavior of our AO system. Control signals from  $u_{min} = 0mV$  to  $u_{max} = 255mV$  are sent to all 96 actuators. The mirror deformation  $\|\mathbf{a} - \mathbf{a}_0\|_2$  is measured and reported here on the  $y$  axis, where  $\mathbf{a}_0$  contains the Zernike coefficients describing the wavefront when  $0mV$  is applied to the actuators. The *OKO trajectory* corresponds to a quadratic curve, which is the modeling assumption made by the OKO SDK in order to make the whole system linear. The plotted result is the mean of 10 different episodes. The standard deviation is minimal and thus barely visible. . . . . 17
- 2-15 Correlation between wavefront modes and temperature. The mirror is maintained in the reference position and the measured wavefront at calibration temperature ( $24^\circ C$ ) is taken as reference. The temperature around the setup is then increased and subsequently set back to the calibration one. Every second, a new wavefront is measured through the SH sensor while actuators are maintained in the bias position. On the left, the mirror deformation with respect to the reference wavefront is plotted. On the right, low-order Zernike modes are reported, named according to Noll's sequential indices [3]. . . . . 18
- 2-16 Closed-Loop reconstruction of a random wavefront for different temperature values. A linear matrix model  $S$  is used to control the system, using a Proportional-Integral approach. Controller's gains are respectively 1 and 0.5. When the operating temperature is different from calibration one ( $23^\circ C$ ) the correction appears more unstable and a steady-state offset arises. . . . . 19
- 2-17 General and specific architecture of a fully connected deep neural network. On the left, the contribution of input and hidden layers' neurons to the output of the network can be appreciated. When the network is fully connected as in this case, the output of each neuron is mapped to the input of all the neurons of the following layer. Trivially, the modeling capabilities of the network are directly correlated to the number of neurons and layers. On the right, instead, the way each neuron computes its output based on the multiple inputs it receives is depicted. Using a nonlinear activation function, it is possible to make the neural network models nonlinearities. In this regard, plenty of different functions can be adopted, depending on the task the network is meant to achieve. . . . . 21
- 2-18 Direct Inverse Control diagram. The controller consists of a neural network that models the inverse of the plant being controlled. When a reference signal  $\mathbf{r}$  is requested, the Artificial Neural Networks (ANN) computes a control action  $\mathbf{u}$ , which is then sent to the actuators of the system. If the network properly models the inverse of the plant, the output variable  $\mathbf{y}$  will then resemble the reference  $\mathbf{r}$ . 24
- 2-19 The convexity problem. The region on the left is the inverse image of the point on the right. The arrow represents the direction in which the mapping is learned by direct inverse modeling. The three points lying inside the inverse image are averaged by the learning procedure, yielding the vector represented by the small circle. This point is not a solution, because the inverse image is not convex. [4] . 24
- 2-20 Diagram of the training of the forward plant network. The network is placed in parallel with the plant and thus receives the same input actions  $\mathbf{u}$ . The prediction error  $e$  between the plant output  $\mathbf{y}$  and the network estimate  $\mathbf{y}_{pred}$  is backpropagated and used to update the weights of the ANN. The training can either be done offline or online, during control operations. . . . . 25

2-21	Diagram of the training of the controller network, following the <i>specialized learning</i> approach. The error between the desired and the predicted output is backpropagated through the plant network, and exploited to update the controller one. In essence, the latter is trained to be the inverse of the plant network and thus make the overall system replicate an identity function. . . . .	25
3-1	Hyperparameters tuning curves. On the left, it is plotted the training and validation loss curves for different hidden layer depths. The loss function is minimized when more than 40 nodes are used in the hidden layer. For each test case, enough training samples are provided to make the system not overfit. On the right, instead, the training and validation loss curves for different training set dimensions are depicted. The network stops overfitting when more than 4000 samples are used for training.	29
3-2	Diagram representing the offline learning architecture of the plant network. The ANN is trained through the backpropagation of the prediction error $\bar{e}$ , here represented by a green dashed line. If the training is successful, the plant network should be able to replicate the WF generation behavior of the system. . . . .	30
3-3	Diagram representing the offline learning architecture of the controller network. Exploiting the pre-trained plant network, the controller model is identified through the backpropagation of the predicted performance error $\bar{e}$ . The plant network is necessary to express the loss gradients as a function of the control actions $\mathbf{u}$ , and thus to make the controller network be able to train on the output error. During this process, the plant network is not updated and it is only used to backpropagate the error. If the training is successful, the overall system composed of the controller and plant networks should resemble an identity function. . . . .	30
3-4	Diagram representing the online learning architecture of both the plant and controller network. Once a new WF measurement is obtained, the forward model is firstly updated using the performance error $\bar{e}$ (green line). Secondly, the predicted performance error $\bar{e}$ is back-propagated through the plant network and exploited to update the inverse model (blue line). . . . .	32
4-1	Reconstruction efficiency of the proposed method, for an L2 regularizing term $\alpha = 10^{-5}$ . The trained identity model $i(\theta_c, \theta_q) = q(\theta_q, c(\theta_c))$ is inspected. Each one of the first 40 Zernike modes is fed to the system with unitary amplitude. The output of the identity model, which represents the theoretical system reconstruction of the requested wavefront in terms of Zernike magnitudes, is reported on the y-axis. Ideally, only the desired mode should be reproduced, with unitary amplitude. The histograms on the right depict the diagonal elements of the grids on the left-hand side. They represent the reconstruction efficiency regarding the desired mode only, that is, not considering the <i>cross-talking</i> between modes. The larger the regularizing gain, the more the reconstruction quality is affected, especially for high-order modes. . . . .	34
4-2	On the left, the singular values of the DM influence matrix $Q$ are plotted. The orange cutoff line represents the threshold under which the values are set to zero and thus the modes get truncated. On the right, instead, the singular gains of the DM influence matrix $Q$ are depicted. Gains are the normalized reciprocal of the singular values. The cutoff line represents the threshold after which the modes get truncated. . . . .	36
4-3	Representation of the similarity matrix $T = S \cdot Q$ and reconstruction efficiency $diag(T)$ , in the case of 4 truncated modes. The gray scale represents the absolute value of the matrix, element-wise. All Zernike modes between 3 (top-left) and 39 (bottom-right) are represented. The histograms on the right depict the elements on the diagonal of the similarity matrix and range from Zernike mode 3 to 39. . .	36

- 4-4 Wavefront prediction performance for different control magnitudes. 500 random control action sets are randomly sampled and sent to the DM. The measured WFs are then compared with the ones predicted by the proposed method and the Least Squares (LS) state-of-the-art. On the y-axis is reported the averaged wavefront prediction rms error, while the x-axis shows the maximum actuator control action used to generate the given wavefronts. In the figure, circles represent the average prediction error for the given control action, while the colored areas depict the error distribution. For small actuator signals, the performances of DIC and LS are comparable. For large phase WFs, instead, the proposed method outperforms the state-of-the-art, both in terms of residual error and error deviation. . . . . 37
- 4-5 Recursive Least Squares performances when persistent excitation is not provided. On the left-hand side, the residual wavefront prediction root-mean-square (rms) error is plotted over time. The x-axis represents time and shows the control iterations. On the y-axis, instead, the rms error is reported, in the wavelength unit. If the forgetting factor  $\gamma$  is set to be strictly lower than 1, the model eventually becomes unreliable and the reconstruction accuracy drops. This is not the case for the proposed DIC method or if Recursive Least Squares (RLS) is set to have  $\gamma = 1$ . On the right-hand side, the maximum covariance matrix element is plotted. The x-axis represents time and shows the prediction steps. On the y-axis, instead, the maximum entry of the covariance matrix  $P$  is shown, in the logarithmic scale. If the forgetting factor  $\gamma$  is set to be strictly lower than 1, the uncertainty increases and the model eventually becomes unreliable. . . . . 40
- 4-6 Reconstruction efficiency of the proposed method, plotted for different L2-regularizing coefficient values. The trained identity model  $i(\theta_c, \theta_q) = q(\theta_q, c(\theta_c))$  is inspected. Each one of the first 40 Zernike modes is fed to the system with unitary amplitude. The output of the identity model, which represents the theoretical system reconstruction of the requested wavefront in terms of Zernike magnitudes, is reported on the y-axis. Ideally, only the desired mode should be reproduced, with unitary amplitude. The histograms on the right depict the diagonal elements of the grids on the left-hand side. They represent the reconstruction efficiency regarding the desired mode only, that is, not considering the *cross-talking* between modes. The larger the regularizing gain, the more the reconstruction quality is affected, especially for high-order modes. . . . . 42
- 4-7 Risk of saturating actuators, as a function of the applied L2 regularizing gain. On the y-axis, it is plotted the maximum actuator input value required to reconstruct the first 40 Zernike modes with unitary amplitudes. On the x-axis are reported the different penalty values  $\alpha$ . This is equivalent to the maximum value of the control matrix  $S$ , analyzed in 2-11. Increasing the regularizing term reduces the magnitude of the control actions and thus the likelihood of actuator saturation. . . . . 43
- 4-8 Maximum amplitude of the first 10 Zernike modes that are reconstructed with a wavefront rms error lower than  $0.3 \lambda$ . The first 3 modes (piston, tip, and tilt) are excluded since the used DM is not able to generate them. The proposed method outperforms the linear state-of-the-art on all the modes. Indeed, DIC allows reconstructing larger mode amplitudes maintaining a low WF error. . . . . 44
- 4-9 Closed-Loop reconstruction of astigmatism mode using Direct Inverse Control. . . . . 46
- 4-10 Closed-Loop reconstruction error of the proposed method, averaged over 50 WF corrections. The reference WFs are randomly sampled from the first 10 Zernike modes, and it is made sure that no actuator saturation occurs. Consistently with the results presented in the Open-Loop section, the proposed method shows lower wavefront rms error during the first control iterations. Nevertheless, the performances of DIC and the benchmarks are comparable once convergence is reached, after 3 – 4 control iterations. . . . . 46

- 4-11 Closed-loop reconstruction error of the first 6 Zernike modes, excluding piston, tip and tilt. For each Zernike, the largest mode amplitude which does not trigger saturation is taken as WF reference. Each WF reconstruction is run 5 times, and the results are then averaged. The proposed method shows lower residual error during the first control iterations for all the analyzed modes. Once convergence is reached, however, the difference between DIC and the linear benchmarks is negligible. 47
- 4-12 Residual wf correction rms error at different temperature levels. The generated WF is Zernike mode 5 with amplitude  $3.14 \lambda$ , and each trajectory is the average of 5 experiment runs. In between episodes, random wavefronts are corrected in order to allow RLS and DIC methods to maintain the models updated. These 2 algorithms indeed present only minor performance degradation regardless of the external condition. On the other hand, the static linear model identified through LS presents slower convergence time and steady-state offset at  $35^{\circ}C$  and  $40^{\circ}C$ . . 48



---

## List of Tables

2-1	Specifications of the used DM from OKO Technologies. . . . .	14
2-2	Specifications of the used Shack-Hartmann (SH) sensor. . . . .	14
2-3	Schematic description of ADAM Algorithm. . . . .	23
3-1	Schematic summary of the Online DIC algorithm for AO. . . . .	31
4-1	Summary of the investigated test cases. LS and RLS respectively stand for Least Squares and Recursive Least Squares. . . . .	35
4-2	Error signals and their sources. $a_r$ , $a$ , and $\bar{a}$ are respectively the reference, measured and predicted WFs. . . . .	35
4-3	WF root-mean-square prediction error $e$ obtained by the proposed method and benchmarks at different temperatures. For each test condition, the prediction of 100 random small-phase wavefronts is considered. . . . .	38
4-4	root-mean-square performance error $e$ obtained by the proposed method and benchmark on 50 wavefronts, randomly generated sampling the first 10 Zernike modes. . . . .	45
4-5	Summary of the Closed-Loop extension of the Online DIC Algorithm. . . . .	49



---

# Acknowledgements

Firstly, I would like to thank my supervisors Carlas Smith, Dimitri Kromm, and Serafim Korovin for their assistance and feedback during the thesis project. I am also grateful to Oleg Soloviev, who has provided me with valuable insights and support at various points during my thesis.

I would like to extend my heartfelt thanks to my family and friends back in Italy for their constant support. They provided me with the necessary attitude and balance throughout this arduous journey. Finally, I reserve a special mention for my beloved partner, Kika, whose unwavering encouragement has been the cornerstone of my success in this project. Her belief in me and her constant motivation has been instrumental in helping me navigate through the most challenging phases of this thesis.

Delft, University of Technology  
May 11, 2023

Alberto Dall'Ora



“There is neither happiness nor misery in the world; there is only the comparison of one state with another, nothing more. He who has felt the deepest grief is best able to experience supreme happiness. We must have felt what it is to die, Morrel, that we may appreciate the enjoyments of life.”

— *Alexandre Dumas*



---

# Chapter 1

---

## Introduction

Optical microscopy has played a crucial role in biological research for hundreds of years now. The possibility of gathering detailed information regarding the structure of specimens in a non-invasive way has tremendously helped scientists in numerous discoveries and research advances [5]. Ideally, provided that high-quality components are employed and a precise calibration process is done, the resolution of conventional microscopes is only constrained by the diffraction limit. However, the quality of the observed image is in practice always impacted by the heterogeneous structure of the analyzed specimen. Indeed, regardless of the quality of the setup, wavefront aberrations are introduced as light passes through sections with different refractive indices [6]. This phenomenon strongly limits the maximum depth at which useful images can be gathered and therefore represents a major hindrance to research advances in microscopy.

Adaptive Optics (AO) was firstly applied to microscopy in the early 2000s to specifically address this issue. Its principle consists of reflecting the incoming light beam using an active dynamical element, like a deformable mirror (DM), in order to compensate for wavefront aberrations. During the last two decades, plenty of research microscopy advances have been made possible by this technique, as thoroughly reported in [7]. As of now, a multitude variety of AO techniques exist, depending on the type of microscope and other use-case-specific requirements. Indeed, currently the most challenging and interesting concern in the AO world is not to prove its potential and applicability, but rather to make it more accessible worldwide [8]. During the first years of 2000s, membrane DMs started being manufactured using micro-lithography technologies, the same ones used to produce electronic chips [9]. This was a major breakthrough, which significantly reduced the cost of AO instruments and thus expanded the market of DMs. Nevertheless, membrane deformable mirrors do present some performance limitations, mostly related to the nonlinearity and temperature-dependency of the mirror's response.

The system's nonlinearities arise when large surface deformations are required, that is when the system operates far from its linearization point. For this reason, the analysis of deep biological tissues, which are characterized by significant wavefront aberrations, is a rather

challenging task. Intuitively, the less accurate the identified system model is, the more control iterations are required to provide a sharp image of the specimen. In other words, wavefront correction becomes slower for large amplitude wavefront (WF)s. On the other hand, however, reducing the exposure time of the specimens and being able to operate at a high speed are extremely desirable features in microscopy, especially when analyzing *in-vivo* samples.

Furthermore, the behavior of membrane DMs is significantly affected by environmental conditions, like temperature and humidity. If these parameters change with respect to the calibration values, the model becomes inaccurate and the control less stable. In order to avoid frequent recalibrations, the system model needs to be continuously updated online, which poses interesting identification challenges.

Enhancing the range of phases that AO can successfully correct, without affecting the speed of WF correction, is currently one of the most active research fields in AO [6]. Finding an effective solution is rather challenging since robustness and reactivity are usually on opposite sides of the control spectrum. This master thesis project aims at improving the homogeneity of correction accuracy while preserving the model's reliability. It does so by investigating the following research question.

**Research question** How to mitigate the impact of DM nonlinearities and temperature changes on correction accuracy, without reducing control speed?

## 1-1 Outline

The thesis report is organized as follows. Chapter 2 introduces Adaptive Optics and all the necessary theoretical concepts needed to comprehend the researched problem. Chapter 3, instead, describes the proposed approach. Finally, Chapter 4 reports and analyzes the obtained results, while Chapter 5 concludes the thesis.

# Theoretical background

This chapter introduces the theory and concepts which the later proposed method is based on. In this regard, Adaptive Optics (AO) theory is firstly explained in Section 2-1. Then, in Section 2-2, the most common AO methodologies are discussed, especially regarding system identification and control. Finally, the core principles of Artificial Neural Networks (ANN) are introduced in Section 2-3, together with the latest applications to AO systems.

## 2-1 Adaptive Optics

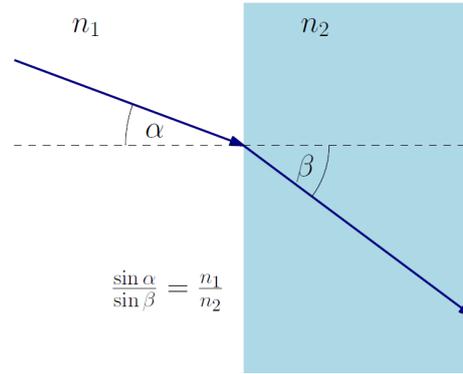
Optical microscopes have been at the core of biomedical research for hundreds of years. Theoretically, the resolution of conventional microscopes is solely defined by the diffraction limit of light. However, in practice, performances are also affected by the optical properties of the specimens, which often present varying refractive indices that introduce aberrations. In the last decade, adaptive optical components like deformable mirrors have been increasingly implemented in microscopy in order to overcome such limitations and improve the quality of microscope images. In this section, AO theory is introduced, by discussing important concepts and methodologies.

### 2-1-1 Image formation

A central concept of all the disciplines which involve optics is the propagation of light. This phenomenon can be described through different models, each of them having its own advantages and disadvantages.

In this regard, geometrical optics is a powerful and straightforward model, which makes use of only a few assumptions. The first simplification considers light as propagating in straight lines with constant speed  $v$ , when in homogeneous mediums. The ratio between the velocity  $v$  and the light speed in vacuum  $c$  is called *refractive index* and it is defined as

$$n = \frac{c}{v}.$$



**Figure 2-1:** Representation of the Law of Refraction [1]

The second assumption is that light rays change direction when transitioning through mediums with different refractive indices. This phenomenon is formalized in the *law of refraction*:

$$\beta = \arcsin \left( \sin(\alpha) \cdot \frac{n_1}{n_2} \right), \quad (2-1)$$

where  $\alpha$  and  $\beta$  are the angles of propagation in a region having refractive indices  $n_1$  and  $n_2$  respectively. In inhomogeneous mediums, the refractive index is dependent on the position  $\mathbf{r}$ ,  $n = n(\mathbf{r})$ . These approximations, together with the idea of representing propagating light as rays, significantly simplify the description of optical imaging. In this regard, some of the most recurring elements in geometrical optics are here briefly described.

- *Point source* - A radiant source of light, which emits light uniformly in all directions and has negligible size.
- *Object* - A collection of independent point sources.
- *Optical Path Length (OPL)* - The integral of the refractive index along the light propagation path  $P$

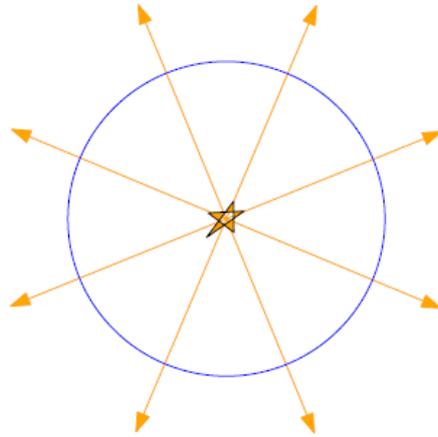
$$\text{OPL} = \int_P n(r) dr.$$

Intuitively, it also represents the travel time of light along the path  $P$ .

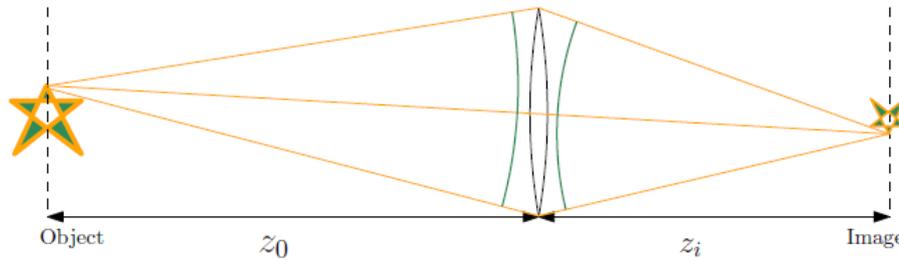
Consistently with the concepts just defined, in geometrical optics a *wavefront* is defined as a surface of a constant OPL from the source. Alternatively, it can also be described as a surface perpendicular to all the light rays emitted by a point source. A visual example of this concept is given in Figure 2-2.

In order to be able to visualize a source of light, all the rays emitted by the point source need to be recollected to a single point, which is called *image* of the source. What lenses and other similar imaging systems do is indeed to exploit the *law of refraction* to make incoming light rays converge to a single point. In Figure 2-3 a descriptive diagram of this process is depicted.

