



HAL
open science

Reconfigurable design of a thermo-optically addressed liquid-crystal phase modulator by a neural network

Stéphane Barland, Loic Ramousse, Gilles Chériaux, Vincent Femy, Cyrille Claudet, Aurélie Jullien

► **To cite this version:**

Stéphane Barland, Loic Ramousse, Gilles Chériaux, Vincent Femy, Cyrille Claudet, et al.. Reconfigurable design of a thermo-optically addressed liquid-crystal phase modulator by a neural network. *Optics Express*, 2023, 31 (8), pp.12597. 10.1364/OE.483141 . hal-04055428

HAL Id: hal-04055428

<https://hal.science/hal-04055428>

Submitted on 2 Apr 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Reconfigurable design of a thermo-optically addressed liquid-crystal phase modulator by a neural network

STÉPHANE BARLAND,¹  LOIC RAMOUSSE,^{1,2}
GILLES CHÉRIAUX,¹ VINCENT FEMY,² CYRILLE CLAUDET,¹
AND AURÉLIE JULLIEN^{1,*} 

¹Université Côte d'Azur, CNRS, Institut de Physique de Nice (INPHYNI), UMR 7010, 17 rue Julien Lauprêtre, 06200 Nice, France

²Fastlite, 165 Route des Cistes, 06600 Antibes, France

*aurelie.jullien@inphyni.cnrs.fr

Abstract: We present a machine learning approach to program the light phase modulation function of an innovative thermo-optically addressed, liquid-crystal based, spatial light modulator (TOA-SLM). The designed neural network is trained with a little amount of experimental data and is enabled to efficiently generate prescribed low-order spatial phase distortions. These results demonstrate the potential of neural network-driven TOA-SLM technology for ultrabroadband and large aperture phase modulation, from adaptive optics to ultrafast pulse shaping.

© 2023 Optica Publishing Group under the terms of the [Optica Open Access Publishing Agreement](#)

1. Introduction

Spatial light modulators (SLMs) are two-dimensional objects, enabling to modulate the intensity, phase or polarization of an incident light beam, at any point of the SLM surface, through a local change of the optical path. Their applications in photonics are too numerous to be enumerated, enough is to say that SLMs are now widely employed in complex optical systems [1]. In particular, the ability to manipulate the light phase is required for applications as varied as holography, optical sensing, femtosecond pulse shaping, singular beams generation and adaptive optics [2–6]. The latter item, adaptive optics, refers to the addition of an active or passive component in an optical system to correct spatial phase aberrations so as to optimize the point spread function for the targeted application, which can belong to many fields, from microscopy to astronomy, including the spatial control of high-power lasers for high-field physics or industrial laser processing [7–11]. A plethora of technologies provide such dynamic components for phase correction, among them mechanically or thermally deformable mirrors, magneto-optic devices or acoustic-optic Bragg cells, and non-mechanical components which exploit the electro-optic anisotropy of liquid crystals (LCs).

Recently, a new SLM concept has been proposed [12], relying on a local thermal modification of a thick LC layer, that is optically-induced through the absorption of a control beam. This innovative thermo-optically addressed LC-SLM, coined TOA-SLM, has shown dynamic phase control capabilities over multi-octave light spectrum, as a promising candidate for spatial or temporal manipulation of ultrafast pulses. In addition to being ultra-broadband and programmable, such a device is low-cost, large-aperture (>50 mm) and un-segmented with a high number of control points (estimated density of 50 control points per cm²).

Physical modeling of the TOA-SLM by finite element analysis provides insight about the relative role of distinct physical mechanisms in the TOA-SLM operation. However, making a model quantitatively predictive is in general difficult, due to poorly known boundary conditions and uncertainties on the thermal coefficients of the different components of the structure. Beyond this difficulty though, physical modeling is at best predictive in terms of the final state of a system

assuming some known input. On the contrary, a practical, reliable and reconfigurable control of the modulator operation requires to guess which control beam profile must be applied in order to obtain a prescribed phase profile, which is known as an *inverse problem*. Even assuming an extremely simple physical model of the phase modulator in terms of the heat equation, solving the inverse problem is a long standing question [13], which can be answered analytically only in very specific conditions (see *eg* [14]). Even in this case, this would still not take into account the nonlinearity of the LC thermal response and other, harder to model, details of the device. To circumvent these difficulties, machine learning can be a suitable approach. We note that the task to be performed (control the TOA-SLM as a nonlocal and nonlinear dynamical system) is in principle a *control* task, for which reinforcement learning for instance could provide an interesting approach. However, since we are not interested in dynamical control but only in the stationary state reached by the TOA-SLM under constant operating condition, the task becomes more of a *design* task for an optical component with prescribed phase retardation properties.