In order to correctly place scientific cameras and obtain high-quality images of the observed



**Figure 2-2:** Spherical wavefront generated by a point source in a homogeneous medium [1]



**Figure 2-3:** Example of object imaging through a thin lens [1]

object, it is of crucial importance to correctly measure the distance  $z_1$  between the lens and the image plane. Considering a perfectly homogeneous medium, the distance  $z_1$  can be theoretically computed using the *thin lens* equation

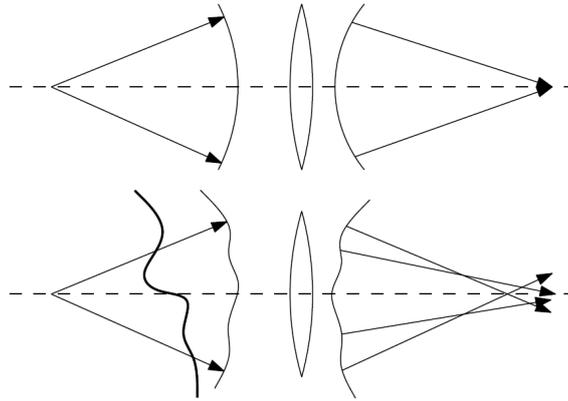
$$\frac{1}{z_0} + \frac{1}{z_1} = \frac{1}{f}, \quad (2-2)$$

where  $f$  is the focal length of the lens.

### 2-1-2 Wavefront aberration

The *thin lens* equation introduced in the previous section truthfully describes the mapping of an object to its image only in the ideal case of perfect alignment and homogeneous medium. In practice, these assumptions are rarely satisfied, and therefore rays do not perfectly converge to a single point on the image plane. As a consequence, the resolution of the captured image, that is, the minimum distance between resolvable points, is reduced.

Static wavefront (WF) aberrations are inherent in optical systems and are usually due to misalignments in the system or inaccuracies in the manufacturing of optical components. The most effective way to compensate for these defects is to adopt more precise instruments



**Figure 2-4:** Example of an unaberrated and aberrated wavefronts when passing through a lens [1]

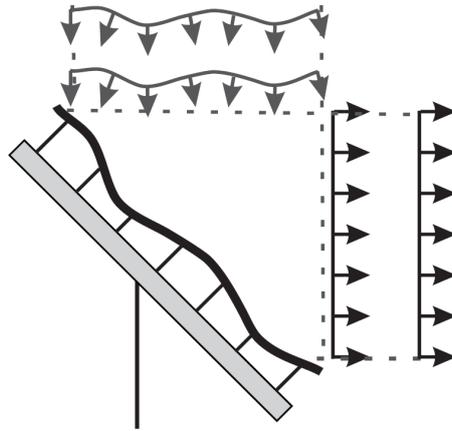
and enhance the complexity of the setup. Nevertheless, even ideal systems can suffer from dynamic wavefront aberrations. Indeed, these defects are due to exogenous sources and cannot be corrected *a priori*.

For example, in microscopy, light usually passes through complex biological specimens having heterogeneous optical properties. The different refractive indices cause the rays to deviate and introduce wavefront aberrations, which are strictly related to the analyzed specimen. Furthermore, when analyzing 3-dimensional *in-vivo* samples, aberrations can also change over time and space following random and unpredictable patterns, which makes these sorts of defects rather challenging to be corrected.

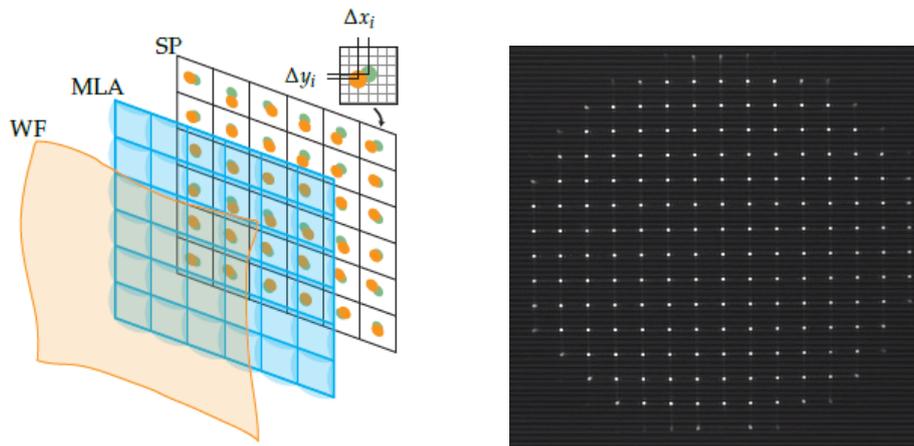
In this regard, Adaptive Optics is a powerful technique able to compensate for dynamic wavefront aberrations. The idea at the base of such a method is to exploit the *phase conjugation* principle [10], which states that if an aberrated wavefront is reflected from a mirror having the same shape and half the amplitude, the reflected beam will have a flat wavefront. Therefore, if the WF is precisely measured and the mirror is promptly deformed, the imaging system will be able to collect a sharp image of the object. However, these two assumptions are not trivially satisfied in practice, and therefore, specific discussions on these matters are reported in Section 2-1-3 and 2-1-4 respectively.

### 2-1-3 Wavefront detection

Scientific cameras, including wavefront sensors, are not able to measure WFs directly, but instead, they can only quantify the brightness intensity of the light they are exposed to. Nevertheless, pupil-plane sensors like the Shack-Hartmann (SH) sensor, are able to retrieve an accurate approximation of the incoming WF only using a single measurement [11]. The SH sensor presents an array of  $M$  lenslets, which divides the pupil plane into  $M$  smaller spatial areas. The incoming light beam traveling through the sub-apertures of the sensor gets projected to  $M$  individual focal spots. Ideally, if the wavefront of the light rays is flat, the spots will be located exactly at the center of the sub-apertures. On the other hand, when a distorted wavefront goes through the lenslets array, the centroids will be shifted with respect to the reference positions. The major simplification that the SH-based WF reconstruction



**Figure 2-5:** Application of the phase conjugation principle through a deformable mirror [1]



**Figure 2-6:** On the left, a schematic of a Shack-Hartmann sensor [1]. WF is the incoming wavefront, MLA is the microlens array, and SP is the projected spot pattern. On the right, is an image of the centroids pattern from the SHWFS used for this project.

applies, is to approximate the wavefront to be locally tilted in each one of the sub-apertures. This allows to reconstruct the overall WF through linear regression and thus to significantly reduce complexity. However, this advantage comes at the cost of losing high-frequency wavefront information, especially when the number of sub-apertures  $M$  is limited [12].

Reconstructing a wavefront through a SH sensor generally consists of three steps: locating the centroids, computing the slopes with respect to the reference pattern, and finally using a zonal or modal approach to obtain the incoming wavefront  $\phi$ .

A possible way to locate the centroids  $(x_{c,m}, y_{c,m})$  is following a center of mass principle based on the measured pixel intensities  $I_{i,j}$ :

$$x_{c,m} = \frac{\sum_{i,j} x_{i,j} I_{i,j}}{\sum_{i,j} I_{i,j}} \quad \text{and} \quad y_{c,k} = \frac{\sum_{i,j} y_{i,j} I_{i,j}}{\sum_{i,j} I_{i,j}}, \quad (2-3)$$

with  $(i, j) \in B_m$ , where  $m$  is the lenslet number and  $B_m$  contains all the pixels belonging to the  $m$ -th lenslet.

Secondly, comparing the measured centroids' locations  $(x_{c,m}, y_{c,m})$  with the reference spots  $(x_{r,m}, y_{r,m})$ , the wavefront slopes are computed :

$$\begin{pmatrix} \langle \partial\phi/\partial x \rangle \\ \langle \partial\phi/\partial y \rangle \end{pmatrix}_m = \begin{pmatrix} s_x \\ s_y \end{pmatrix}_m \approx \frac{1}{z} \begin{pmatrix} x_c - x_r \\ y_c - y_r \end{pmatrix}_m + \begin{pmatrix} \eta_x \\ \eta_y \end{pmatrix}_m, \quad (2-4)$$

where  $z$  is the propagation distance between the aperture detection planes. Measurement noise and approximation errors are both included in the term  $\eta$ .

Once the slopes  $\mathbf{s}$  have been computed, the wavefront can finally be reconstructed. Here, a modal approach is described, which consists of approximating the measured WF as a weighted sum of a set of basis functions, for instance, Zernikes polynomials [13]:

$$\mathbf{s} = \mathbf{Z}\mathbf{a}$$

The vector  $\mathbf{a}$  represents the weighting coefficients, while the matrix  $\mathbf{Z}$  contains the slopes of the modal functions:

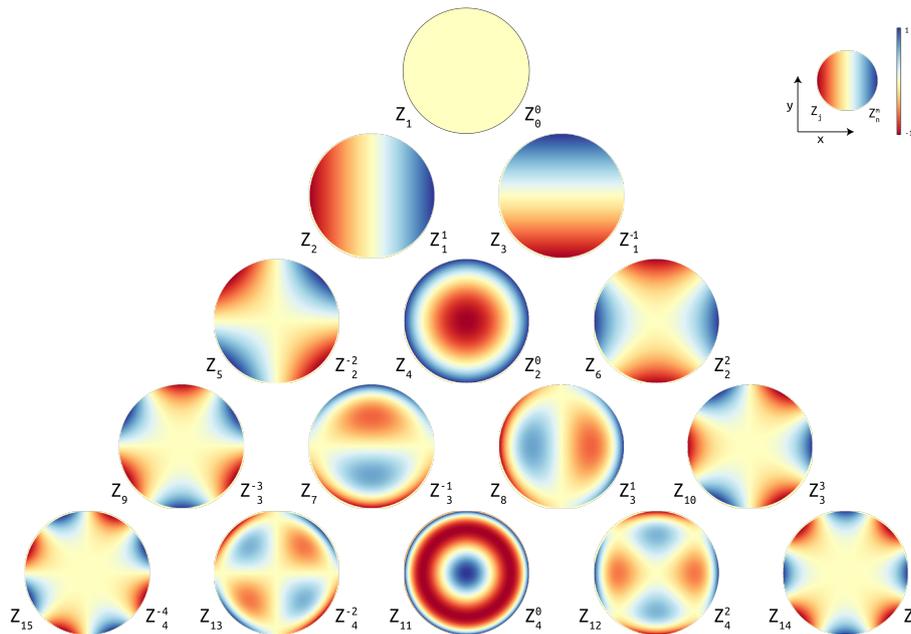
$$\mathbf{s} = \begin{bmatrix} s_1 \\ s_2 \\ s_3 \\ \vdots \\ \vdots \\ \vdots \\ s_M \end{bmatrix} \quad \mathbf{Z} = \begin{bmatrix} \frac{\partial \bar{Z}_1(x_1, y_1)}{\partial x} & \frac{\partial \bar{Z}_2(x_1, y_1)}{\partial x} & \dots & \frac{\partial \bar{Z}_D(x_1, y_1)}{\partial x} \\ \frac{\partial \bar{Z}_1(x_2, y_2)}{\partial x} & \frac{\partial \bar{Z}_2(x_2, y_2)}{\partial x} & \dots & \frac{\partial \bar{Z}_D(x_2, y_2)}{\partial x} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial \bar{Z}_1(x_M, y_M)}{\partial x} & \frac{\partial \bar{Z}_2(x_M, y_M)}{\partial x} & \dots & \frac{\partial \bar{Z}_D(x_M, y_M)}{\partial x} \\ \frac{\partial \bar{Z}_1(x_1, y_1)}{\partial y} & \frac{\partial \bar{Z}_2(x_1, y_1)}{\partial y} & \dots & \frac{\partial \bar{Z}_D(x_1, y_1)}{\partial y} \\ \frac{\partial \bar{Z}_1(x_2, y_2)}{\partial y} & \frac{\partial \bar{Z}_2(x_2, y_2)}{\partial y} & \dots & \frac{\partial \bar{Z}_D(x_2, y_2)}{\partial y} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial \bar{Z}_1(x_M, y_M)}{\partial y} & \frac{\partial \bar{Z}_2(x_M, y_M)}{\partial y} & \dots & \frac{\partial \bar{Z}_D(x_M, y_M)}{\partial y} \end{bmatrix} \quad \mathbf{a} = \begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ \vdots \\ \vdots \\ \vdots \\ a_D \end{bmatrix}.$$

Here,  $D$  is the number of modes that are used to describe the wavefront, and  $M$  is the number of centroids in the SH pattern.

Zernike polynomials are frequently adopted in optics due to various reasons. Firstly, they resemble common optical aberrations like coma, astigmatism, and defocus. Secondly, they can be written through simple closed-form expressions in Cartesian coordinates. Last, they are orthogonal on a unit disk, meaning that when they are applied to a circular aperture, no information redundancy is present between modes. A visual representation of the first 21 modes is given in Figure 2-7.

#### 2-1-4 Wavefront correction

As it was first introduced in Section 2-1-2, the principle at the base of Adaptive Optics consists of correcting wavefront aberrations reflecting light with a surface having the same shape and twice less amplitude. Once the incoming wavefront is detected and reconstructed, an adaptive instrument needs to be promptly deformed. In most corrective imaging systems, the adaptive element is a deformable mirror (DM), which generally speaking consists of a thin deformable layer usually made of glass or plastic, coated with reflective material. The surface



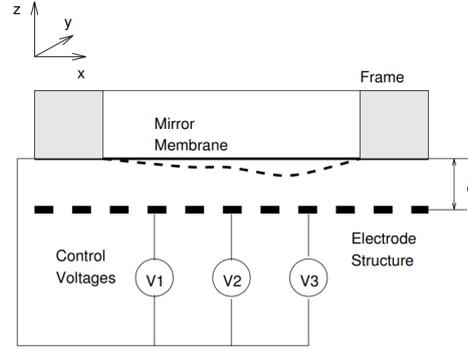
**Figure 2-7:** Graphic representation of the first 15 Zernike modes. This basis is commonly used in optics due to various interesting properties. Zernike polynomials are orthogonal on the unit disk, have simple formulations in the cartesian domain, and resemble typical optical deformations like coma and astigmatism.

of such a mirror is changed by triggering  $C$  actuators, placed beneath the reflective layer. The performances of DMs are thus characterized not only by the number of actuators, but also by the maximal displacement of the surface, called *stroke*, and by the inter-actuators' distance, which influences the precision of the correction.

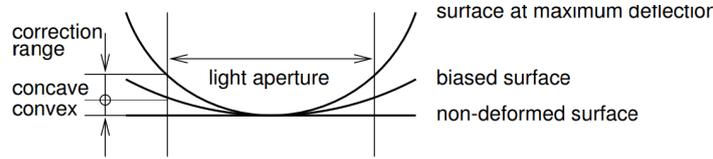
Throughout the years, different designs and manufacturing technologies have been developed, which have given birth to various DM types. A thorough review on the topic can be found at [14]. In this following section, micromachined membrane deformable mirrors (MMDMs) are discussed, consistently with the equipment at our disposal in the Delft Center for Systems and Control (DCSC) Adaptive Optics lab.

### Micromachined Membrane Deformable Mirrors

The first example of a DM consisting of a thin plastic membrane, deformed by an array of electrostatic or magnetic actuators was introduced in 1977 by Grosso and Yellin [15] (Figure 2-8). Twenty years later, membrane DMs started being produced using micromachine technologies, in a way similar to electronic chips [16]. This was a major breakthrough that significantly reduced the production cost of DMs and consequently made AO solutions more appealing. MMDMs present a reflective membrane fixed at the edges that can only be attracted to the electrode structure, thus assuming only concave shapes. Nevertheless, if a slightly concave deformation is taken as a reference, operations are allowed in both directions. An explicative diagram on this matter can be found in Figure 2-9. Electrostatic electrodes can easily be placed in a high-density structure, however, the deformation they are able to trigger is



**Figure 2-8:** Schematic of a Membrane DM. The shape of the reflective layer is changed through electrostatic force, applied by electrodes beneath the surface. [2]



**Figure 2-9:** Schematic of biased operation of a Membrane DM. Bi-directional operations are enabled by a slightly concave reference shape. [2]

directly proportional to their dimension. This poses a design trade-off between correction accuracy and deformation stroke, which usually leads to membrane DMs able to reproduce large deformations for low-order WF modes, and low deformations for high-order modes. Finally, the relatively low production cost, together with good accuracy and lack of hysteresis, makes this kind of deformable mirror one of the most commonly used in Adaptive Optics systems [17].

## 2-2 AO identification and control

In order to correctly compensate for wavefront aberrations, an accurate model of the AO system needs to be identified. In this work, Zernike modes are used, due to their simplicity, robustness, and similarity with real optical aberrations. The inputs of the model are then defined as the electrical signals of the actuators, while the outputs are set as the Zernike coefficients describing the measured wavefront.

Most deformable mirrors present a linear or linearizable relationship between the voltages applied to the electrodes and the deformation of the reflective surface. Furthermore, AO components usually do not present dynamics, and a change in actuators' commands has an immediate consequence on the shape of the mirror. For these reasons, AO systems are usually modeled as an influence matrix  $Q$ :

$$\mathbf{a}_{k+1} = Q \cdot \mathbf{u}_k + \eta_{k+1} \quad (2-5)$$

where  $\mathbf{a} \in \mathbb{R}^D$ ,  $\mathbf{a} = [a_1, a_2, \dots, a_D]^T$  are the WF slopes,  $Q \in \mathbb{R}^{D \times C}$  is the AO model,  $\mathbf{u} \in \mathbb{R}^C$ ,  $\mathbf{u} = [u_1, u_2, \dots, u_C]^T$  are the actuators' commands, and finally  $\eta \in \mathbb{R}^D$ ,  $\eta =$

$[\eta_1, \eta_2, \dots, \eta_D]$  represents measurement noise. Intuitively, the elements of the matrix  $Q$  describe the influence that each actuator has on every WF mode.

The most straightforward method to retrieve such a model is based on the superposition principle. It consists of poking each actuator individually while keeping the rest of the mirror in a flat position. The influence of the  $i$ -th actuator on the mirror deformation, measured as a vector of  $M$  slopes, forms the  $i$ -th column of the influence matrix  $Q$ . Despite the simplicity and data efficiency of such an identification technique, data-driven approaches are usually preferred, since they better compensate for modeling errors that arise due to nonlinearities and coupling effects between actuators [18].

### 2-2-1 Least Squares method

Assuming the measurement noise  $\eta$  is white and Gaussian, the optimal linear estimator of the matrix  $Q$ , based on  $N$  measurement samples, is computed through the *Least Squares* approach:

$$Q = \arg \min_Q \sum_{k=1}^N (\mathbf{a}_k - Q\mathbf{u}_k)^2 \quad (2-6)$$

$$= \arg \min_Q \|\mathbf{A}_N - Q\mathbf{U}_N\|^2. \quad (2-7)$$

with  $\mathbf{A}_N \in \mathbb{R}^{D \times N}$ ,  $\mathbf{A}_N = [\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_N]$ ,  $\mathbf{U}_N \in \mathbb{R}^{C \times N}$ ,  $\mathbf{U}_N = [\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_N]$  being  $N$  random input-output samples. Imposing the derivative of the loss function to be equal to 0, a closed-form solution can be retrieved:

$$Q = \mathbf{A}_N \mathbf{U}_N^T (\mathbf{U}_N \mathbf{U}_N^T)^{-1}. \quad (2-8)$$

In order to make the matrix inversion process well-posed, the dimension of the dataset  $N$  needs to be larger than the number of actuators  $C$ .

Once the model of the system has been identified, the influence matrix  $Q$  is inverted in order to obtain a control matrix  $T$ , which maps desired wavefronts to the corresponding actuators' commands:

$$\mathbf{u}_k = S \cdot \mathbf{a}_r \quad (2-9)$$

Since the matrix  $Q$  is usually not square, the Moore-Penrose pseudo-inverse is usually computed, using Singular-Value-Decomposition (SVD):

$$Q = U \Sigma V^T \quad (2-10)$$

$$S = V \Sigma^{-1} U^T \quad (2-11)$$

In this case, SVD breaks down the influence matrix  $Q$ , such that the mapping between actuators commands and WF modes is done by passing through a set of orthogonal modes, also called *singular* or *eigen modes*. These functions describe the patterns the mirror follows when changing shape, and thus truthfully represent the reconstruction abilities of the system. Understanding the connection between Zernikes and such modes is then crucial for WF correction tasks.

In particular, the rows of  $V^T$  matrix maps the influence of each actuator commands to the eigen modes, while the columns of  $U$  represent the modes in sensor space. The singular values stored on the diagonal of  $\Lambda$  represent the ability of the mirror to reproduce the corresponding mode, and therefore the importance of that mode for WF reconstruction. Inverting the order of the matrices as done in Eq. 2-11 allows to achieve the opposite task, that is, to compute the linear combination of eigen modes, and thus of actuators commands, that is able to create a desired wavefront shape.

Pseudo-inversion is a powerful tool, which however requires some extra attention when applied to real systems [19]. Indeed, the signal-to-noise ratio (SNR) of high-order parameters of  $Q$  is in practice often low, and as a consequence, a full inversion of the matrix can lead to extremely large entries in the control model  $S$ . In other words, the wavefronts that the DM finds harder to reproduce, require larger control actions. Since this might cause actuators to saturate, a regularization process is usually implemented. Singular-Value-Decomposition provides a straightforward way to maintain the inversion process well-posed, which consists of setting some of the lowest singular modes of  $\Lambda$  to zero. Naturally, this comes at the expense of a reduced reconstruction ability of the mirror. Indeed, discarding some of the singular modes will force the mirror to not use its full potential, and thus some high-order wavefronts will not be controllable anymore. Nevertheless, this is usually a worthy trade-off, since most common optical aberrations can be accurately approximated through low-order modes.

In this regard, a clear metric of the reconstruction abilities of the mirror is given by the *similarity matrix*  $T = S \cdot Q$ , for which it holds:

$$\mathbf{a} = T \cdot \mathbf{a}_r \quad (2-12)$$

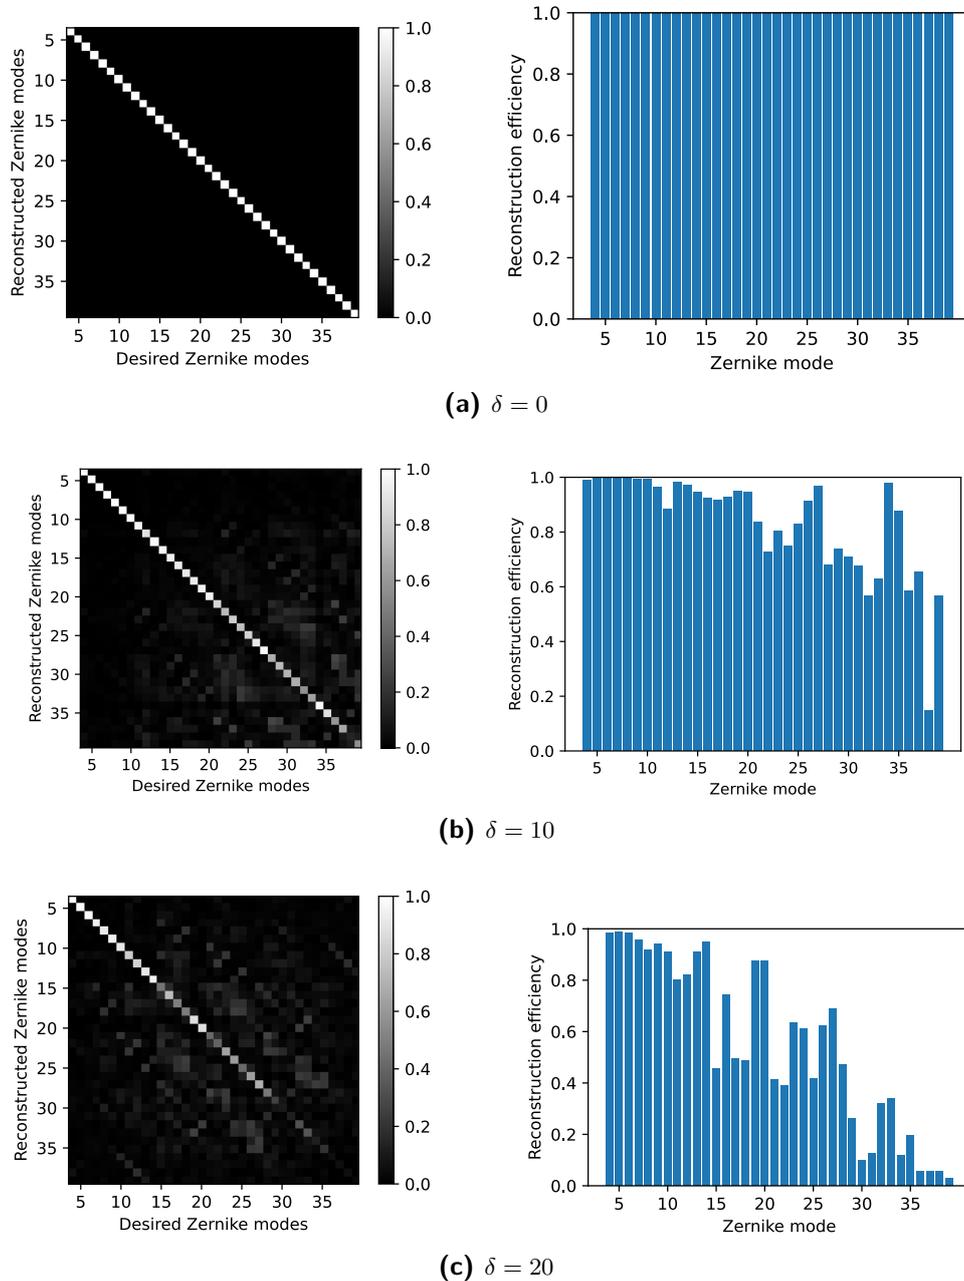
where  $\mathbf{a}_r$  are the desired modal coefficients and  $\mathbf{a}$  are the ideally obtained ones, according to the influence and control matrices [20].

In this way, not only the reconstruction efficiency of the reference mode can be evaluated, but also the impact of undesired ones (*cross-talking*). Ideally,  $T$  should indeed be a diagonal matrix, however, SVD truncation reduces the mirror's ability to reconstruct higher-order modes and thus it inevitably impairs performances. On the other hand, letting the control matrix  $S$  ignore the modes that require larger actuators signals, minimizes the risk of saturation and thus of nonlinear behaviours.

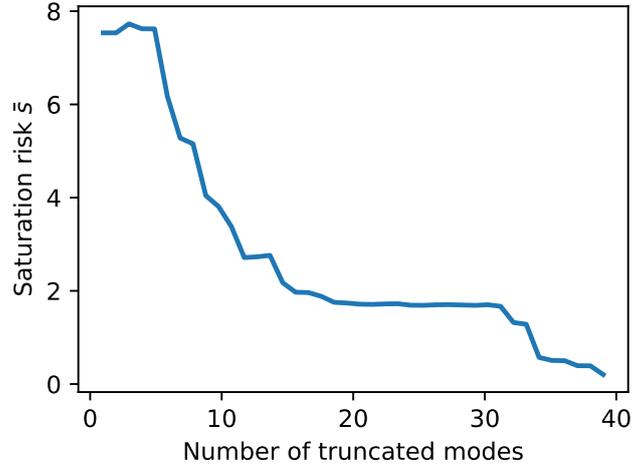
Although the onset of saturation depends on the amplitude of the desired wavefront  $\mathbf{a}_r$ , an accurate metric of the likelihood of such behaviour is represented by the largest value of the control matrix  $S$  [20]. In this regard,

$$\bar{s} = \max(S)$$

is also analyzed and plotted in Fig. 2-11 to give a more complete picture of the effects of the truncation parameter  $\delta$ .



**Figure 2-10:** Representation of the similarity matrix  $T$  and reconstruction efficiency  $diag(T)$ , for different number of truncated modes  $\delta$ . The gray scale represents the absolute value of the matrix, element-wise. All Zernike modes between 4 (top-left) and 39 (bottom-right) are represented. The histograms on the right depict the elements on the diagonal of the similarity matrix and range from Zernike mode 4 to 39. Truncating the lowest singular values decreases the reconstruction efficiency of highest-order modes, and enhances the cross-talking with other, off-diagonal, elements.



**Figure 2-11:** Maximum value of the control matrix  $S$  for different number of truncated singular modes  $\delta$ .

### 2-2-2 Challenges and limitations

In this section, the main hindrances of AO identification and control are discussed. In this regard, analyses and considerations are supported by experimental results obtained through an AO setup present in the AOlab of the TU Delft DCSC.

#### Description of the AO setup

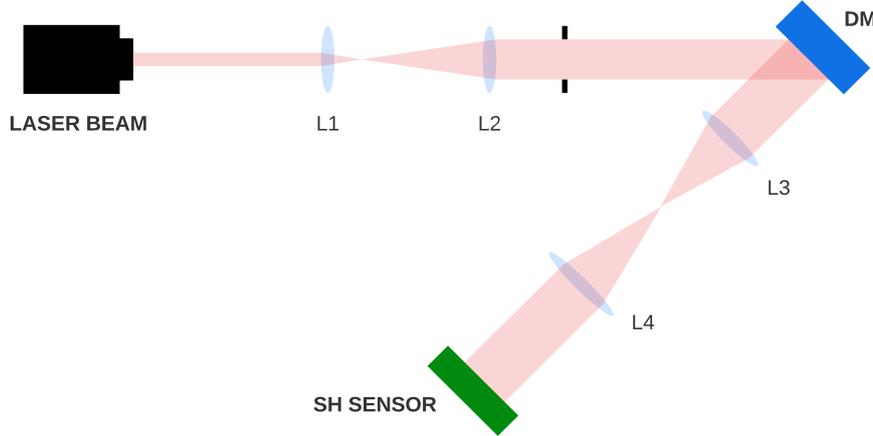
As can be seen in Figure 2-12, the experimental AO setup is essentially composed of a deformable mirror and a SH sensor. Collimated light is emitted by a laser as a red beam having a wavelength equal to  $640 \text{ nm}$ . The beam is then magnified through a couple of positive bi-convex lenses, and filtered by a  $20 \text{ mm}$  sub-aperture. A 96-channel membrane DM produced by OKO Technologies is used, whose characteristics are reported in Table 2-1 [21]. After being reflected and diminished, the light beam enters the SH sensor. The specifications of this last component are provided in Table 2-2.

Number of actuators	Aperture	Maximum surface deflection
96	25.4 mm	$19 \mu\text{m}$

**Table 2-1:** Specifications of the used DM from OKO Technologies.

Aperture	AR coating range (ARC)	Pitch	Focal length
10 mm	400 nm - 900 nm	$300 \mu\text{m}$	14.2 mm

**Table 2-2:** Specifications of the used SH sensor.



**Figure 2-12:** Experimental AO setup used for this thesis project. AO Lab, DCSC, TU Delft. The deformable mirror is an OKO 96-channel MMDM with 25.4 mm aperture. The SHWFS consists of a Thorlabs  $1024 \times 1024$  CMOS Camera and a  $10 \text{ mm} \times 10 \text{ mm}$  lens array. The laser beam has 640 nm wavelength, a diameter of 2.9 mm, and 1.2 mW power.

## Nonlinearity

As introduced in the previous section, linear models are often preferred over nonlinear approaches for the modeling and control of DMs. However, the relationship between actuators' commands and mirror deformation is often quadratic in membrane DMs [2], and, as a consequence, some linearization process needs to be implemented. Two different approaches are often followed: the first simply consists of choosing a linearization point and assuming the mirror performs linearly around it:

$$\hat{\mathbf{a}} = \mathbf{Q} \cdot \hat{\mathbf{u}}$$

$$\hat{\mathbf{u}} = \mathbf{u} - \mathbf{u}_R, \quad \hat{\mathbf{a}} = \mathbf{a} - \mathbf{a}_R$$

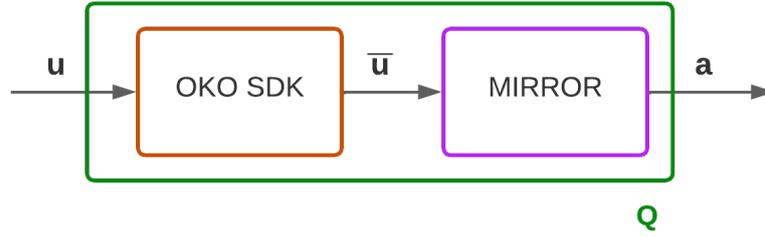
where  $\mathbf{Q}$  is the influence matrix of the mirror,  $\mathbf{a}$  and  $\mathbf{a}_R$  are respectively the measured and reference wavefronts, and finally  $\mathbf{u}$  and  $\mathbf{u}_R$  are the actuators action sets that generate  $\mathbf{a}$  and  $\mathbf{a}_R$  respectively.

Despite its clear simplicity, this method is only suitable for the correction of small wavefront aberrations. Indeed, the approximation error becomes more significant the more the desired wavefront is distant from the linearization point  $\mathbf{a}_R$ .

Alternatively, the model can directly be retrieved in a linear fashion, taking as inputs the squares of the applied electrodes' voltages. This is the case of our AO setup, where an OKO deformable mirror is implemented, together with its Software Development Kit (SDK). Every time an input command is sent to the mirror, the OKO SDK applies the following transformation to each element of the control array:

$$\bar{u}_i = \sqrt{\frac{u_i + 1}{2}} \cdot u_{max} \quad (2-13)$$

In this way, the control commands  $u \in [-1, 1]$  get nonlinearly mapped to  $\bar{u} \in [0, u_{max}] \text{mV}$ , which are the voltage signals effectively sent to the actuators. Due to the quadratic relationship between the latter and the mirror deformation  $\mathbf{s}$ , an overall linear model  $\mathbf{Q}$  is obtained.



**Figure 2-13:** Modelling diagram of the DM implemented in our setup. The mirror presents a quadratic relationship between the applied voltages contained in  $\bar{\mathbf{u}}$  and the wavefront  $s$ . The OKO SDK applies the square root operator, together with scaling and shifting, to create a linear relationship between the control commands  $\mathbf{u}$  passed to the mirror and the wavefront  $s$ .

Opposed to the first linearization approach, this second method theoretically allows obtaining a linear model of the system with no approximation error. However, it has been so far assumed that the mirror shows a perfectly quadratic response, which might not be the case in reality. To validate such an approach, all the mirror's actuators are poked with increasing control signals and the measured wavefronts are evaluated with respect to the flat mirror deformation:

$$\|\mathbf{a} - \mathbf{a}_0\|_2$$

where  $\mathbf{a}_0$  is the wavefront obtained for  $\bar{\mathbf{u}} = 0mV$ .

As shown in figure 2-14, the quadratic assumption proves to be satisfactory for low and mid voltages. However, consistently with the documentation provided by OKO [2], the relationship between large control actions and mirror deflection is better represented by a polynomial of higher order. As a consequence, linearized models will inevitably present inaccuracies for large deformations.

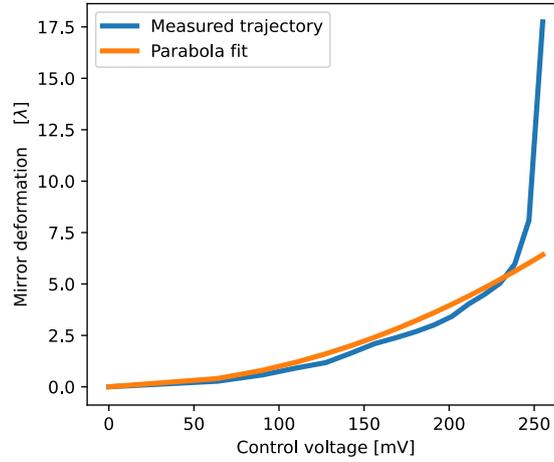
### Time-varying dynamics

The second major simplification that the traditional Least Squares method implements, is to assume that the AO system remains constant over time. In practice, this condition is hardly ever satisfied, due to mechanical drift and changes in external conditions, like humidity and temperature. This last issue is particularly interesting to analyze as it can significantly impact the response of the system in a rather short period of time. Mechanical drift, instead, has a negligible influence on AO performances if the system is periodically calibrated.

The metric used here to quantify the correction performances is the root-mean-square (rms) wavefront error, in units of the wavelength of light:

$$\sigma = \text{RMS}_w = \left( \frac{1}{A} \int_{\text{pupil}} (W(x, y) - \bar{W})^2 dx dy \right)^{1/2} \quad (2-14)$$

where  $A$  is the pupil area,  $W$  is the measured aberration in cartesian coordinates, and  $\bar{W}$  is the mean wavefront. The rms is a common indicator of WF flatness, as it indicates how much



**Figure 2-14:** Nonlinear behavior of our AO system. Control signals from  $u_{min} = 0mV$  to  $u_{max} = 255mV$  are sent to all 96 actuators. The mirror deformation  $\|\mathbf{a} - \mathbf{a}_0\|_2$  is measured and reported here on the  $y$  axis, where  $\mathbf{a}_0$  contains the Zernike coefficients describing the wavefront when  $0mV$  is applied to the actuators. The *OKO trajectory* corresponds to a quadratic curve, which is the modeling assumption made by the OKO SDK in order to make the whole system linear. The plotted result is the mean of 10 different episodes. The standard deviation is minimal and thus barely visible.

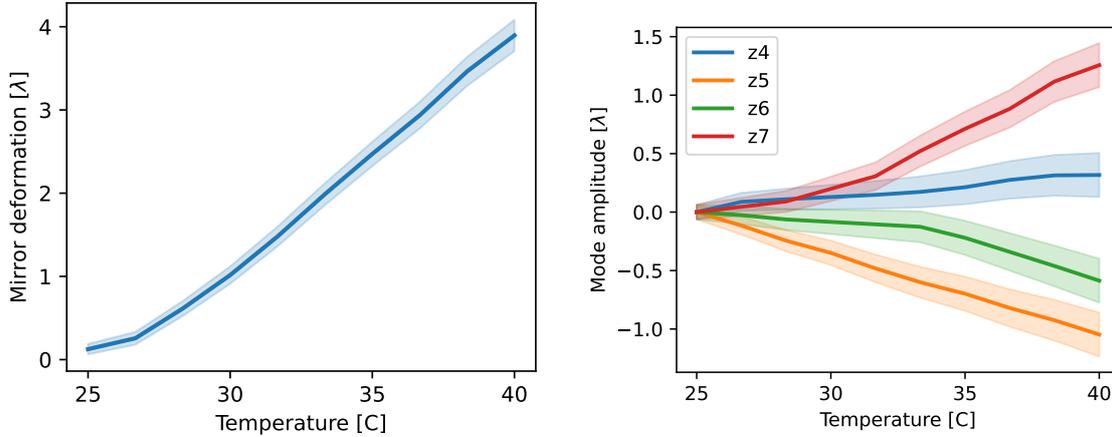
the given WF resembles its mean value. In AO it is often applied to the residual WF error and it provides information regarding the correction accuracy. Additionally, the rms can also be computed from a vector of Zernike coefficients:

$$\sigma = \left( \sum_{j=3}^N a_j^2 \right)^{1/2} \quad (2-15)$$

As can be seen in Figure 2-15, the static aberration of the AO system is significantly impacted by changes in temperature. This is mostly due to the dependency of the mirror's tension on temperature and it is thus inherent to AO systems. In Figure 2-16, the error trajectory of a closed-loop reconstruction of a random WF is plotted for different temperature values. Consistently with previous considerations regarding static aberration, the accuracy after the first control step becomes worse, the further the working point is away from calibration conditions. Furthermore, what is also interesting to notice in this plot, is that the reconstruction stability is also affected by temperature. This means that the actuators' influence on the measured mirror deformation, i.e. the model  $Q$ , is also impacted by thermal effects.

### 2-2-3 Extensions

Least Squares (LS) is a powerful and straightforward method, which is currently still largely implemented in plenty of scientific fields. As long as the assumptions it takes are not significantly unfulfilled, the modeling accuracy is usually satisfactory for most control tasks. Nevertheless, in order to overcome modeling errors and limitations, many extensions to this method have been proposed over the years.



**Figure 2-15:** Correlation between wavefront modes and temperature. The mirror is maintained in the reference position and the measured wavefront at calibration temperature ( $24^{\circ}\text{C}$ ) is taken as reference. The temperature around the setup is then increased and subsequently set back to the calibration one. Every second, a new wavefront is measured through the SH sensor while actuators are maintained in the bias position. On the left, the mirror deformation with respect to the reference wavefront is plotted. On the right, low-order Zernike modes are reported, named according to Noll's sequential indices [3].

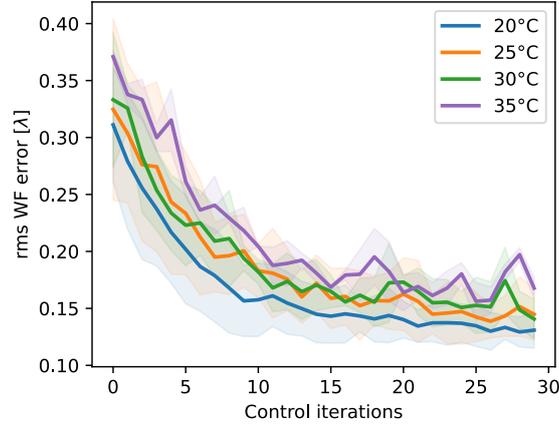
The Nonlinear Least Squares (NLS) method, for instance, adapts the traditional LS approach to the estimation of a generic modeling function  $h(\cdot)$ . A closed-form solution like the one reported in Eq. 2-8 cannot be retrieved in this case, as the loss function's derivative generally includes both independent variables and model parameters. As a consequence, the terms of the model function  $h$  are estimated iteratively. The main drawback of such an approach is that the structure of the function  $h$  has to be defined *a priori* and it then remains static, which, in a time-varying system, is not an ideal solution.

Regarding this last matter, Recursive Least Squares (RLS) is the most common and effective LS extension to identify systems that change over time [22]. Instead of collecting  $N$  samples and estimating a static matrix  $Q$  based on them, RLS updates the model with every newly gathered sample, which afterward gets discarded. Therefore, memory requirements are relaxed and the traditional LS approach is enhanced with time-tracking abilities. This method has already been implemented in AO with promising results [18][23], and in the next section, a more detailed description of it is given.