Machine learning algorithms are currently being developed to address many topics in photonics, in order to add new functionalities, solve inverse problems and enhance performances, with applications in ultrafast photonics, photonic sensing or design of photonic structures, to name a few [15–21]. While the first applications of machine learning to adaptive optics can be dated back to the early 90's [22,23], these tools might now be used in every stage of an adaptive optics system, i.e. intelligent wavefront sensing, wavefront prediction and post-processing [24].

In this paper, we propose the construction and training of a neural network-based statistical model suitable to provide configurable design of a prototype TOA-SLM [25]. The network is trained on available experimental data and tested by predicting parameters to generate typical low-order spatial phase aberrations. The paper is organized as follows: section 2 presents the experimental methods, section 3 describes the neural network design and training, and the last section gathers the experimental results, assessing the capacities of our device for low-order aberrations correction.

2. Experimental methods

2.1. Principle and prototype description

The device is an ultra-broadband reflective spatial light modulator, whose design is issued from previous investigations on the use of thick cells of thermotropic nematic LCs [26–28]. Unlike traditional SLMs, no electrode is required, as the modulation of the optical index of the LC layer is introduced by a local modification of its temperature, itself induced by the absorption of a control beam. The absorption process generates heat and establishes a localized thermal gradient in the LC layer by diffusion. The LC being thermotropic and birefringent, both the ordinary and extraordinary optical refractive indices are locally modulated, with a non linear dependence on the temperature [26]. As a result, the device is ultra-broadband (validated over 500nm–2600nm spectral bandwidth) and unsegmented. The phase modulation range can reach three wave-numbers even at the longest wavelengths. The active area can reach 50 mm, with a density of control points around 50 pts/cm². Spatial resolution is limited to a few hundreds of μm by thermal diffusion, which also limits the commutation time (between 30-60s depending of the LC layer thickness) and refresh time (about 100ms) [12].

All the tests presented here are realized with a prototype of TOA-SLM with 50 mm lateral aperture, depicted in Fig. 1(a). The multi-layer structure is composed of: (i) a BK7 window (2 mm), coated on one side by a 100 nm gold layer and by an AR-coating at 465 nm on the other side, (ii) a SiO₂ window (2 mm) with a broadband AR-coated on one side and uncoated on the other side. A thin film of polyvinyl alcohol (PVA) is spin-coated on both the gold layer and the uncoated SiO₂ surface, and then rubbed to anchor the molecules in a plane parallel to the substrate. Finally, a layer of 80 μm of the nematic mixture E7 (Merck) is inserted between the substrates, in contact with the PVA layers. The gold layer achieves simultaneously partial

absorption ($\approx 60\%$ absorption) of a control beam at 465 nm and reflection of the probe light. As shown in Fig. 1(b), thermo-optical addressing is applied from the back of the TOA-SLM, with a 465 nm fiber-coupled laser diode (2.4 W) propagating through the BK7 substrate. Arbitrary phase modulation is then achieved by spatially shaping the output of the 465 nm laser with digital light processing (DLP6500, Texas Instruments). The DLP diffracts two modes: a mode that reproduces the loaded image and its complementary image. The latter, the complementary mode, is relay-imaged from the DLP plane onto the metallic layer by a two-lens telescope.