### Recursive Least Squares method

The most natural and straightforward way to implement the Least Squares method iteratively is to extend the regression dataset with every new input-output measurement, and estimate a new model  $Q_i$  at each time step. However, this causes the regression dataset dimension  $N$  to rapidly increase, inducing the system to eventually run out of space.

Recursive Least Squares adopts a more data-efficient way to keep the model  $Q$  updated. Whenever a new sample  $(\mathbf{u}_i, \mathbf{a}_i)$  is gathered, the RLS exploits it to update the entries of the matrix  $Q_i$ . In this way, past samples do not need to be stored anymore, which is particularly



**Figure 2-16:** Closed-Loop reconstruction of a random wavefront for different temperature values. A linear matrix model  $S$  is used to control the system, using a Proportional-Integral approach. Controller's gains are respectively 1 and 0.5. When the operating temperature is different from calibration one ( $23^{\circ}C$ ) the correction appears more unstable and a steady-state offset arises.

useful for high-dimensional systems like AO. Furthermore, when the system is known to be time-varying, the loss function that RLS aims to minimize is usually slightly different from the traditional LS one (Eq. 2-6):

$$Q = \arg \min_Q \sum_{k=1}^N \gamma^{N-k} (\mathbf{a}_k - Q\mathbf{u}_k)^2. \quad (2-16)$$

The weight  $\gamma$  is usually a scalar term constrained between 0 and 1, and since it is exponentially decayed, it controls the forgetting rate of the algorithm. Indeed, the closer the sample time  $k$  is to the current value  $N$ , the larger the weight will be, and as a consequence, the prediction error will be taken into higher consideration. In other words, recent measurements influence the estimation of the model  $Q$  exponentially more than older samples. In this way, the model remains sensitive to system changes, regardless of the number of samples already measured. It is then trivial to prove that setting  $\gamma = 1$ , makes the minimization problem identical to the standard LS.

The RLS algorithm can easily be applied to the AO use-case as follows [23]: Let us define the measured wavefront  $\mathbf{a}_k$  as

$$\mathbf{a}_k = G_k \mathbf{q}_k, \quad (2-17)$$

where the control commands are passed as a Kronecker matrix

$$G_k = \mathbf{u}_k^T \otimes I,$$

and the DM influence matrix is vectorized

$$\mathbf{q}_k = \text{vec}(Q_k).$$

The model is then updated following an integral approach, that is, adding to the previous matrix a weighted version of the prediction error, which is here defined as

$$\varepsilon_k = \mathbf{a}_k - G_k \mathbf{q}_k. \quad (2-18)$$

The update recursion then looks like:

$$\mathbf{q}_{k+1} = \mathbf{q}_k + L_{k+1} \cdot \varepsilon_{k+1}, \quad (2-19)$$

with

$$L_{k+1} = P_k G_k^T T_k \quad (2-20)$$

$$P_{k+1} = \frac{1}{\gamma} P_k - \frac{1}{\gamma} L_{k+1} G_k P_k \quad (2-21)$$

$$T_k = \left( \gamma I + G_k P_k G_k^T \right)^{-1} \quad (2-22)$$

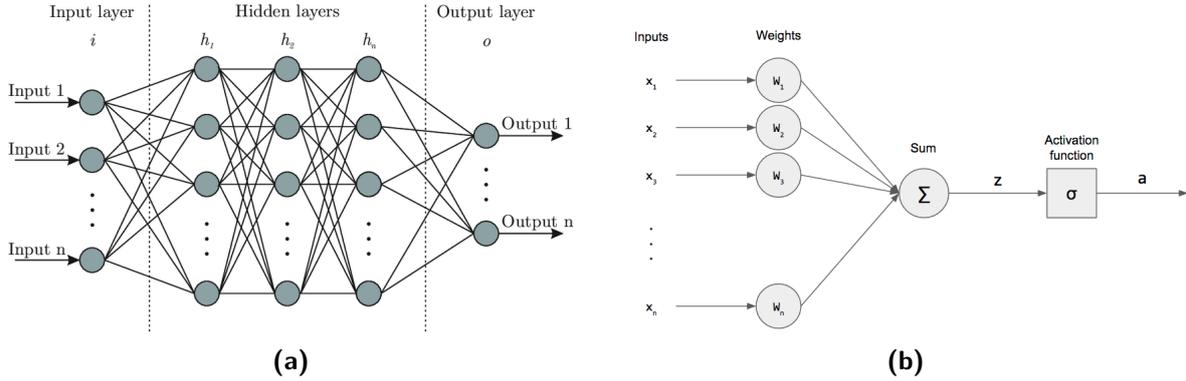
$$0 < \gamma \leq 1 \quad (2-23)$$

as optimal gain parameters. The matrix  $P$  represents the parameters' uncertainty and it is usually initialized as a diagonal matrix. Similarly to the standard LS method, the model parameters are updated based on the prediction error, the model uncertainty, and the control actions. However, in this case, the entries of the covariance matrix  $P$  do not necessarily converge to 0. Indeed, if  $\gamma < 1$ , the term  $\frac{1}{\gamma}$  in 2-21 will force the matrix  $P$  not to shrink excessively, which is a crucial point to keep the model sensitive to changes.

Recursive Least Squares usually presents great convergence speed and good tracking abilities. For these reasons, it has been for years state-of-the-art for the estimation of time-varying systems. Nevertheless, it also has its drawbacks. Firstly, as the traditional non-recursive LS, it does not model nonlinearities and, as a consequence, it inevitably introduces approximation errors when applied to real systems. Secondly, despite the more effective usage of memory storage, RLS presents a higher computational complexity than most other linear regressive methods [22]. This is due to the update process of the covariance matrix, which, for high-dimensional systems like AO, can be a challenging task. Finally, whenever persistent excitation is not provided, the accuracy of RLS drops significantly [24]. Indeed, if the most recent measurements do not excite all system's dynamics, the uncertainty of the parameters increases which eventually makes the model unreliable. This last issue is usually fixed with *ad-hoc* solutions, for instance, setting an upper boundary to the matrix  $P$ , or implementing a method to set  $\gamma = 1$  whenever persistent excitation is not provided [24].

## 2-3 Artificial Neural Networks

ANN are nonlinear adaptive function approximators, which are able to learn specific input-output relationships from data. They include multiple interconnected processing units, called *neurons*, which contribute to the output computation [25]. ANN have proven to be an extremely powerful tool for high-dimensional nonlinear systems, and, over the years, they have been successfully implemented for system identification, control, and classification tasks in plenty of different fields [26]. In the following subsections, a concise description of ANN is given.



**Figure 2-17:** General and specific architecture of a fully connected deep neural network. On the left, the contribution of input and hidden layers' neurons to the output of the network can be appreciated. When the network is fully connected as in this case, the output of each neuron is mapped to the input of all the neurons of the following layer. Trivially, the modeling capabilities of the network are directly correlated to the number of neurons and layers. On the right, instead, the way each neuron computes its output based on the multiple inputs it receives is depicted. Using a nonlinear activation function, it is possible to make the neural network models nonlinearities. In this regard, plenty of different functions can be adopted, depending on the task the network is meant to achieve.

### 2-3-1 Working principle

A neural network can be defined as a function  $q$ , computing an output  $y$  based on an input  $x$ , and internal parameters  $\theta$ :

$$y = q(\theta, x).$$

Each internal neuron also represents a function,  $h(x, \theta, \sigma)$ , defined as:

$$h(x, w) = \sigma \left( \sum_{i=0}^N \theta_i x_i \right). \quad (2-24)$$

Indeed, every single neuron takes a set of inputs  $x_i$ , computes a weighted sum of them based on the current weights  $\theta_i$ , applies an activation function  $\sigma$  on the calculated value, and thus finally obtains a single output (Figure 2-17a). The activation function introduces nonlinearity in the model and ensures that the neural network is able to identify nonlinear behaviors. Based on the implementation task, different functions can be adopted, such as the Rectified Linear Unit, the hyperbolic tangent, or the sigmoid.

Neurons can then be grouped in parallel to form a *layer*. The latter might also be stacked in series with other layers to form what is called a *deep* neural network (Figure 2-17b). The output of the network is then the result of the given set of inputs  $x_i$  passing through concatenated neurons' functions  $h_k(x_i, \theta_i)$ .

### 2-3-2 Training

In order to properly approximate a nonlinear input-output map, the weights of each neuron need to be accurately estimated from the available data. This process is called *training* and it usually follows the backpropagation method [27]. The algorithm first computes the error

gradient with respect to all the neurons' weights. Then, it updates the network parameters following a gradient descent approach:

$$\theta_k = \theta_{k-1} - \alpha \cdot \nabla_{\theta} f_k(\theta, x, y),$$

with  $\alpha$  being the update step size and  $f(\theta, x, y)$  the loss function to be minimized. These two steps are then repeated iteratively until convergence. Therefore, the way the error function is defined has a significant impact on the accuracy and convergence speed of the training process. For regression problems like system identification, the mean squared error loss function is usually chosen [28]

$$f(\theta, x, y) = \frac{1}{K} \sum_{k=1}^K (y_k - q(\theta, x_k))^2$$

where  $K$  is the number of training samples.

### ADAM optimizer

Over the years, plenty of different gradient descent variations have been developed regarding the update of networks' weights. In this regard, Adaptive moment estimation (ADAM) optimizer is currently considered state-of-the-art for the training of ANN, given its fast convergence speed, good accuracy, and noise robustness [29].

The loss function to be minimized is defined here as  $f(\theta)$ , where the dependency on input-output data is made implicit. The ADAM optimizer updates the parameters  $\theta$  at every stochastic realization of the function  $f_k(\theta)$ , exploiting the gradient  $g_k = \nabla_{\theta} f_k(\theta)$ . Differently from other Stochastic Gradient Descent (SGD) algorithms, however, ADAM does not directly use the stochastic realization of the gradient but instead exploits estimates of the 1<sup>st</sup> and 2<sup>nd</sup> order moments of it, that is, the mean ( $m_k$ ) and the uncentered variance ( $v_k$ ). In Table 2-3, a schematic description of the algorithm flow is shown, taken from [29].

Three main features uniquely define the ADAM optimizer. Firstly, the learning rate is adapted to each individual parameter, which is crucially beneficial for sparse systems where neurons have different changing rates. Secondly, using an estimate of the mean of the gradient  $m_k$ , instead of the real stochastic one  $g_k$ , reduces the sensitivity to noise and significantly enhances training stability. Finally, scaling the learning rate  $\alpha$  with the ratio between the first and the second-order moment of the gradient guarantees fast convergence speed, and good noise robustness. Indeed, when the network is continuously acquiring new relevant information, we will have that  $\frac{m_k}{\sqrt{v_k}} \approx \pm 1$  and the learning rate will not be reduced. On the other hand, when an optimum is reached, the gradient  $g_k$  will oscillate around 0, and as a consequence, the ratio  $\frac{m_k}{\sqrt{v_k}}$  will tend to 0, *de facto* not updating the parameter anymore.

### 2-3-3 Direct Inverse Control

As shown in section 2-2, accurate WF correction is obtained if and only if the controller model is a good approximation of the inverse of the AO system. If the latter is nonlinear, computing its inverse can be a challenging task. In this regard, plenty of methods have been published over the years on the control of dynamical systems using neural networks. In particular,

**ADAM Algorithm:**  $g_k^2$  indicates the elementwise square  $g_k \odot g_k$ . All operations on vectors are element-wise. With  $\beta_1^k$  and  $\beta_2^k$  we denote  $\beta_1$  and  $\beta_2$  to the power  $k$ .

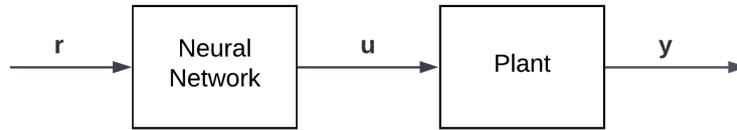
Step	Description
<b>Require:</b> $\alpha$	Stepsize
<b>Require:</b> $\beta_1, \beta_2 \in [0, 1)$	Exponential decay rates for the moment estimates
<b>Require:</b> $f(\theta)$	Stochastic objective function with parameters $\theta$
<b>Require:</b> $\theta_0$	Initial parameter vector
$m_0 \leftarrow 0$	Initialize 1 <sup>st</sup> moment vector
$v_0 \leftarrow 0$	Initialize 2 <sup>nd</sup> moment vector
$k \leftarrow 0$	Initialize timestep
<b>while</b> $\theta_k$ not converged <b>do</b>	
$k \leftarrow k + 1$	
$g_k \leftarrow \nabla_{\theta} f_k(\theta_{k-1})$	Get gradients w.r.t. stochastic objective at timestep $k$
$m_k \leftarrow \beta_1 \cdot m_{k-1} + (1 - \beta_1) \cdot g_k$	Update biased first moment estimate
$v_k \leftarrow \beta_2 \cdot v_{k-1} + (1 - \beta_2) \cdot g_k^2$	Update biased second raw moment estimate
$\hat{m}_k \leftarrow m_k / (1 - \beta_1^k)$	Compute bias-corrected first moment estimate
$\hat{v}_k \leftarrow v_k / (1 - \beta_2^k)$	Compute bias-corrected second raw moment estimate
$\theta_k \leftarrow \theta_{k-1} - \alpha \cdot \hat{m}_k / (\sqrt{\hat{v}_k} + \epsilon)$	Update parameters
<b>end while</b>	
<b>return</b> $\theta_k$	Resulting parameters

**Table 2-3:** Schematic description of ADAM Algorithm.

Direct Inverse Control (DIC) [30][31] is an intuitive and straightforward technique that has been successfully implemented in multiple fields like aerospace [32], unmanned vehicles [33], mechanics [34][35], and chemistry [36].

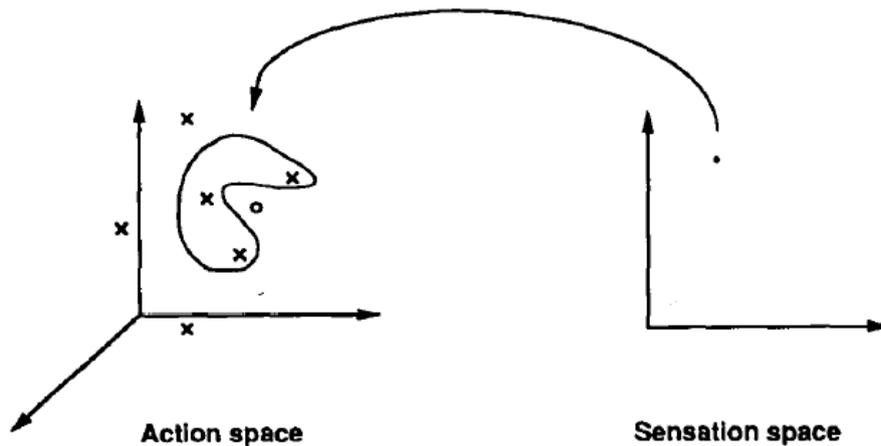
Its principle consists of training a neural network as the inverse of a system and using it to control the plant (Figure 2-18). Consistently with traditional control theory notation [37], the term *plant* is here used to indicate the combination of process and actuators, and it is thus often used as a synonym for *system*. Depending on the loss function being used to train the controller network, two distinct DIC methods can be followed.

**Generalized learning** The most straightforward approach consists of directly training the controller on historical data such as to predict the input actions  $\mathbf{u}$  which generated the output  $\mathbf{y}$ . Despite its intuitive and effective structure, this approach presents two main drawbacks which might hinder its success. Firstly, when the system is characterized by a many-to-one mapping between actions and output, DIC might not be able to find a correct inverse of the plant. Indeed, when multiple exact inverse models of the system exist, as is the case for underdetermined systems, the DIC method provides an average of them, accordingly to the



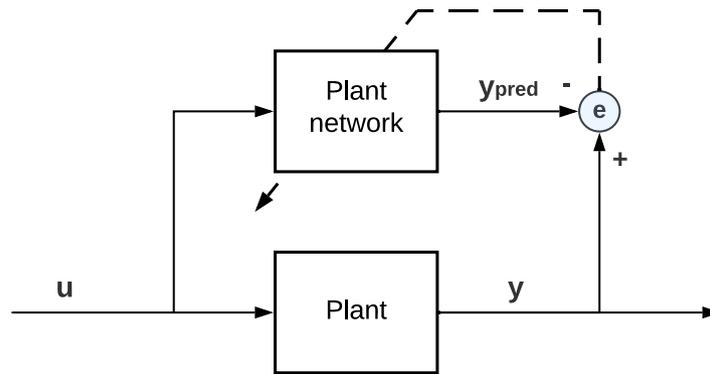
**Figure 2-18:** Direct Inverse Control diagram. The controller consists of a neural network that models the inverse of the plant being controlled. When a reference signal  $r$  is requested, the ANN computes a control action  $u$ , which is then sent to the actuators of the system. If the network properly models the inverse of the plant, the output variable  $y$  will then resemble the reference  $r$ .

least squares principle. However, when the system is nonlinear, inverse images can belong to a non-convex set, in other words, the average of the solutions might not necessarily be an exact inverse of the system [4]. An explicative example of this issue is reported in Figure 2-19.

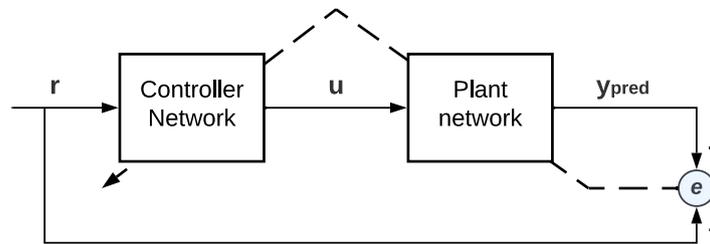


**Figure 2-19:** The convexity problem. The region on the left is the inverse image of the point on the right. The arrow represents the direction in which the mapping is learned by direct inverse modeling. The three points lying inside the inverse image are averaged by the learning procedure, yielding the vector represented by the small circle. This point is not a solution, because the inverse image is not convex. [4]

The second drawback of generalized learning is that the learning process is not goal-directed, since the minimized training error does not consider the actual output of the plant. As a consequence, the controller cannot be optimized for a specific state trajectory, and the effectiveness of the method relies on a thorough and uniform sampling of the state space during training. This is a major hindrance mostly for the control of systems having large action and state sets, for which specialized learning is beneficial.



**Figure 2-20:** Diagram of the training of the forward plant network. The network is placed in parallel with the plant and thus receives the same input actions  $\mathbf{u}$ . The prediction error  $e$  between the plant output  $\mathbf{y}$  and the network estimate  $\mathbf{y}_{pred}$  is backpropagated and used to update the weights of the ANN. The training can either be done offline or online, during control operations.



**Figure 2-21:** Diagram of the training of the controller network, following the *specialized learning* approach. The error between the desired and the predicted output is backpropagated through the plant network, and exploited to update the controller one. In essence, the latter is trained to be the inverse of the plant network and thus make the overall system replicate an identity function.

**Specialized learning** This second alternative method consists of training the controller network to minimize the output error, exploiting a forward model of the plant. This allows to solve the convexity issue and make the learning process goal-directed. Firstly, a shallow neural network is trained to predict the output  $\mathbf{y}$ , based on an input  $\mathbf{u}$ . Secondly, the controller network is placed in series with the plant one and the overall system is trained such as to replicate an identity transfer function. During the error backpropagation, only the controller weights are updated, while the plant network is exclusively used to translate the output loss gradient in the control action domain. In this way, the overall model learns to track the reference signal  $\mathbf{r}$ , and, as a consequence, the controller becomes a good approximation of the inverse of the plant. Therefore, the quality of the controller, and thus of the reference tracking operation, is closely related to the accuracy of the forward plant model. Indeed, if the latter presents modelization errors, these will inevitably be present in the controller as well.

#### 2-3-4 ANN in Adaptive Optics

Given the complexity and high-dimensionality of Adaptive Optics systems, ANN have represented an appealing modeling technique since the early stages of AO in 1990 [38]. Over the years, ANN have been successfully implemented for various challenging tasks like wavefront

sensing [39], WF reconstruction [40], and turbulence prediction [41]. Regarding the identification and control of deformable mirrors, ANN have been mostly implemented on Wavefront Sensorless (WFSless) applications, where time-consuming random-search optimization is required to retrieve WFs, and thus neural networks can significantly speed up control [42]. In particular, Reinforcement Learning (RL) techniques have shown excellent robustness and good converging speed, even in model-free implementations [43].

However, one of the major drawbacks of Machine Learning (ML) methods is their inherent complexity with respect to traditional linear algorithms. Indeed, in order to outperform the existing state-of-the-art, non-trivial network structures are often adopted like Long-Short Term Memory Networks [44], autoencoders [45] and Convolutional Neural Network [40]. Nevertheless, shallow ANNs have also been researched, in particular for the identification of DM influence models, which often present nonlinearities and sometimes time-varying behaviors [46]. In [47], a DIC-inspired approach is proven to be able to accurately control a simulated AO system in the presence of misalignments. However, in order to train the ANN using the output reconstruction error without implementing a plant model network, the proposed framework runs a model-free estimation of the unknown gradient of the loss function on the control actions. This is a rather noisy and time-consuming computation, which requires additional modules to run efficiently, such as a model-actor parameter-sharing framework. Furthermore, the gradient estimation process has to be initialized in every new control loop and requires multiple measurements to reach convergence. Therefore, it is not a suitable method to pursue fast, possibly open-loop, WF correction objectives.

---

## Chapter 3

---

# Method proposal

The main idea behind this thesis work is to exploit neural networks' ability to identify non-linear systems while training online, to improve the control of Adaptive Optics (AO) that present nonlinearities and time-varying behavior. High importance is given to the complexity of the proposed algorithm, which should not represent an obstacle for the practical implementation of the method. Indeed, the given approach should be a valid and appealing alternative to traditional linear control algorithms. In this regard, an AO-specific version of the Direct Inverse Control (DIC) method, described in Section 2-3-3, is proposed here.

### 3-1 Adaptive Direct Inverse Control for Adaptive Optics

The approach presented here is based on the DIC *specialized learning* method, which consists of identifying a forward model of the plant, and then computing the inverse of it through the backpropagation of the output error. The proposed method can thus be analyzed as the combination of two main processes, described more in detail here below:

#### 3-1-1 Forward model identification

Similarly to the Least Squares method introduced in Section 2-2, the plant model is initially estimated using random input-output calibration samples, collected offline. The actuators' signals are sampled following a uniform distribution over the whole action space. The loss function used to train the network is the mean squared error between the predicted wavefront

and the measured one:

$$\mathcal{L}_q(\theta_q, \mathbf{a}, \mathbf{u}) = \frac{1}{K} \sum_{k=1}^K (\bar{e}_k)^2 \quad (3-1)$$

$$= \frac{1}{K} \sum_{k=1}^K (\mathbf{a}_k - \bar{\mathbf{a}})^2 \quad (3-2)$$

$$= \frac{1}{K} \sum_{k=1}^K (\mathbf{a}_k - q(\theta_q, \mathbf{u}_k))^2 \quad (3-3)$$

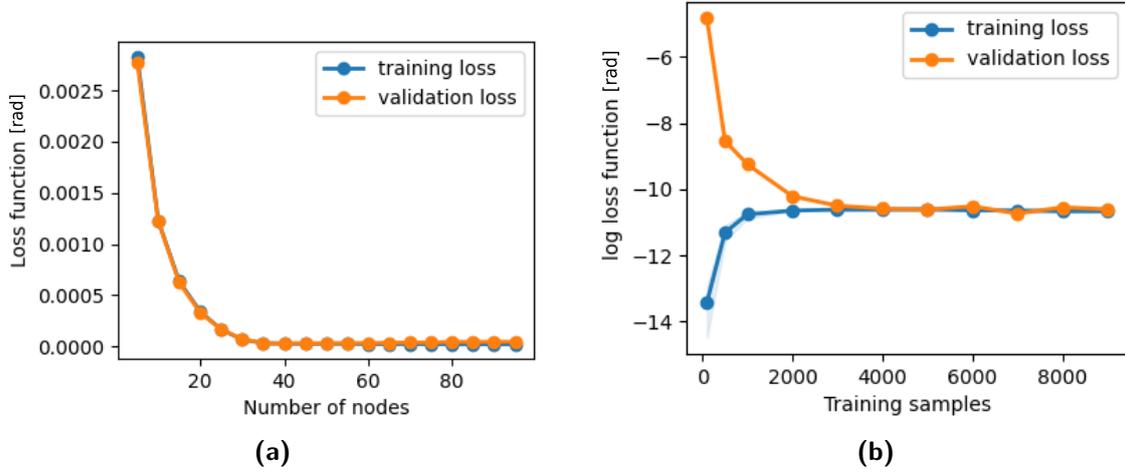
where  $q$  is the plant network and  $\bar{\mathbf{a}}$  is the predicted wavefront.

The implemented shallow Neural Network has  $C$  inputs and  $D$  outputs, respectively equal to the number of mirror actuators and wavefront (WF) modes. The difference with the linear Least Squares method exclusively stands in the internal structure of the identified model. Consistently with the objective of this project, the complexity of the network is set to be as low as possible. Following Occam's Razor, the number of hidden layers is set to one since adding more complexity proved not to result in any performance improvement. Wavefront prediction tests are then run for different numbers of layer nodes, which showed that using the same depth as the output layer represents the most effective choice (Figure 3-1a). Nonlinear activation functions are only applied in the hidden layer, and in this regard, the hyperbolic tangent is applied as it proves to obtain the lowest *rms* WF prediction error  $e$ . For more details regarding this metric, the reader can refer to Section 2-2-2.

The number of samples used to train the network also significantly affects the performances of the identified model, and it is thus optimized through iterative tests. In particular, the network is trained on random sets of different sizes, while the other hyperparameters such as the learning rate and the number of epochs are always set to guarantee stable training convergence. As can be seen in Figure 3-1b, the network stops overfitting when the training set consists of more than 4000 samples, which is thus taken as training size. Finally, the ADAM algorithm described in Section 2-3-2 is chosen as the update method for the weights of the network.

### 3-1-2 Forward model inversion

Finding the inverse of a nonlinear function is not a trivial task. The most useful feature of DIC using *specialized learning* is that it approaches this matter with a rather straightforward solution. The controller network is placed in series with the plant one, and the overall system is trained to replicate an identity transfer function, that is, to obtain an output  $\mathbf{a}$  as much similar to the reference  $\mathbf{a}_r$  as possible (Figure 3-3). During training, the plant network is not updated and it is indeed only used to back-propagate the output wavefront error to the controller. In this regard, the minimized loss function is chosen as the root-mean-squared error between the reference WF and the WF predicted by the plant network, on the base of the control actions computed by the controller. Jordan and Rumelhart [4] refer to this metric



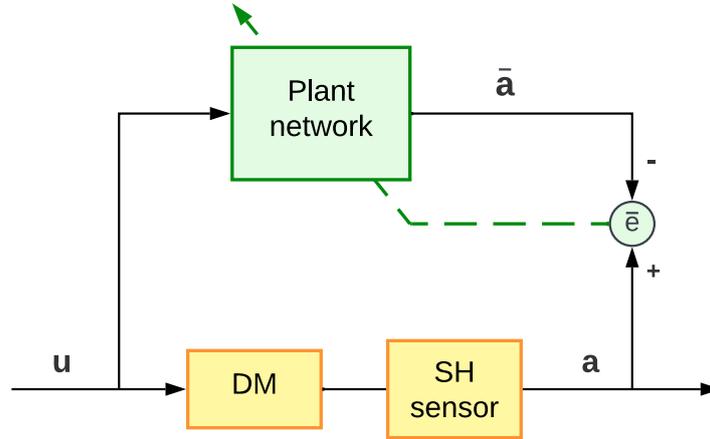
**Figure 3-1:** Hyperparameters tuning curves. On the left, it is plotted the training and validation loss curves for different hidden layer depths. The loss function is minimized when more than 40 nodes are used in the hidden layer. For each test case, enough training samples are provided to make the system not overfit. On the right, instead, the training and validation loss curves for different training set dimensions are depicted. The network stops overfitting when more than 4000 samples are used for training.

as the *predicted performance error*:

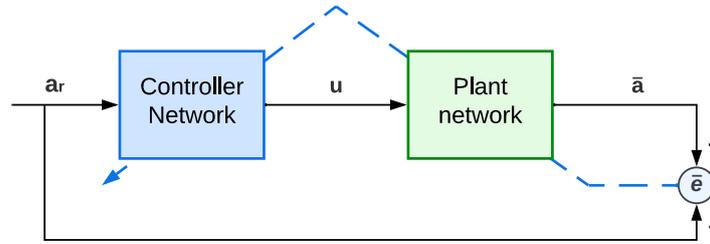
$$\begin{aligned}
 \hat{\mathcal{L}}_c(\theta_c, \theta_q, \mathbf{a}_r) &= \frac{1}{K} \sum_{k=1}^K (\bar{\mathbf{e}}_k)^2 \\
 &= \frac{1}{K} \sum_{k=1}^K (\mathbf{a}_{r,k} - \bar{\mathbf{a}})^2 \\
 &= \frac{1}{K} \sum_{k=1}^K (\mathbf{a}_{r,k} - q(\theta_q, \mathbf{u}_k))^2 \\
 &= \frac{1}{K} \sum_{k=1}^K (\mathbf{a}_{r,k} - q(\theta_q, c(\theta_c, \mathbf{a}_r)))^2
 \end{aligned}$$

where  $c$  is the controller network. Computing the error using the predicted WF rather than the measured one allows us to train the controller offline, and thus significantly speed up the inversion process. On the other hand, it must be said that in this way the controller network is not able to compensate for inaccuracies in the plant model, as it is simply an inverted version of it.

As it will be also mentioned in the dedicated Section 2-2-2, the AO setup used to validate the proposed method consists of a deformable mirror having 96 channels, and a Shack-Hartmann sensor characterized by 180 centroids, which suggests not using more than 60 wavefront modes to represent the measured aberrations [48]. As a consequence, the inverse model of the plant is an underdetermined system, that is, it contains more unknown variables (control actions), than equations (WF modes). A visual representation of such a many-to-one mapping is reported in Figure 2-19. It is here important to remark that since the measured wavefronts



**Figure 3-2:** Diagram representing the offline learning architecture of the plant network. The ANN is trained through the backpropagation of the prediction error  $\bar{e}$ , here represented by a green dashed line. If the training is successful, the plant network should be able to replicate the WF generation behavior of the system.



**Figure 3-3:** Diagram representing the offline learning architecture of the controller network. Exploiting the pre-trained plant network, the controller model is identified through the backpropagation of the predicted performance error  $\bar{e}$ . The plant network is necessary to express the loss gradients as a function of the control actions  $\mathbf{u}$ , and thus to make the controller network be able to train on the output error. During this process, the plant network is not updated and it is only used to backpropagate the error. If the training is successful, the overall system composed of the controller and plant networks should resemble an identity function.

are reconstructed with a finite number of modes, in particular lower than the number of actuators, different control action commands trigger mirror's deformations that get measured as the exact same WF. This poses a serious issue for model inversion since every wavefront has to be mapped to only one control action. In the linear Least Squares case described in Section 2-2, Singular-Value-Decomposition (SVD) can be used, which maps each output to its minimum norm image. Following a similar approach, the training loss function of the controller network is here expanded with an L2 regularizing term:

$$\mathcal{L}_c(\theta_c, \theta_q, \mathbf{a}_r) = \hat{\mathcal{L}}_c(\theta_c, \theta_q, \mathbf{a}_r) + \alpha \cdot \sum_{j=1}^T \theta_{c,j}^2 \quad (3-4)$$

where  $\theta_{c,j}$  are the weights of the controller network and  $\alpha$  is the L2 regularizing gain. In this way, if  $\alpha$  is set to a positive number, the controller network is induced to map each WF to the control solution having the lowest norm.

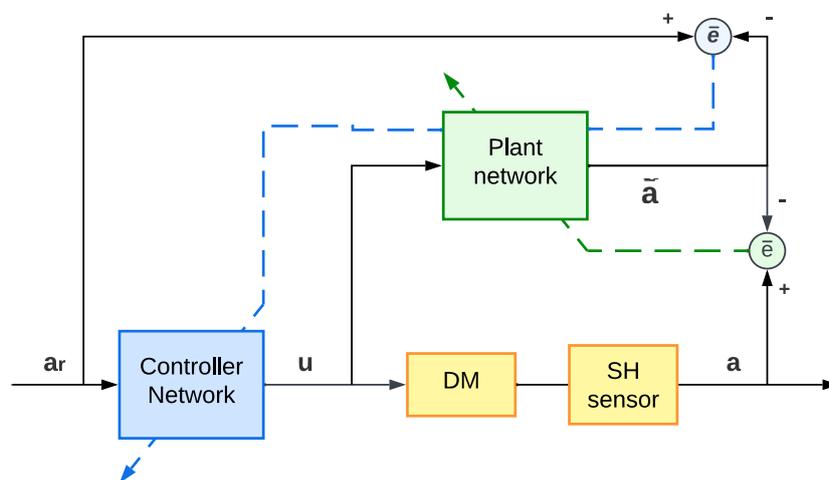
The hyperparameters of the controller network are optimized following the same tests as for the forward model. Since similar results are obtained, the structure of the two networks is made identical, except for the dimension of the input and output layers, which are inverted. Regarding the training process, it is important to remark that the controller network guarantees good reconstruction abilities only for those types of WFs observed during training, therefore, it is extremely important to show to the network those WFs samples that are more likely to be encountered in practice. The optimal number of training samples is significantly dependent on the range of wavefronts that are meant to be controlled. More insights on this regard are reported in the following chapter.

### 3-1-3 Online update of forward and inverse models

The proposed approach can also be used online, updating both the forward and the inverse model during control operations. The training procedure is essentially the same as the one described in the previous sections, with the difference that both networks are updated with every new gathered measurement, through ADAM Stochastic Gradient Descent (SGD). In Table 3-1, a summary of the online algorithm is provided.

Step	Description
<b>Require:</b> $q(\theta_{q,0})$	Pretrained plant network
<b>Require:</b> $c(\theta_{c,0})$	Pretrained controller network
<b>Require:</b> $\alpha_q$	ADAM SGD learning rate for the plant network
<b>Require:</b> $\alpha_c$	ADAM SGD learning rate for the controller network
$k \leftarrow 0$	Initialize timestep
<b>while</b> wavefront control is active <b>do</b>	
$k \leftarrow k + 1$	
$\mathbf{a}_{r,k} \leftarrow$	Obtain new reference wavefront
$\mathbf{u}_k \leftarrow c(\theta_c, \mathbf{a}_{r,k})$	Compute control actions through inverse model
$\mathbf{a}_k \leftarrow$	Send control actions to DM and measure reproduced wavefront
$\bar{e}_k \leftarrow \mathbf{a}_k - \bar{\mathbf{a}}_k$	Compute prediction error
$\theta_{q,k} \leftarrow \text{ADAM}(\theta_{q,k-1}, \bar{e}_k)$	Update forward model
$\bar{\mathbf{e}}_k \leftarrow \mathbf{a}_{r,k} - \bar{\mathbf{a}}_k$	Compute predicted performance error
$\theta_{c,k} \leftarrow \text{ADAM}(\theta_{c,k-1}, \bar{\mathbf{e}}_k)$	Update inverse model
<b>end while</b>	

**Table 3-1:** Schematic summary of the Online DIC algorithm for AO.



**Figure 3-4:** Diagram representing the online learning architecture of both the plant and controller network. Once a new WF measurement is obtained, the forward model is firstly updated using the performance error  $\bar{e}$  (green line). Secondly, the predicted performance error  $\hat{e}$  is back-propagated through the plant network and exploited to update the inverse model (blue line).

---

# Chapter 4

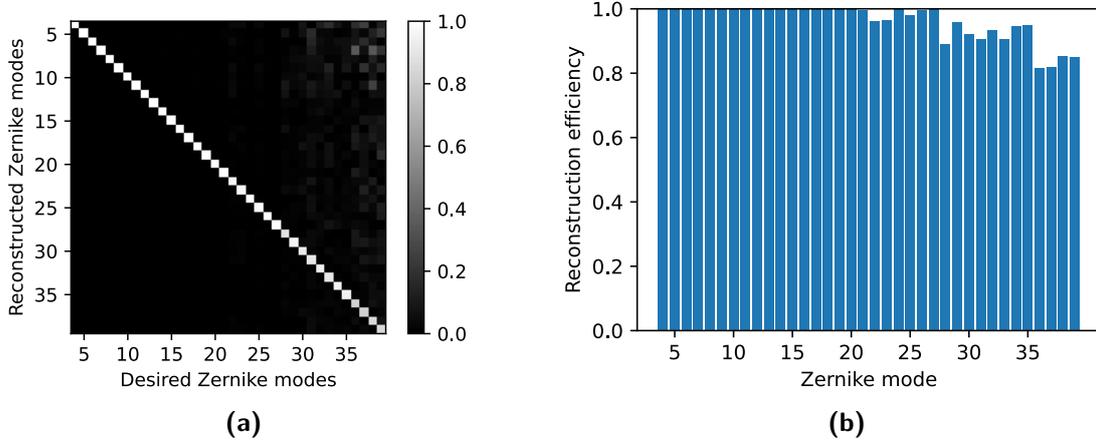
---

## Results

The performances of the proposed method are assessed using an experimental Adaptive Optics (AO) setup, located in the AO Lab of the Delft Center for Systems and Control (DCSC) at TU Delft. In Section 4-1, the details regarding the implementation of the control algorithm on the AO equipment are given, together with a description of the investigated test cases and the methods used as a benchmark. The results obtained by the proposed approach are reported in three different sections, based on the issue being highlighted. Wavefront prediction performances are discussed in Section 4-2, while insights into the model inversion process and the impact of regularization are reported in Section 4-3. In Section 4-4, results concerning WFs correction are then presented, with a particular focus on nonlinearities and time-varying conditions. Finally, in Section 4-5, considerations regarding the obtained results are reported.

### 4-1 Experimental conditions

A technical description of the AO setup used to validate the method can be found in section 2-2-2 (Figure 2-12). The plant network used to model the system presents an input layer with 96 nodes, equal to the number of deformable mirror (DM) actuators, one hidden layer having 40 nodes and  $\tanh$  activation function, and an output layer also with 40 nodes, equal to the number of Zernike modes used to reconstruct wavefronts. Consistently with the test results presented in Section 3-1-1, the plant network is firstly trained offline for 10000 epochs on 4096 random input-output samples gathered with the setup. The learning rate is set to 0.005. The structure of the controller network is set to be exactly symmetric to the plant network, and additionally, all the weights are regularized during training with an L2 coefficient  $\alpha = 10^{-5}$ . This is to enforce the network to map wavefronts to their minimum norm solution, given the system is underdetermined. The impact that the regularization has on the reconstruction efficiency is not significant for most of the modes, as shown in Figure 4-1. The controller network is trained for 20000 epochs on 10240 wavefront samples, generated by randomly exciting the actuators of the mirror. This number is the smallest able to guarantee an ideal reconstruction efficiency of at least 98% on the first 20 Zernike modes. Finally, in this case, the learning rate of the ADAM optimizer is set to 0.001.



**Figure 4-1:** Reconstruction efficiency of the proposed method, for an L2 regularizing term  $\alpha = 10^{-5}$ . The trained identity model  $i(\theta_c, \theta_q) = q(\theta_q, c(\theta_c))$  is inspected. Each one of the first 40 Zernike modes is fed to the system with unitary amplitude. The output of the identity model, which represents the theoretical system reconstruction of the requested wavefront in terms of Zernike magnitudes, is reported on the y-axis. Ideally, only the desired mode should be reproduced, with unitary amplitude. The histograms on the right depict the diagonal elements of the grids on the left-hand side. They represent the reconstruction efficiency regarding the desired mode only, that is, not considering the *cross-talking* between modes. The larger the regularizing gain, the more the reconstruction quality is affected, especially for high-order modes.

#### 4-1-1 Investigated test cases

Wavefront reconstruction performances are essentially dependent on two matters: the quality of the identified system model and the efficiency of the model inversion. The first point is investigated in Section 4-2, where the wavefront (WF) prediction accuracy of the proposed method is compared with the state-of-the-art. The second topic, that is, the identification of the inverse model is instead discussed in Section 4-3. In particular, the impact of L2 regularization on the ideal reconstruction efficiency is analyzed, in order to understand how the risk of saturation can be mitigated, and how this technique relates to the Singular-Value-Decomposition (SVD) truncation practice.

Once these two preliminary points have been discussed, results regarding WF reconstruction are reported. Generally speaking, wavefronts can be corrected by either following an open or closed-loop approach, depending if the residual aberration is fed back to the controller or not. Trivially, the quality of the model plays a crucial role in open-loop implementations, where linearization errors and neglected dynamics are not compensated for. As a consequence, in order to better highlight methods' performances in nonlinear environments, the open-loop reconstruction of small and large-phase WFs is tested, and the results are reported in Section 4-4-1. Nevertheless, in practice, closed-loop control is usually implemented to improve the quality and stability of the reconstruction. Therefore, large-phase Zernike modes are also generated in closed-loop, and insights regarding the convergence speed and accuracy of the proposed method are obtained. Finally, in Section 4-4-2, the tracking abilities of online Direct Inverse Control (DIC) are evaluated. In order to keep the system's response linear,

only low-phase wavefronts are reconstructed. Indeed, the possible performance degradation should here be exclusively caused by a change in the plant's dynamics. Experimentally, this is induced by increasing the environmental temperature through a heating element.

In table 4-1, a brief schematic summary of the investigated test cases is given:

Task	Triggered condition	Benchmarks	Section
WF prediction	Nonlinearity	LS	4-2-1
WF prediction	Temperature changes	LS, RLS	4-2-2
WF prediction	Lack of persistent excitation	LS, RLS	4-2-3
Model inversion		SVD	4-3
WF correction	Nonlinearity	LS	4-4-1
WF correction	Temperature changes	LS, RLS	4-4-2

**Table 4-1:** Summary of the investigated test cases. LS and RLS respectively stand for Least Squares and Recursive Least Squares.