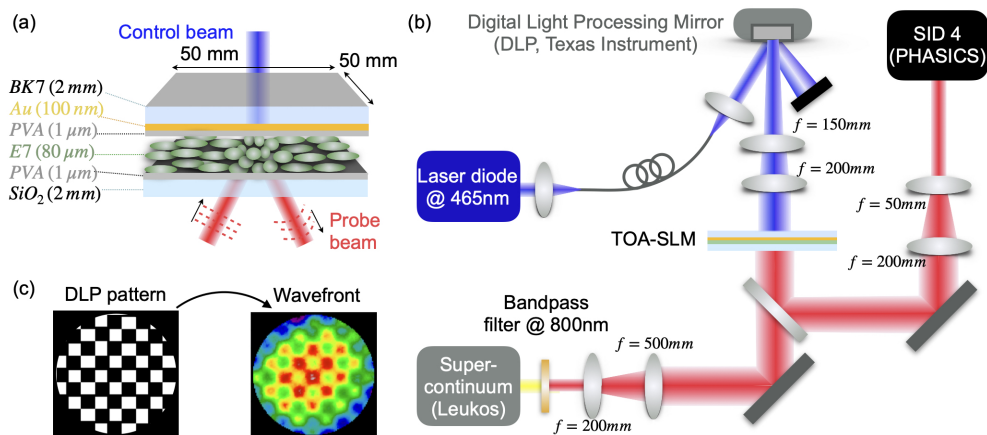


Fig. 1. (a) Design of the TOA-SLM and cross-section of the multi-layer structure (b) Schematic of the adaptive optical setup characterization. (c) Example of how a DLP pattern (black parts correspond to a maximum of blue light) enables spatial phase modulation.

2.2. Phase detection

The induced spatial phase modulation is probed by a super-continuum laser (ElectroVisir, Leukos), with a bandpass filter selecting a spectral bandwidth of a few tens of nm around 800 nm, but could be replaced by a CW or femtosecond laser source. The beam size is fixed to 10 mm and the laser polarization is chosen to match the LC extraordinary axis. A beam-splitter enables pick-up of the reflected beam and the TOA-SLM plane is then relay-imaged with ≈ 3.1 reduction factor onto a wavefront sensor based on quadri-wave lateral shearing interferometry (SID4-GE, Phasics) [29]. We underline that an homogeneous phase offset (piston) cannot be measured with the wavefront sensor. The minimum phase shift that can be detected by the sensor is 0.02λ (corresponding to an optical path difference of 16 nm). The wavefront acquired without any blue light applied to the TOA-SLM (phase shift throughout the pupil lower than 0.05λ) is considered as the reference wavefront and subtracted to forthcoming measurements. For this study, the maximum stroke of the TOA-SLM is $\approx 0.8\lambda$, i.e. an optical phase difference (OPD) of 640 nm, which corresponds to a LC refractive index change of 0.005 and a temperature change of 3-4 K, limited by the power of the available laser diode (here less than 1 W/cm^2). For each acquisition, the wavefront is recorded when the thermal steady state is reached, ie when the wavefront no longer evolves within the detection limit of the sensor (approximately 1 minute). Between each illumination, the TOA-SLM returns to the equilibrium (lab. temperature). Figure 1(c) illustrates how a pattern loaded onto the DLP modulates the introduced spatial phase. As the DLP complementary mode is imaged onto the TOA-SLM, a local minimum in the pattern (in black) corresponds to a local maximum of control blue light, which generates heat. The extraordinary index of the LC material then decreases, which translates as a positive phase-shift detected by the wavefront sensor. As can be seen in Fig. 1(c), the wavefront pattern in the steady-state

regime reproduces the applied chess-like amplitude modulation, convolved with a large-scale background modulation, originating from heat diffusion in the multi-layer structure. The phase inhomogeneity observed between the center and the edges of the image may result from heat diffusion at the edges and from the residual inhomogeneity of the writing beam incident on the TOA-SLM, due to the imperfect reflectivity of the DLP and collection of diffraction orders.

3. Neural network model and training

3.1. Model design

The proposed model is expected to be able to suggest which parametric condition (*ie* which DLP profile) must be applied in order to achieve a desired phase profile correction through the thermally induced refractive index variation. Thus, the network's task is to solve the inverse problem of the experiment by outputting (starting from a prescribed two-dimensional wavefront) the numerical bidimensional pattern which should be uploaded to the DLP device. This pattern will then be used to spatially modulate the (blue) control beam profile which in turn will locally heat the liquid crystal modulator to obtain experimentally the prescribed phase modulation. From empirical observations, it appears that the phase modulation imposed by the patterning beam depends both on the locally applied power but also on the total applied power in the patterning beam. Thus, a locally connected network would imply connections over a very large neighborhood, leading to a huge number of coefficients. We choose instead to build a convolutional network [25], considering that at least some approximate spatial invariance exist (rotations in particular). We use the Keras library [30] (a deep learning application programming interface written in Python, running on top of the machine learning platform TensorFlow) to build a relatively small network inspired by the U-Net architecture [31] and adapted to the specifics of our data as detailed in the following. At the input of the network, a layer computes the gradient of the phase image in order to get rid of the arbitrary homogeneous offset implied by the lack of piston data in the phase images (see section 2). The following stages of the network consist of a contracting stack of convolutional and pooling layers followed by an expanding stack of transposed-convolution [32] and convolutional layers. The contracting and the expanding parts of the network are connected at two spatial scales by concatenating the output of a convolution layer of the contracting part with the output of a convolutional layer with the same size in the expanding part (see Fig. 2). All convolution layers use a Rectifier as nonlinear activation function except the last one which uses a sigmoid function since the signal sent to the DLP must be between 0 and 1. The kernel size at each layer is 3 and the number of kernels grows towards the center of the network before decreasing again so as to obtain a simple bidimensional image between 0 and 1 at the output of the network. Only a very moderate amount of hyperparameter tuning [33] has been performed since we aim at a proof of operation and not pushing the network performance to the limits. Thus, the number of kernels (much smaller than the original U-net [31]) has been chosen so as to minimize the total number of trainable parameters ($2.4 * 10^5$) without appreciable degradation of the error during training. Two dropout layers (with 20% drop probability) are used in the contracting part.

3.2. Training

The experimental data available for training consists of 138 wavefront images together with the DLP pattern which was used to produce the phase images. The amount of training data is limited by the time needed to reach thermal equilibrium for each acquisition. The DLP patterns consist of periodic square and line patterns (with different wavelengths and different levels of gray) and of randomly generated patterns of the superposition of three gaussian bumps with different maximum value and standard deviation. These two sets of images are randomly mixed and a subset of 108 couples is kept for training while the remaining 30 will be used for model

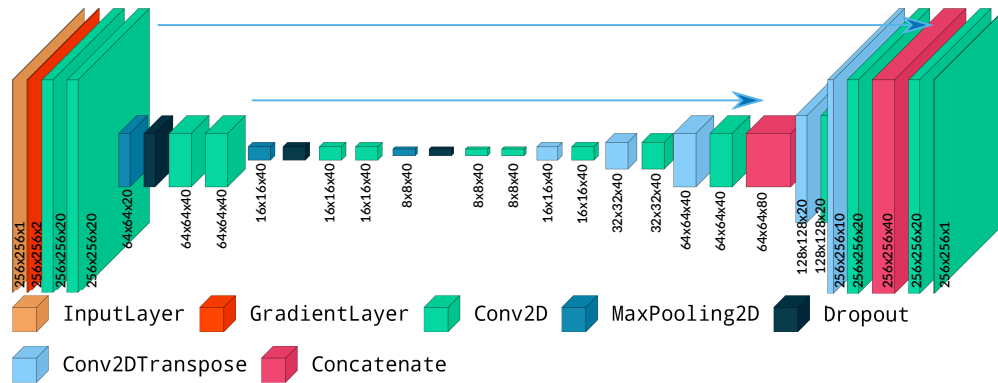


Fig. 2. Network architecture. A contracting-expanding stack of convolutional layers, both sides being connected via concatenation at two different scales. The output shape is shown after each convolutional layer, which all use 0-padding to preserve image size. Visualization done with [34].

validation. Before any other processing, the phase images are downsampled to 256×256 pixels, which is sufficient to accurately sample the spatial phase variations. The DLP patterns are cropped and shifted so that the central point in the phase images correspond to the central spot in the DLP pattern. The patterns also undergo a 45° rotation to compensate for the positioning of the DLP device with respect to the wavefront sensor in the experiment. The DLP patterns are then downsampled to 256×256 pixels as the phase images and finally a magnification factor matching the experimental magnification is applied to DLP patterns.