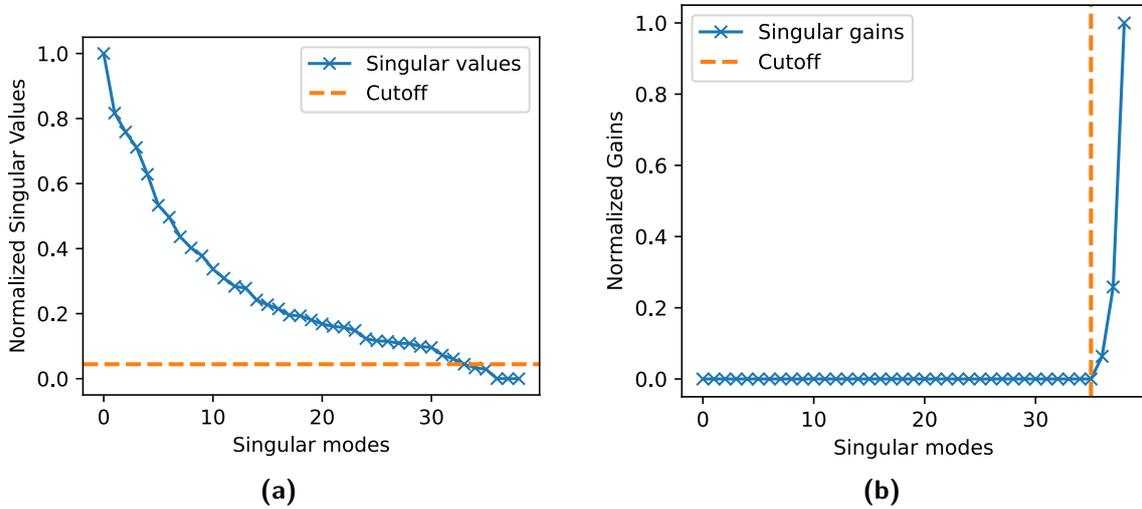
Throughout the chapter, three main error signals will be used to present and discuss the results: the *performance error*, the *prediction error*, and the *predicted performance error* (Table 4-2).

Error	Error in WF domain	Name	Source
$\mathbf{e}$	$\mathbf{a}_r - \mathbf{a}$	performance error	environment, environment
$\bar{e}$	$\mathbf{a} - \bar{\mathbf{a}}$	prediction error	environment, model
$\bar{\mathbf{e}}$	$\mathbf{a}_r - \bar{\mathbf{a}}$	predicted performance error	environment, model

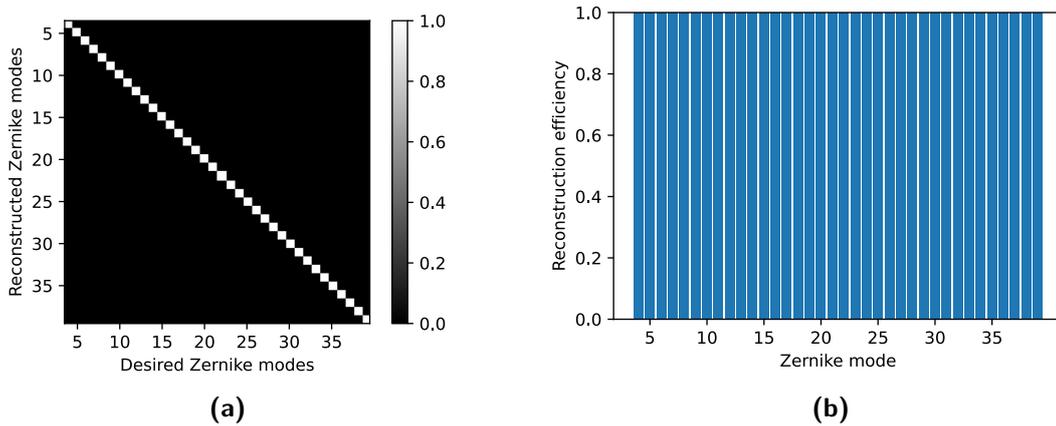
**Table 4-2:** Error signals and their sources.  $\mathbf{a}_r$ ,  $\mathbf{a}$ , and  $\bar{\mathbf{a}}$  are respectively the reference, measured and predicted WFs.

## 4-1-2 Benchmarks

As discussed in Section 2-2, the linear influence matrix approach (Eq. 2-5) represents the state-of-the-art for the identification and control of DM-based Adaptive Optics, and therefore, the proposed algorithm is benchmarked against it. In the tests which specifically aim to assess the tracking abilities of the controller in case of time-varying conditions, the proposed approach is also compared against Recursive Least Squares (RLS). Regardless of the method being used, the linear influence matrix is identified using the same 4096 random samples used to train the plant network. Furthermore, in order to maintain the pseudoinverse matrix  $Q^\dagger = S$  well-conditioned, some of the lowest singular modes are set to zero (Figure 4-2). In particular, as shown in Figure 4-2b, a significant magnitude increase happens after the 36<sup>th</sup> normalized gain, and, therefore, the last 4 modes are truncated. It is important to remark that the regularization applied here only concerns noise and does not have a significant impact on the reconstruction efficiency, as can be seen in Figure 4-3.



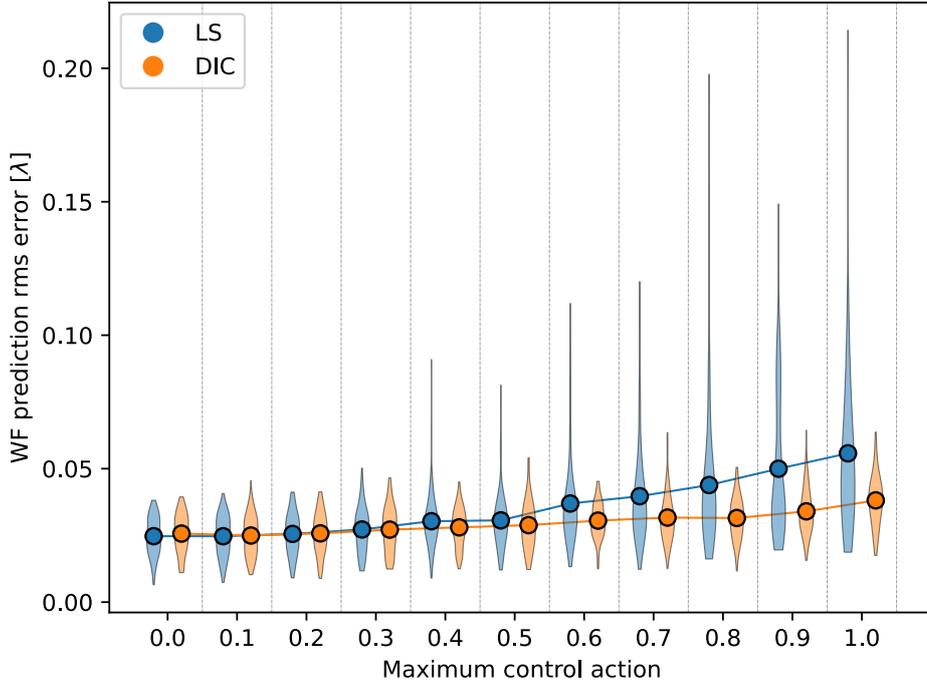
**Figure 4-2:** On the left, the singular values of the DM influence matrix  $Q$  are plotted. The orange cutoff line represents the threshold under which the values are set to zero and thus the modes get truncated. On the right, instead, the singular gains of the DM influence matrix  $Q$  are depicted. Gains are the normalized reciprocal of the singular values. The cutoff line represents the threshold after which the modes get truncated.



**Figure 4-3:** Representation of the similarity matrix  $T = S \cdot Q$  and reconstruction efficiency  $diag(T)$ , in the case of 4 truncated modes. The gray scale represents the absolute value of the matrix, element-wise. All Zernike modes between 3 (top-left) and 39 (bottom-right) are represented. The histograms on the right depict the elements on the diagonal of the similarity matrix and range from Zernike mode 3 to 39.

## 4-2 Wavefront prediction

### 4-2-1 Nonlinear conditions



**Figure 4-4:** Wavefront prediction performance for different control magnitudes. 500 random control action sets are randomly sampled and sent to the DM. The measured WFs are then compared with the ones predicted by the proposed method and the Least Squares (LS) state-of-the-art. On the y-axis is reported the averaged wavefront prediction rms error, while the x-axis shows the maximum actuator control action used to generate the given wavefronts. In the figure, circles represent the average prediction error for the given control action, while the colored areas depict the error distribution. For small actuator signals, the performances of DIC and LS are comparable. For large phase WFs, instead, the proposed method outperforms the state-of-the-art, both in terms of residual error and error deviation.

The main difference between the proposed DIC-based method and the linear LS approach consists of the well-known ability of Artificial Neural Networks (ANN) to model nonlinearities. In order to quantify the advantage gained from this characteristic, the forward plant network is tested against the linear influence matrix on a WF prediction task. The metric used to evaluate performances is the root-mean-square (rms) of the prediction error  $\bar{e}$ :

$$\bar{e} = \|\mathbf{a} - \bar{\mathbf{a}}\|_2,$$

where  $\mathbf{a}$  and  $\bar{\mathbf{a}}$  are vectors containing the modal coefficients of the measured and predicted WFs respectively. For more insights regarding the rms metric, the reader can refer to Section 2-2-2.

The experiment is rather straightforward and consists of sending a random set of control actions to the mirror, predicting the WF that will be measured by the Shack-Hartmann (SH) sensor, and then comparing the estimated WF with the measured one. The process is then repeated for 500 random control actions of various magnitudes. The obtained results are reported in Figure 4-4.

The plant network proves to outperform the linear influence matrix in predicting WFs that require medium and especially large control magnitudes. Small phase WFs, on the other hand, are predicted with comparable accuracy. This behavior is consistent with the DM response curve reported in Figure 2-14, which shows that the quadratic assumption used to linearize the system is not satisfied for large control magnitudes. Furthermore, the Direct Inverse Control method also shows a more stable error deviation  $var(\bar{e})$  than the benchmark, which instead presents rather unstable predicting performances for very large control magnitudes.

#### 4-2-2 Time-varying conditions

As explained in Section 2-2-2, membrane DMs are sensitive to external conditions like humidity and temperature, both in terms of bias deformation and response behavior. In this regard, the tracking abilities of the proposed method are here assessed. The experiment is rather similar to the one discussed in the previous section and indeed consists of sending random control actions to the DM and predicting the measured WFs. In this case, however, the maximal magnitude of the random control actions is set to 0.3, in order not to trigger nonlinearities. Furthermore, the temperature around the AO setup is increased up to  $40^{\circ}C$ . The results of the test are reported in Table 4-3.

One first important remark is that the linear time-invariant model identified through Least-Squares becomes unreliable once the working point has moved from the calibration one. The prediction error  $\bar{e}$  relative to LS increases by a factor of 20, while on the other hand, RLS and the proposed DIC-based method both present very limited effects on performances, which are most likely due to the lower signal-to-noise ratio induced by the increase in temperature. Indeed, the predicted error  $\bar{e}$  only increases by a factor of 1.9, for both methods.

Model identification method	$23^{\circ}C$	$30^{\circ}C$	$35^{\circ}C$	$40^{\circ}C$
LS	0.311 $\lambda$	2.120 $\lambda$	4.512 $\lambda$	6.523 $\lambda$
RLS	0.311 $\lambda$	0.398 $\lambda$	0.510 $\lambda$	0.585 $\lambda$
DIC	0.315 $\lambda$	0.399 $\lambda$	0.512 $\lambda$	0.591 $\lambda$

**Table 4-3:** WF root-mean-square prediction error  $e$  obtained by the proposed method and benchmarks at different temperatures. For each test condition, the prediction of 100 random small-phase wavefronts is considered.

### 4-2-3 Lack of persistent excitation

As presented in Section 2-2-3, Recursive Least Squares iteratively updates the influence matrix  $Q$  accordingly with the following equation:

$$Q = \arg \min_Q \sum_{k=1}^N \gamma^{N-k} (\mathbf{a}_k - Q\mathbf{u}_k)^2. \quad (4-1)$$

where  $\gamma$  is a positive forgetting factor lower or equal to 1. Here below, it is assumed that the reader is familiar with the mathematical formulation of the method, which can be found in Section 2-2-3.

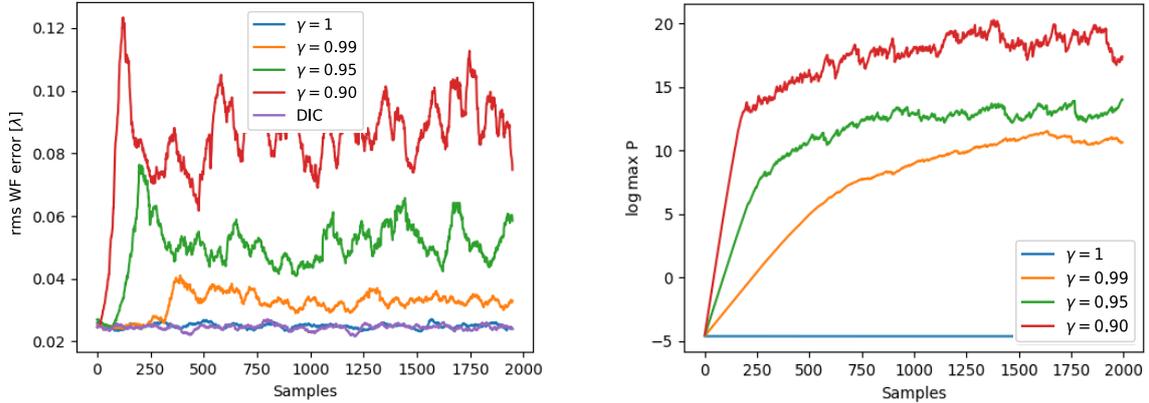
Setting  $\gamma$  to be strictly lower than 1 enhances the priority given to the last measurements with respect to older ones. In particular, doing so helps the covariance matrix  $P$  not to shrink excessively, given its entries are reduced with every new gathered sample. This fact can be better understood by looking at Eq.2-21. Since the gain  $L$  used to update the model  $Q$  is strictly dependent on the uncertainty of the model, setting  $\gamma < 1$  is necessary to maintain the tracking abilities of the system over time. However, when the most recent measurements do not provide the same amount of information as the older data, the covariance matrix  $P$  tends to diverge, thus increasing the update gain  $L$  and eventually making the model unstable.

This matter is here studied through the following experiment. A constant control action set is repeatedly sent to the mirror, and the prediction error  $\bar{e}$  obtained by the proposed method and benchmarks is analyzed. This test case is meant to represent the scenario of the system being in an idle state, which could happen, for instance, when no sample is being analyzed or when the aberration has already been corrected. In this case, DIC is compared against 4 RLS algorithms, each having a different exponential forgetting factor  $\gamma$ . Trivially, the case of  $\gamma = 1$  is equivalent to the traditional LS. Results are reported in Figure 4-5 and clearly show that when persistent excitation is not provided, using  $\gamma < 1$  eventually makes the RLS-estimated model unreliable. On the other hand, the proposed method and the RLS with  $\lambda = 1$  both present stable prediction performances regardless of the information contained in the gathered samples.

This is a well-known problem of RLS, which however does not have a unique, reliable solution. Plenty of extensions have been developed over the years, like periodically resetting the matrix  $P$  [49], using directional forgetting [50], or setting  $\gamma = 1$  whenever persistent excitation is not guaranteed [51]. Nevertheless, all these methods require additional development effort and are rather specific to the considered use case.

## 4-3 Model inversion and ideal reconstruction efficiency

As firstly introduced in Section 2-2, high-quality WF correction is achieved when the controller accurately represents the inverse of the DM model. Assuming the forward network is properly trained to replicate the AO system response, the WF correction performances are then solely dependent on the quality of the inversion process. In this regard, following the framework presented in Section 3-1-2, the controller network is trained on 10240 random



**Figure 4-5:** Recursive Least Squares performances when persistent excitation is not provided. On the left-hand side, the residual wavefront prediction rms error is plotted over time. The x-axis represents time and shows the control iterations. On the y-axis, instead, the rms error is reported, in the wavelength unit. If the forgetting factor  $\gamma$  is set to be strictly lower than 1, the model eventually becomes unreliable and the reconstruction accuracy drops. This is not the case for the proposed DIC method or if RLS is set to have  $\gamma = 1$ . On the right-hand side, the maximum covariance matrix element is plotted. The x-axis represents time and shows the prediction steps. On the y-axis, instead, the maximum entry of the covariance matrix  $P$  is shown, in the logarithmic scale. If the forgetting factor  $\gamma$  is set to be strictly lower than 1, the uncertainty increases and the model eventually becomes unreliable.

wavefronts, without regularization. The identity model  $i(\theta_c, \theta_q) = q(\theta_q, c(\theta_c))$  is then analyzed, inspecting the predicted reconstruction efficiency for all the 40 Zernike modes used to describe WFs. According to the results reported in Figure 4-6a, WF modes are ideally reconstructable with an average rms error of only  $0.044 \lambda$ , which is significantly below the noise level ( $0.126 \lambda$ ) and thus indicates that the controller well models the inverse of the plant network.

Nevertheless, regularizing the controller network can sometimes be beneficial. Indeed, the actuators of membrane DMs can sustain only a limited amount of voltage, after which the actuators saturate, thus introducing nonlinearity in the system's response. In order to avoid this phenomenon, the controller model can be regularized, that is, the magnitude of its gains can be reduced, together with the correction abilities of the controller. In the linear case, introduced in Section 2-2, the SVD technique allows to regularize the system truncating the lowest modes while pseudo-inverting the plant model. In this way, controlling the trade-off between reconstruction efficiency and the risk of saturation is a rather straightforward task, and in this regard, the reader can refer to Figure 2-10.

Similarly to SVD, the L2 regularization process (Eq. 3-4) can also be exploited in the proposed method to reduce the risk of saturating actuators. Trivially, this comes at the expense of the reconstruction ability of the system, which inevitably gets worse the more the system is regularized. In order to gather insights regarding the impact of the penalizing term  $\alpha$ , experimental tests are run. Maintaining constant hyperparameters and training conditions, the network is trained on different L2 gain values, and the corresponding performances are

compared. As can be seen in Figure 4-6, increasing the magnitude of the regularizing term worsens the reconstruction efficiency and increases cross-talk between modes.

To give an example, if the controller network is trained with  $\alpha = 0.05$ , and the wavefront to be reconstructed is a coma with amplitude  $1 \lambda$ , the WF generated in open-loop by the system would ideally be composed by coma with amplitude  $0.89 \lambda$ , vertical trefoil with  $0.03 \lambda$  and other minor modes like  $0.01 \lambda$  of astigmatism. On the other hand, the likelihood of triggering actuator saturation would also be reduced, and in the former example, the maximum actuator signal would change from 125% of the maximum allowed voltage to 74%, thus avoiding saturation. In this regard, Fig. 4-7 provides insights into the impact that the regularizing term  $\alpha$  has on the maximum control action required to generate Zernike modes. In conclusion, similarly to the SVD truncation technique described in section 2-2, it has been shown that the L2 regularization of the weights of the controller network allows controlling the trade-off between efficiency and stability of the WF reconstruction process.