Since the number of training images is very small (due to the rather time-consuming acquisition process), data augmentation is required. Although we know that the physics is not perfectly symmetric with respect to rotations (because of tiny residual anisotropies during fabrication and imperfection of the relay imaging system) we apply during training random arbitrary rotations around the geometric center of the images to both the DLP patterns and their corresponding phase images. As noted above, the system is clearly not invariant by translation since boundary effects are expected to be strong due to heat flow at the edges, therefore we do not apply synthetic translations during training. To augment the variability of the training data, we randomly apply small "swirl transformations" (a local rotation whose angle exponentially decreases with the distance from the rotation center) to wavefronts and their corresponding DLP patterns. Thanks to the data augmentation, 1024 examples are shown to the network at each epoch of the training (and a new set of 1024 is generated at each epoch).

We train the network by minimizing a weighted loss function based on mean squared difference between the actual DLP pattern and the network-generated pattern, tailored to the physics of the system and to the data augmentation we use. Since we apply rotations which imply unknown values at the corners of some images, the mean squared error is weighted by a supergaussian of order 4 and standard deviation $\sigma = 0.8$ in units of image size. As noted in 2, the blue beam reaching the phase modulator is actually the complementary image of the numerical DLP profile. Therefore, low values in the DLP control matrix translate into high intensity control beam reaching the liquid crystal device. Since the refractive index variation is expected to depend faster than linearly on the control beam intensity, we choose to confer a larger weight to the regions of large control beam intensity. We achieve this by dividing locally the mean squared error value by $\max(0.1, DLP(x, y))$ so as to emphasize the regions where the DLP matrix is near 0. Two examples of training data are shown on the top row of Fig. 3. On both images, we show the DLP pattern in scales of blue and, on top of it, the resulting wavefront in scales of green. The

annotated circles outline the supergaussian which is used to weight the quadratic error between prediction and truth during training. On the left picture, one observes that the three distant bumps in the DLP matrix (which result in three low intensity bumps over a homogeneous background in the patterning beam) result in a rather smooth gaussian-like wavefront, the rightmost bump having barely any effect on the phase. On the right image, the diffusion process at play is also clearly apparent but in this case the phase modulation caused by the DLP pattern is much clearer, especially in the central area.