## 4-4 Wavefront correction

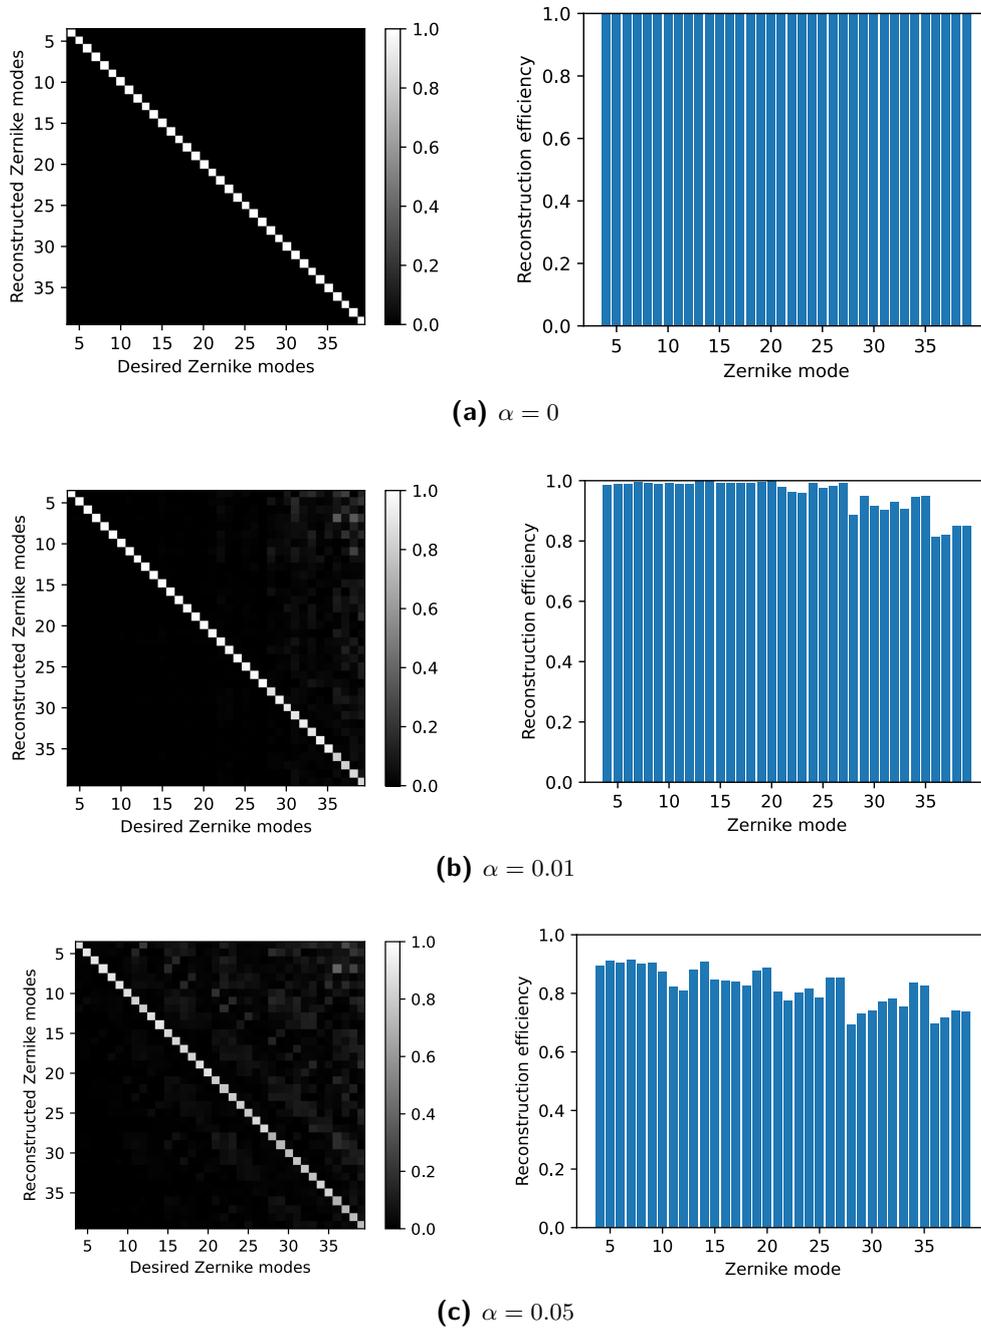
### 4-4-1 Nonlinear conditions

In this Section, the proposed method is compared with the linear state-of-the-art approaches on a various range of wavefronts. Particular interest is posed on the performance error for large phase aberrations, that is, when significant control magnitudes are required. As introduced in Section 2-2-2, this is indeed the working range where most linearization errors arise.

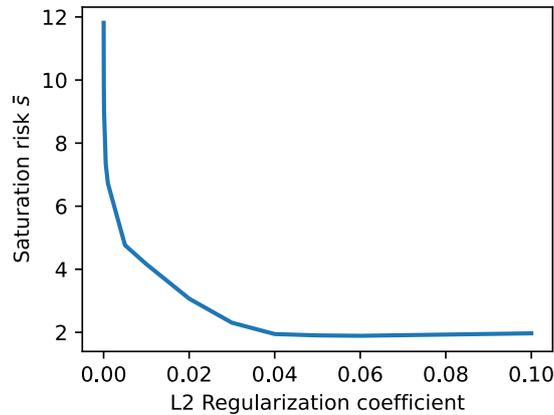
#### Open-Loop reconstruction

Theoretically, if the model used to control the DM exactly represents the inverse of the AO system, a perfect WF reconstruction can be achieved in a feed-forward fashion. Despite multiple measurements and closed-loop iterations being in practice often necessary to minimize the residual WF error, analyzing open-loop reconstruction performances provides useful insights regarding the accuracy of the model being used in the controller. In particular, a metric often used in AO to evaluate control techniques is the maximal wavefront amplitude that can be reproduced with a satisfactory accuracy, where the latter usually stands for a threshold either on the residual rms WF error or the variance of it [20]. This gives a clear indication of the limits of the control algorithm, and the operating range in which satisfactory performances can be achieved.

In this regard, a test is run here which consists in measuring the maximal Zernike mode amplitude that can be reconstructed with a residual rms error lower than  $0.3 \lambda$ . The experiment is run 10 times per mode and the results are then averaged. Figure 4-8 shows the performances of the proposed method and the LS benchmark for the first 10 Zernike modes, excluding piston, tip, and tilt. In order to set up a fair comparison between the proposed method and the linear benchmark, the DIC framework is not updated online, and, in both cases, the regularization tool is set up such as not to impact the reconstruction efficiency of the first 20 Zernike modes, as can be seen in Figures 4-1 and 4-3.



**Figure 4-6:** Reconstruction efficiency of the proposed method, plotted for different L2-regularizing coefficient values. The trained identity model  $i(\theta_c, \theta_q) = q(\theta_q, c(\theta_c))$  is inspected. Each one of the first 40 Zernike modes is fed to the system with unitary amplitude. The output of the identity model, which represents the theoretical system reconstruction of the requested wavefront in terms of Zernike magnitudes, is reported on the y-axis. Ideally, only the desired mode should be reproduced, with unitary amplitude. The histograms on the right depict the diagonal elements of the grids on the left-hand side. They represent the reconstruction efficiency regarding the desired mode only, that is, not considering the *cross-talking* between modes. The larger the regularizing gain, the more the reconstruction quality is affected, especially for high-order modes.



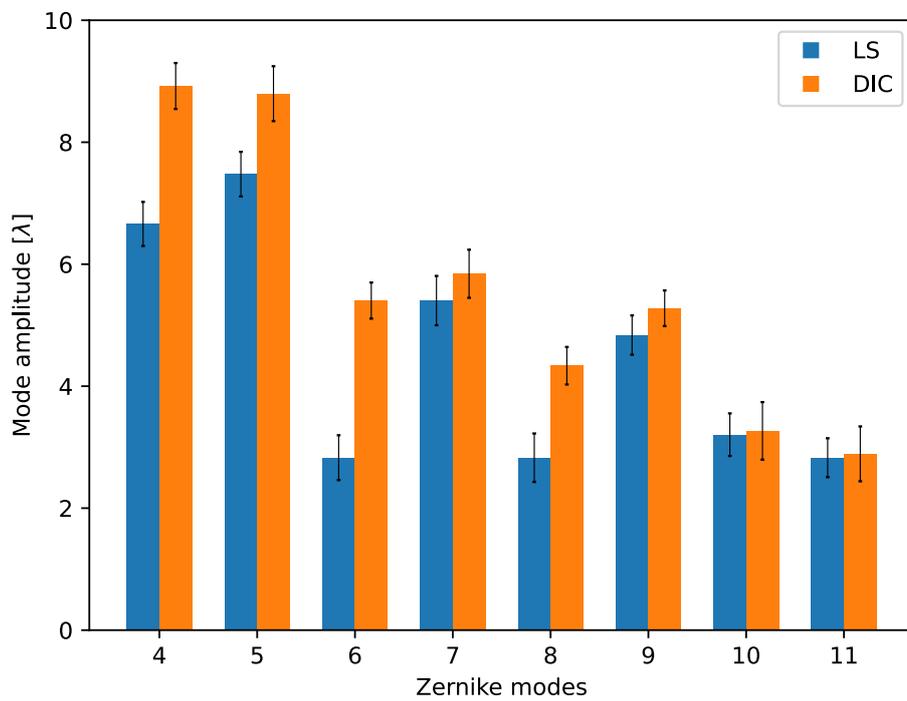
**Figure 4-7:** Risk of saturating actuators, as a function of the applied L2 regularizing gain. On the y-axis, it is plotted the maximum actuator input value required to reconstruct the first 40 Zernike modes with unitary amplitudes. On the x-axis are reported the different penalty values  $\alpha$ . This is equivalent to the maximum value of the control matrix  $S$ , analyzed in 2-11. Increasing the regularizing term reduces the magnitude of the control actions and thus the likelihood of actuator saturation.

As shown in Figure 4-8, the proposed DIC-based method proves to be able to generate larger Zernike amplitudes, maintaining a low rms residual WF error. Since the inversion process does not affect the reconstruction efficiency of the analyzed modes, the open-loop control performances are uniquely dependent on the quality of the plant model. Therefore, given the WF prediction results shown in Figure 4-4, it is not surprising that the proposed method outperforms the linear benchmark in this open-loop reconstruction task.

### Closed-Loop reconstruction

The control of AO systems is in practice usually run in closed-loop, in order to enhance performances and stability [52]. In particular, this is done not only to compensate for nonlinearities but also to improve the residual reconstruction error when the model is significantly regularized. As a consequence, multiple WF measurements need to be gathered, which inevitably slows down the WF reconstruction process. Since the research question of this thesis clearly aims at compensating for DM nonidealities causing little impact on the correction speed, investigating the control trajectory and convergence of the proposed method is extremely important.

Generally speaking, frameworks like DIC and RLS can be used online to iteratively update the model and correct residual WF errors, *de facto* performing closed-loop reconstruction. However, this technique is able to correct for modeling errors only and does not guarantee satisfactory results if the inversion process is regularized, which is usually the case in practice. Therefore, with *closed-loop correction* I refer here to the traditional proportional controller,



**Figure 4-8:** Maximum amplitude of the first 10 Zernike modes that are reconstructed with a wavefront rms error lower than  $0.3 \lambda$ . The first 3 modes (piston, tip, and tilt) are excluded since the used DM is not able to generate them. The proposed method outperforms the linear state-of-the-art on all the modes. Indeed, DIC allows reconstructing larger mode amplitudes maintaining a low WF error.

which can be formulated as follows:

$$\begin{aligned}\mathbf{u}_{k+1} &= Q^\dagger \cdot \mathbf{a}_{r,k+1} \\ &= c(\theta_{c,k+1}, \mathbf{a}_{r,k+1}),\end{aligned}$$

with

$$\mathbf{a}_{r,k+1} = \mathbf{e}_k + \mathbf{a}_{r,k}.$$

In Figure 4-9, an example of astigmatism mode generation is reported, where the proposed DIC method is implemented in closed-loop. As can be seen from the plot at the bottom right of the figure, convergence is immediately reached, and following control iterations do not improve the residual WF error. The proposed method is then also compared against LS and RLS on a batch of 50 wavefronts generated by randomly sampling the first 10 Zernike modes. The averaged results are reported in Figure 4-10 and in Table 4-4. Finally, the closed-loop correction performances relative to the first 6 Zernike modes are presented in Figure 4-11.

Consistently with what has been shown in the previous section, the proposed method out-

Control model	rmse in Open-Loop	rmse in Closed-Loop
LS	0.478 $\lambda$	0.205 $\lambda$
DIC	0.323 $\lambda$	0.204 $\lambda$

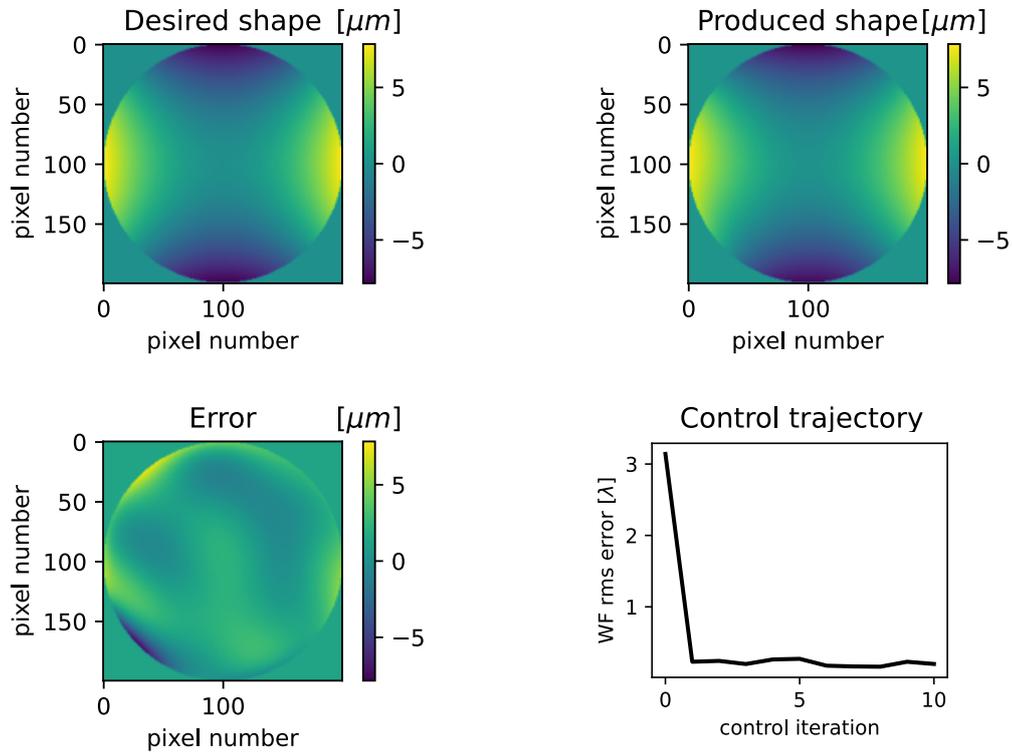
**Table 4-4:** root-mean-square performance error  $\epsilon$  obtained by the proposed method and benchmark on 50 wavefronts, randomly generated sampling the first 10 Zernike modes.

performs the linear benchmark after the first control step. Nevertheless, LS rapidly reaches convergence and in 3 iterations the difference between the control methods is negligible. For those wavefronts where the DIC-based approach does not present any initial advantage, the converge speed is comparable with the state-of-the-art one.

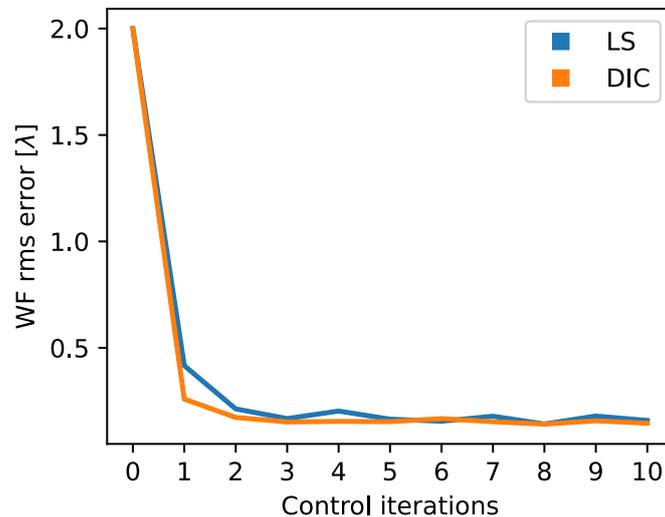
#### 4-4-2 Time-varying conditions

Finally, the WF reconstruction performances of the proposed method are tested in a time-varying environment. In particular, the temperature is increased from the calibration level of  $23^\circ C$  up to  $40^\circ C$ , while the forward and inverse networks are updated online. In this regard, a closed-loop extension of the online framework described in Section 3-1-3 is implemented, which is reported in table 4-5.

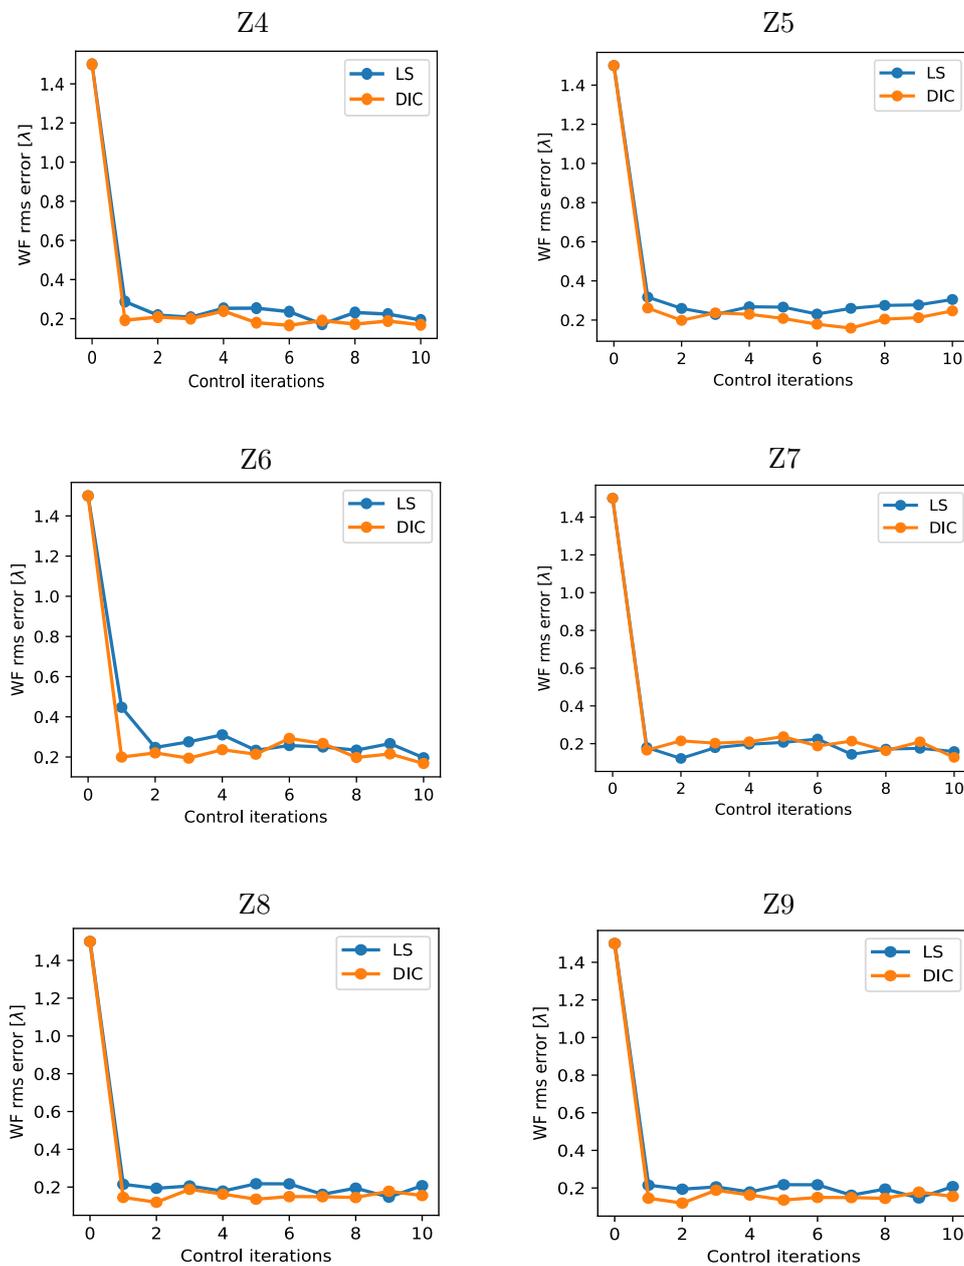
The closed-loop correction of Zernike mode 5 for different temperature levels is reported in Figure 4-12. Consistently with the prediction results shown in Section 4-2, the static model identified through LS becomes unreliable when the working point moves from the calibration one, and as a consequence, its correction performances are also hindered. In particular, the closed-loop reconstruction presents a slower convergence and an offset at steady-state. On the other hand, Recursive Least Squares and the proposed method are both able to maintain the model reliable and indeed they show only minor performance degradation for large temperature changes.



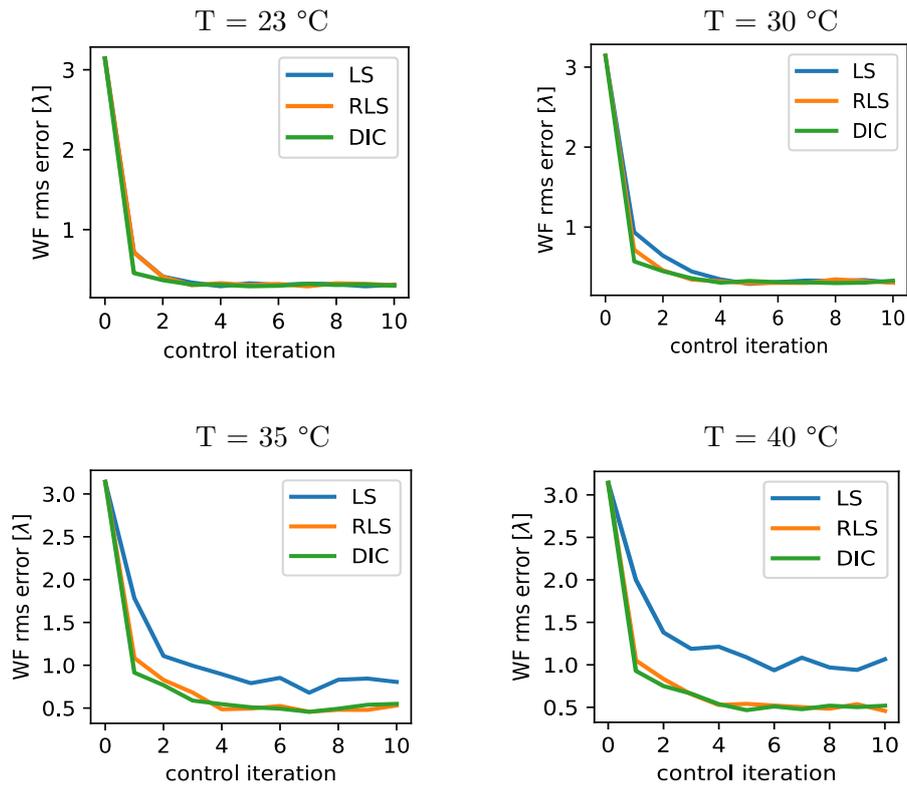
**Figure 4-9:** Closed-Loop reconstruction of astigmatism mode using Direct Inverse Control.