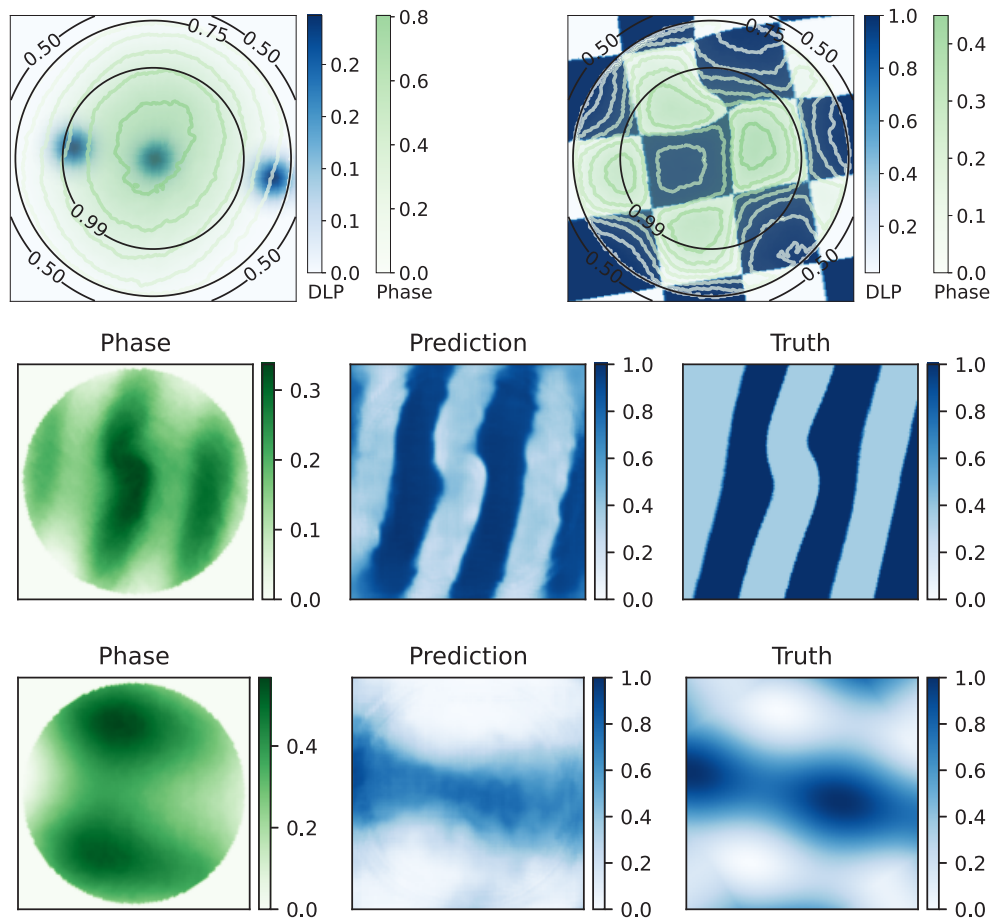


Fig. 3. Training data and training results. Top row: two examples of DLP patterns and the resulting phase profile obtained in the experiment. The annotated circles outline the weighting used for the loss function during training. Middle row: example of how the model reproduces training data (i.e. data set seen during the training): from the left phase image, the model proposes the middle DLP image. The actual DLP image is shown on the right, with the same color scale. Bottom row: example of how the model predicts validation data (i.e. data set not seen during the training).

After training (during 160 epochs by batches of 8), the model is able to reproduce very well the data presented in the training set (middle row of Fig. 3). The weighted RMS error is down to 0.10 on the validation set and correspondingly the model predicts also very well the data used for validation, *ie* data which was never shown to the model during training (bottom row of Fig. 3). Attempts of training with a smaller number of kernels in order to reduce potential overfitting led to a worsening of the RMS error on the validation set, *ie* poorer generalization. On a consumer

grade GPU card (GeForce GTX 1060), the network training takes about one hour. On a desktop CPU, producing a DLP image from a phase profile takes 130 ms, including the computation time dedicated to scaling the predicted image and saving it to disk in a format amenable for upload to the DLP. The actual inference time per image is about 85 ms.

We note that in principle there might be more than one way to achieve a specified phase profile, so that the problem we try to solve with the neural network is *ill-posed*. This may make it impossible for the training process to converge since the network may be exposed to two very different DLP profiles leading to a single (or very similar) phase profile(s). Trying to minimize an error in this case would lead to an "average" solution of the two true solutions, which is not a solution at all. This difficulty, if it arises, can be overcome by a clever approach in form of *tandem networks* [20,35]. Here, our data set is so small that we did not have to resort to that method and the training converged easily because we had no such conflicting solutions in our training set.

4. Results

Section 3 has shown the model ability to predict, after training, with reasonable accuracy the control beam profile which was used to obtain some wavefront profiles. We now assess the performance of the whole phase control system (*ie* the experimental apparatus driven with guidance of the model) in achieving some prescribed wavefronts, totally outside of the training and validation sets.

To do so, quantification of the phase data is required. The first metric to quantify a wavefront is the peak-to-valley value (PtV), that is the maximum phase shift present throughout the pupil. If the phase is written $\varphi(x, y)$, then the PtV value is simply expressed as:

$$PtV = \varphi_{max}(x, y) - \varphi_{min}(x, y) \quad (1)$$

To go further, a usual method consists in decomposing the wavefront as a sum of Zernike polynomials. The Zernike polynomials form an orthogonal basis over the unit circular aperture. They are commonly used to describe the low-spatial-frequency wavefront distortions, as the first polynomials correspond to the principal aberrations occurring in optical systems [36]. The wavefront is then written as :

$$\varphi(x, y) = \sum_{(i,j)=1}^N \alpha_i^j Z_i^j \quad (2)$$

with α_i^j the coefficient of the Zernike mode i, j (i is the polynomial degree and j the azimuthal frequency), and N the number of modes used for the decomposition (here $N=15$). The decomposition allows to better appreciate the contribution of each aberration to the phase profile. The sensor does not measure the piston, this term (Z_0^0) is then neglected.

In order to test the efficiency of the neural network, the prescribed target wavefronts are chosen among the Zernike polynomials, with a given variable phase amplitude (PtV) and different sign. The quality of the experimental wavefront is then quantified in the following way. For each target, we visualize the residual phase error by plotting: $\varphi_{target}(x, y) - \varphi_{measured}(x, y)$. From this image, we deduce the residual RMS defined as (with N_p the number of pixels):

$$\sigma = \sqrt{\frac{1}{N_p} \sum [\varphi_{target}(x, y) - \varphi_{measured}(x, y)]^2} \quad (3)$$

From this value, the Strehl Ratio (S_r), which quantifies the focal spot degradation due to the error on the wavefront correction, can be estimated using the Maréchal approximation [37]:

$$S_r \approx 1 - \left(\frac{2\pi}{\lambda}\right)^2 \times \sigma^2 \quad (4)$$

Finally, each target and measured wavefront is projected on the Zernike polynomials (the tilt terms, Z_1^{-1} and Z_1^1 , are not shown). The projection allows us to assess the functional error of the TOA-SLM, through the weight of the unwanted Zernike modes. We thus define a generation efficiency factor for a given polynomial [9]:

$$A = \frac{\alpha_k^l}{\sqrt{\sum_{(i,j)=1}^N (\alpha_i^j)^2}} \quad (5)$$

With k, l the labels of the targeted polynomial.

We first consider the simple case of a parabolic (positive and negative) wavefront. Such adaptive focusing or defocusing optics can then be exploited for slow thermal lens fluctuations correction in high average power lasers systems [10]. The results are displayed in Fig. 4, the wavefront is featured as an optical phase difference (OPD, in nm). The target corresponds to a convex, respectively concave, mirror with $PtV = 0.4\lambda = 320$ nm. The wavefronts, measured when applying the predicted DLP pattern, are in good agreement with the targets, with a residual RMS error of 51 nm and 26 nm respectively. In both cases, the error is minimal at the center of the pupil, due to the weighted optimization function applied in the algorithm (Fig. 3 top row). In both cases, the experimental PtV is slightly different from the target. This is explained by the lack of piston measurement forcing for the network to optimize the phase gradient instead of the phase itself. For the negative wavefront curvature (i.e. convex mirror, Fig. 4(a)), PtV value is slightly overreached in the experimental case. Zernike projection shows a dominant Z_2^0 mode, as expected, and minor dispersion in the other modes with $A = 0.98$. For the positive wavefront curvature (i.e. concave mirror, Fig. 4(b)), PtV value is slightly lower than the target and an unwanted Z_2^{-2} term appears, with a coefficient ratio of 20% of the target mode. This parasitic term corresponds to a 45° astigmatism, visible on the edges of the experimental wavefront. The asymmetry between the negative and positive curvatures can be understood and originates from the TOA-SLM and optical setup. First, realizing a temperature profile with positive curvature in TOA-SLM is more difficult than a negative one, given the wide spatial diffusion of the thermal gradient [12]. Second, the concave mirror case is more likely to be affected by heat flow at the edges, as already seen in Fig. 1(c). Nevertheless, a good function efficiency is obtained in both cases.

Figure 5 displays the results obtained for targeted more complex higher-order aberrations. Astigmatism (Z_2^2) is generated with a target $PtV = 0.2\lambda = 160$ nm (Fig. 5(a)). The measured wavefront after prediction is of very good quality with a RMS error of only 10 nm and $A = 0.98$. Trefoil (Z_3^3) is also obtained with a low error (20 nm) and good efficiency ($A = 0.96$, Fig. 5(b)). The production of such phase pattern with low symmetry orders validates the predictive efficiency of the network and the training procedure. Finally, the results obtained for spherical aberration (Z_4^0) are shown in Fig. 5(c). This aberration remains usually challenging to produce, due to the deep phase modulation across the pupil. The RMS is 54 nm for a target $PtV = 0.4\lambda = 240$ nm and $A = 0.77$. In all three instances, the projection on the Zernike polynomials shows the generation of a parasitic term of non-negligible relative amplitude: Z_2^0 term, which stands for focus/defocus and comes from the phase error at the edges of the pupil. This error is due on the one hand to the weighted optimization function described in the former section, but also to the device itself, with its intrinsic difficulty in keeping the heat generated at the pupil edges.