**Figure 4-10:** Closed-Loop reconstruction error of the proposed method, averaged over 50 WF corrections. The reference WFs are randomly sampled from the first 10 Zernike modes, and it is made sure that no actuator saturation occurs. Consistently with the results presented in the Open-Loop section, the proposed method shows lower wavefront rms error during the first control iterations. Nevertheless, the performances of DIC and the benchmarks are comparable once convergence is reached, after 3 – 4 control iterations.



**Figure 4-11:** Closed-loop reconstruction error of the first 6 Zernike modes, excluding piston, tip and tilt. For each Zernike, the largest mode amplitude which does not trigger saturation is taken as WF reference. Each WF reconstruction is run 5 times, and the results are then averaged. The proposed method shows lower residual error during the first control iterations for all the analyzed modes. Once convergence is reached, however, the difference between DIC and the linear benchmarks is negligible.



**Figure 4-12:** Residual wf correction rms error at different temperature levels. The generated WF is Zernike mode 5 with amplitude  $3.14\lambda$ , and each trajectory is the average of 5 experiment runs. In between episodes, random wavefronts are corrected in order to allow RLS and DIC methods to maintain the models updated. These 2 algorithms indeed present only minor performance degradation regardless of the external condition. On the other hand, the static linear model identified through LS presents slower convergence time and steady-state offset at  $35^{\circ}\text{C}$  and  $40^{\circ}\text{C}$ .

---

**Online Closed-Loop DIC Algorithm for AO** Closed-loop extension of the framework reported in Figure 3-4.

---

Step	Description
<b>Require:</b> $q(\theta_{q,0})$	Pretrained plant network
<b>Require:</b> $c(\theta_{c,0})$	Pretrained controller network
<b>Require:</b> $\alpha_q$	ADAM SGD learning rate for the plant network
<b>Require:</b> $\alpha_c$	ADAM SGD learning rate for the controller network
<b>For</b> episode $t = 1, T$ <b>do</b>	
$k \leftarrow 0$	Initialize timestep
$\mathbf{a}_{r,0} \leftarrow$	Obtain new reference wavefront
<b>For</b> control iteration $k = 1, K$ <b>do</b>	
$k \leftarrow k + 1$	
$\mathbf{u}_k \leftarrow c(\theta_c, \mathbf{a}_{r,k})$	Compute control actions through inverse model
$\mathbf{a}_k \leftarrow \text{SH}(\text{DM}(\mathbf{u}_k))$	Send control actions to DM and measure wavefront
$\mathbf{e}_k = \mathbf{a}_{r,0} - \mathbf{a}_k$	Compute performance error
$\mathbf{a}_{r,k+1} = \mathbf{e}_k + \mathbf{a}_{r,k}$	Set new reference wavefront
$\mathbf{R} \leftarrow (\mathbf{u}_k, \mathbf{a}_k)$	Store samples
<b>End for</b>	
$\theta_{q,t+1} \leftarrow \text{ADAM}(\theta_{q,t}, \mathbf{R})$	Update forward model using latest data
$\theta_{c,t+1} \leftarrow \text{ADAM}(\theta_{c,t}, \mathbf{R})$	Update inverse model using latest data
<b>End for</b>	

---

**Table 4-5:** Summary of the Closed-Loop extension of the Online DIC Algorithm.

## 4-5 Discussion

In this chapter, the performances of the proposed method have been compared against state-of-the-art techniques. The wavefront correction process has been studied by analyzing the two main subproblems that compose it: forward model identification and model inversion. Here below, considerations regarding the obtained results are reported and grouped according to the investigated test cases.

### WF prediction

The model identified through the proposed method proves to be more accurate than the linear influence matrix  $Q$  in predicting the behavior of the AO system for large control magnitudes (Figure 4-4). This fact can be explained by the presence of nonlinearities in this specific working range (Figure 2-14), which are not modeled by the LS identification method. Given the current ferment in microscopy research around analyzing deeper tissues, this is an extremely important remark.

Secondly, the proposed method proves to be able to maintain the model reliable regardless of time-varying conditions (Table 4-3). This clearly represents an advantage over the traditional

LS approach. On the other hand, RLS is able to achieve comparable results. The tracking sensitivity of the proposed method can be tuned by changing the learning rate of the network, however, further quantitative analyses in this regard are left for future work.

Finally, it is shown that the stability of the forward model is not hindered by the lack of persistent excitation in the measured data (Figure 4-5). This is an important remark given that the RLS requires specific extensions to maintain the model reliable when this condition arises [24].

### **Ideal reconstruction efficiency**

Regarding the identification of the controller, the first and most important point that has been shown is that the proposed framework is able to invert the forward model with negligible losses. This is an important result since it makes the WF reconstruction process solely dependent on the quality of the forward model, which has been proven to outperform benchmarks in multiple aspects.

In practice, however, the controller is often purposely regularized, in order to reduce the risk of saturation and make the correction process more stable. As it is shown in Fig.4-1, augmenting the loss function of the controller with a penalization term on the weights of the network, allows to control the trade-off between reconstruction ability and saturation risk. Despite the process being less trivial than truncating modes through Singular-Value-Decomposition, L2 regularization is a well-documented method, that can easily be implemented through available libraries. To the writer's knowledge, model regularization is often not discussed when proposing new Machine Learning (ML) approaches for the control of AO, and for this reason, I believe these findings are particularly worth mentioning.

### **WF correction**

Consistently with the prediction performances mentioned before, the proposed method proves to outperform the linear state-of-the-art on the open-loop reconstruction of large-phase wavefronts (Fig.4-8). Direct Inverse Control presents an advantage only during the first 2 control iterations, after which convergence is usually reached and the performance errors are comparable. In this regard, it is worth noting that the proposed method often reaches convergence after just 1 step and indeed does not improve the residual WF error through closed-loop iterations. Given the research question of this thesis is to address system nonlinearities without affecting control speed, this is an extremely important result.

Furthermore, the proposed method shows solid tracking abilities also with regards to the controller network, which indeed guarantees stable WF correction regardless of the temperature the system is working at (Fig.4-12). These considerations are also valid for the controller based on RLS, while on the other hand, they do not hold for the static LS framework, which presents slower convergence speed and a steady-state offset.

---

## Chapter 5

---

# Conclusion

This thesis aimed to investigate the problem of identifying and controlling an MMDM-based Adaptive Optics (AO) system which exhibits nonlinearities and time-varying response. The primary objective of this study was to develop an approach that compensates for these non-idealities while affecting correction speed less than linear state-of-the-art approaches. The proposed approach relies on the Direct Inverse Control framework, which exploits neural networks to model the AO system. Specifically, the controller is computed by inverting the plant network following the *specialized learning* Direct Inverse Control (DIC) architecture.

The results showed that the proposed approach provides better modeling accuracy than benchmarks in the working ranges where nonlinearities are present, enabling faster control convergence than the state-of-the-art when generating large-phase wavefronts. Moreover, the DIC-based method demonstrated better stability and similar tracking abilities to Recursive Least Squares. The analysis of model regularization yielded positive results. However, a more comprehensive understanding of the method's capabilities could be obtained by further testing and comparisons with Singular-Value-Decomposition.

In this regard, other future development possibilities of this approach include:

- Using a Convolutional Neural Network (CNN) for the controller network.
- Benchmarking the proposed approach against DIC based on *generalized learning*.
- Running extensive tests to study the long-term tracking abilities of the algorithm.
- Automating the tuning of networks' hyperparameters.

In conclusion, this study provides evidence that the proposed method based on the DIC framework and neural network modeling can effectively compensate for the non-idealities of MMDM-based AO systems. The findings suggest that this approach has significant potential for further development and application in the field of Adaptive Optics.



---

# Nomenclature

Symbol	Quantity	Unit
$\alpha$	L2 regularizing coefficient	-
$a$	Individual Zernike mode's coefficient	-
$\mathbf{a}$	Wavefront's Zernike modes coefficients	-
$\bar{\mathbf{a}}$	Predicted wavefront's Zernike modes coefficients	-
$\mathbf{A}$	Series of wavefronts' Zernike modes coefficients	-
$C$	Number of mirror actuators	-
$c$	Controller neural network	-
$D$	Number of Zernike modes	-
$\delta$	Number of truncated modes during SVD	-
$e$	Performance error	-
$\bar{e}$	Predicted performance error	-
$\bar{e}$	Prediction error	-
$\epsilon$	Wavefront prediction error	-
$\eta$	Measurement noise	-
$f$	Focal length	m
$G$	Input vector of Deformable Mirror as Kronecker matrix	-
$I$	Pixel intensity	-
$i$	Overall neural network (controller+plant)	-
$\gamma$	RLS forgetting factor	-
$L$	RLS update gain matrix	-
$\lambda$	Light wavelength	-
$M$	Number of wavefront slopes	-
$N$	Number of calibration samples	-
$n$	Refractive index	-
$\phi$	Wavefront	rad
$P$	RLS uncertainty matrix	-

Symbol	Quantity	Unit
$Q$	Influence matrix of a deformable mirror	-
$\mathbf{q}$	Vectorized influence matrix of a deformable mirror	-
$q$	Neural network function	-
$S$	Control matrix of a Deformable Mirror	-
$s$	Wavefront slopes	rad/m
$\Sigma$	Singular-Value-Decomposition (SVD) diagonal matrix	-
$\sigma$	Activation function of neural network	-
$T$	Reconstruction efficiency matrix	-
$\theta$	Set of neural network weights	-
$u$	Individual actuator's command	-
$\mathbf{u}$	Actuators' commands vector	-
$\mathbf{U}$	Series of actuators' commands vector	-
$U$	SVD left matrix	-
$v$	Light speed in a homogeneous medium	m/s
$V$	SVD right matrix	-
$x_c$	Centroid position on the $x$ -direction	m
$y_c$	Centroid position on the $y$ -direction	m
$x_r$	Reference centroid position on the $x$ -direction	m
$y_r$	Reference centroid position on the $y$ -direction	m

**DCSC** Delft Center for Systems and Control

**WF** wavefront

**DM** deformable mirror

**AO** Adaptive Optics

**OPL** Optical Path Length

**SH** Shack-Hartmann

**SVD** Singular-Value-Decomposition

**SNR** signal-to-noise ratio

**SDK** Software Development Kit

**CNN** Convolutional Neural Network

**LS** Least Squares

**RLS** Recursive Least Squares

**NLS** Nonlinear Least Squares

**ANN** Artificial Neural Networks

**WFSless** Wavefront Sensorless

**RL** Reinforcement Learning

**ML** Machine Learning

**DIC** Direct Inverse Control

**ADAM** Adaptive moment estimation

<b>SGD</b>	Stochastic Gradient Descent
<b>MMDMs</b>	micromachined membrane deformable mirrors
<b>rms</b>	root-mean-square



---

# Bibliography

- [1] M. Verahegen, P. Pozzi, O. Soloviev, G. Vdovin, and D. Wilding. *Lecture notes on control for high resolution imaging*. Delft Center for Systems and Control, 2017.
- [2] Gleb Vdovin, Oleg Soloviev, M Loktev, and V. Patlan. *OKO Guide to Adaptive Optics*. Flexible Optical BV, 2013.
- [3] Robert J Noll. Zernike polynomials and atmospheric turbulence. *JOsA*, 66(3):207–211, 1976.
- [4] Michael I Jordan and David E Rumelhart. Forward models: Supervised learning with a distal teacher. *Cognitive science*, 16(3):307–354, 1992.
- [5] Siegfried Weisenburger and Vahid Sandoghdar. Light microscopy: an ongoing contemporary revolution. *Contemporary Physics*, 56(2):123–143, 2015.
- [6] M. J. Booth. Adaptive optical microscopy: the ongoing quest for a perfect image. *Light: Science & Applications*, 2014.
- [7] Joel A Kubby. *Adaptive optics for biological imaging*. CRC press, 2013.
- [8] Martin Booth. A universal framework for adaptive optics microscopy. *Young researchers biophotonics summer meeting*, 2022.
- [9] Thomas Bifano. Mems deformable mirrors. *Nature photonics*, 5(1):21–23, 2011.
- [10] Guang S He. Optical phase conjugation: principles, techniques, and applications. *Progress in Quantum Electronics*, 26(3):131–191, 2002.
- [11] Ben C Platt and Roland Shack. History and principles of shack-hartmann wavefront sensing, 2001.
- [12] Daniel R Neal, James Copland, and David A Neal. Shack-hartmann wavefront sensor precision and accuracy. In *Advanced Characterization Techniques for Optical, Semiconductor, and Data Storage Components*, volume 4779, pages 148–160. SPIE, 2002.

- [13] JY Wang and D Es Silva. Wave-front interpretation with zernike polynomials. *Applied optics*, 19(9):1510–1518, 1980.
- [14] G Vdovin, M Loktev, and A Simonov. Low-cost deformable mirrors: technologies and goals. In *Advanced Wavefront Control: Methods, Devices, and Applications III*, volume 5894, pages 103–112. SPIE, 2005.
- [15] Ronald P Grosso and Martin Yellin. The membrane mirror as an adaptive optical element. *JOSA*, 67(3):399–406, 1977.
- [16] Gleb Vdovin. *Micromachined membrane deformable mirrors*. Marcel Dekker Inc, New York, 1998.
- [17] Rachel E Morgan, Ewan S Douglas, Gregory W Allan, Paul Bierden, Supriya Chakrabarti, Timothy Cook, Mark Egan, Gabor Furesz, Jennifer N Gubner, Tyler D Groff, et al. Mems deformable mirrors for space-based high-contrast imaging. *Micromachines*, 10(6):366, 2019.
- [18] Lei Huang, Xingkun Ma, Qi Bian, Tenghao Li, Chenlu Zhou, and Mali Gong. High-precision system identification method for a deformable mirror in wavefront control. *Applied Optics*, 54(14):4313–4317, 2015.
- [19] Peiliang Xu. Truncated svd methods for discrete linear ill-posed problems. *Geophysical Journal International*, 135(2):505–514, 1998.
- [20] Martin Booth, Tony Wilson, Hong-Bo Sun, Taisuke Ota, and Satoshi Kawata. Methods for the characterization of deformable membrane mirrors. *Journal of the Optical Society of America*, 2005.
- [21] Flexible Optical BV. Datasheet of 96-channel deformable mirror with embedded electronics. [http://www.okotech.com/images/pdfs/mdm96\\_25.4\\_datasheet.pdf](http://www.okotech.com/images/pdfs/mdm96_25.4_datasheet.pdf).
- [22] Ali H Sayed. *Adaptive filters*. John Wiley & Sons, 2011.
- [23] Aleksandar Haber and Thomas Bifano. Dual-update data-driven control of deformable mirrors using walsh basis functions. *JOSA A*, 39(3):459–469, 2022.
- [24] Ardalan Vahidi, Anna Stefanopoulou, and Hwei Peng. Recursive least squares with forgetting for online estimation of vehicle mass and road grade: theory and experiments. *Vehicle System Dynamics*, 43(1):31–55, 2005.
- [25] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.
- [26] Sheena Angra and Sachin Ahuja. Machine learning and its applications: A review. In *2017 International Conference on Big Data Analytics and Computational Intelligence (ICBDAC)*, pages 57–60. IEEE, 2017.
- [27] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning internal representations by error propagation. Technical report, California Univ San Diego La Jolla Inst for Cognitive Science, 1985.
- [28] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.

- 
- [29] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [30] Bernard Widrow and Samuel D. Stearns. *Adaptive Signal Processing*. Prentice-Hall, 1985.
- [31] Demetri Psaltis, Athanasios Sideris, and Alan A Yamamura. A multilayered neural network controller. *IEEE control systems magazine*, 8(2):17–21, 1988.
- [32] Valdemir Carrara and Atair Rios Neto. A neural network satellite attitude controller with error based reference trajectory. In *XIV International Symposium on Space Flight Dynamics, Brazil*. Citeseer, 1999.
- [33] Herwin Suprijono and Benyamin Kusumoputro. Direct inverse control based on neural network for unmanned small helicopter attitude and altitude control. *Journal of Telecommunication, Electronic and Computer Engineering (JTEC)*, 9(2-2):99–102, 2017.
- [34] GCM De Abreu, Rafael Luís Teixeira, and José F Ribeiro. A neural network-based direct inverse control for active control of vibrations of mechanical systems. In *Proceedings. Vol. 1. Sixth Brazilian Symposium on Neural Networks*, pages 107–112. IEEE, 2000.
- [35] William Camilo Ariza-Zambrano and Alberto Luiz Serpa. Direct inverse control for active vibration suppression using artificial neural networks. *Journal of Vibration and Control*, 27(1-2):31–42, 2021.
- [36] DB Anuradha, G Prabhaker Reddy, and JSN Murthy. Direct inverse neural network control of a continuous stirred tank reactor (cstr). In *Proceedings of the International MultiConference of Engineers and Computer Scientists*, volume 3, 2009.
- [37] Frederick Walker Fairman. *Linear control theory: the state space approach*. John Wiley & Sons, 1998.
- [38] Youming Guo, Libo Zhong, Lei Min, Jiaying Wang, Yu Wu, Kele Chen, Kai Wei, Changhui Rao, et al. Adaptive optics based on machine learning: a review. *Opto-Electronic Advances*, 5(7):200082–1, 2022.
- [39] Peter L Wizinowich, Michael Lloyd-Hart, Brian A McLeod, D'nardo Colucci, Richard G Dekany, David M Wittman, James Roger P Angel, Donald W McCarthy Jr, William G Hulburd, and David G Sandler. Neural network adaptive optics for the multiple-mirror telescope. In *Active and adaptive optical systems*, volume 1542, pages 148–158. SPIE, 1991.
- [40] Robin Swanson, Masen Lamb, Carlos Correia, Suresh Sivanandam, and Kiriakos Kukulakos. Wavefront reconstruction and prediction with convolutional neural networks. *SPIE Adaptive Optics Systems*, 2018.
- [41] Dennis A Montera, Byron M Welsh, Michael C Roggemann, and Dennis W Ruck. Prediction of wave-front sensor slope measurements with artificial neural networks. *Applied optics*, 36(3):675–681, 1997.

- [42] Hu Ke, Bing Xu, Zhenxing Xu, Lianghua Wen, Ping Yang, Shuai Wang, and Lizhi Dong. Self-learning control for wavefront sensorless adaptive optics system through deep reinforcement learning. *Optik*, 178:785–793, 2019.
- [43] Rico Landman, Sebastiaan Y Haffert, Vikram M Radhakrishnan, and Christoph U Keller. Self-optimizing adaptive optics control with reinforcement learning. In *Adaptive Optics Systems VII*, volume 11448, pages 842–856. SPIE, 2020.
- [44] Robin Swanson, Masen Lamb, Carlos M Correia, Suresh Sivanandam, and Kiriakos Kukulakos. Closed loop predictive control of adaptive optics systems with convolutional neural networks. *Monthly Notices of the Royal Astronomical Society*, 503(2):2944–2954, 2021.
- [45] B Pou, Florian Ferreira, Eduardo Quinones, Damien Gratadour, and Mario Martin. Adaptive optics control with multi-agent model-free reinforcement learning. *Optics express*, 30(2):2991–3015, 2022.
- [46] Aleksandar Haber. Steady-state control and machine learning of large-scale deformable mirror models. *arXiv preprint arXiv:1911.07456*, 2019.
- [47] Zhenxing Xu, Ping Yang, Ke Hu, Bing Xu, and Heping Li. Deep learning control model for adaptive optics systems. *Applied Optics*, 2019.
- [48] Alfredo Dubra. Wavefront sensor and wavefront corrector matching in adaptive optics. *Optics express*, 15(6):2762–2769, 2007.
- [49] Mario E Salgado, Graham C Goodwin, and Richard H Middleton. Modified least squares algorithm incorporating exponential resetting and forgetting. *International Journal of Control*, 47(2):477–491, 1988.
- [50] Tore Hägglund. Recursive estimation of slowly time-varying parameters. *IFAC Proceedings Volumes*, 18(5):1137–1142, 1985.
- [51] N Rao Sripada and D Grant Fisher. Improved least squares identification. *International Journal of Control*, 46(6):1889–1913, 1987.
- [52] Karen Hampson and Martin Booth. Calibration and closed-loop control of deformable mirrors using direct sensing, 2020.