Table 1 gathers the metrics (PtV , σ , S_r , A) corresponding to all the wavefronts tested in this study. For all of them, the residual phase error fulfills the Maréchal criterion ($S_r > 0.8$). We have demonstrated that the developed neural network for TOA-SLM can efficiently generate Zernike polynomials up to the fourth-order, opening the way for aberration compensation using this technology, with a correction up to a few hundreds of nm across a 10 mm pupil, limited by the available blue light power. The generation performance is comparable to other types of thermally

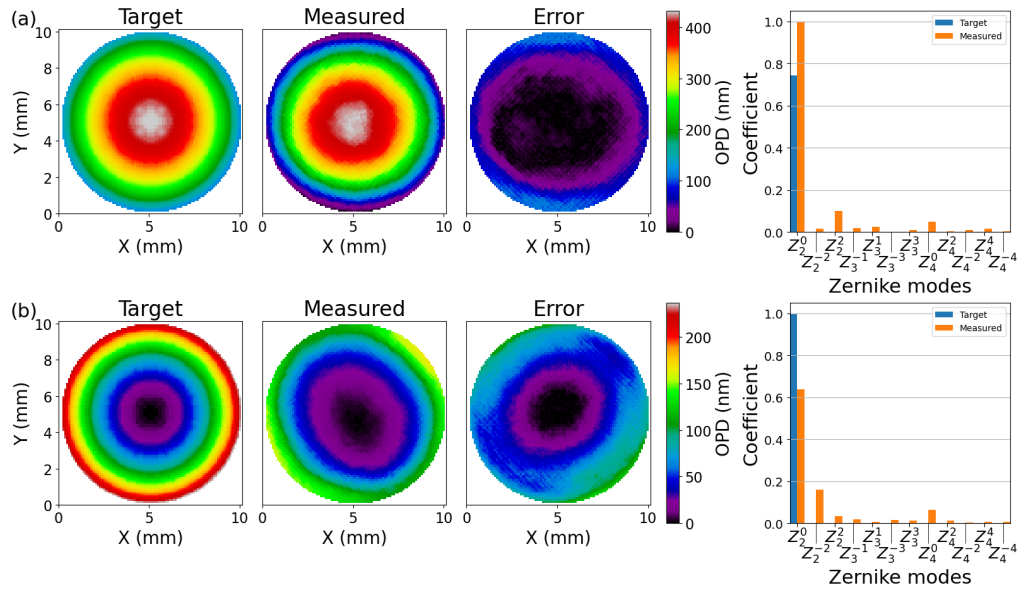


Fig. 4. Focusing (a) and defocusing (b) mirror (Z_2^0): targeted wavefront, measured wavefront, residual phase difference and wavefronts projection of the Zernike modes.

deformable mirrors [9]. Higher-order modes may be compensated by a larger beam aperture, so as to overcome the limited spatial resolution. Thus, these results endorse the performance of the neural network as an efficient phase designer, as well as the performances of the TOA-SLM in terms of wavefront corrections.

Table 1. Metrics from all the wavefronts tested in this study.

Target	PtV_{target} (nm)	PtV_{exp} (nm)	σ (nm) (Eq. (3))	S_r (Eq. (4))	A (Eq. (5))
Z_2^0 (foc.) (Fig. 4(a))	320	432	35	0.94	0.98
Z_2^0 (defoc.) (Fig. 4(b))	320	225	30	0.95	0.96
Z_2^0 (foc.)	240	430	51	0.86	0.98
Z_2^0 (defoc.)	240	169	26	0.96	0.95
Z_2^0 (foc.)	160	329	40	0.90	0.98
Z_2^0 (defoc.)	160	107	20	0.97	0.94
Z_2^2 (Fig. 5(a))	160	190	10	0.99	0.97
Z_3^1	160	220	48	0.86	0.9
Z_3^3 (Fig. 5(b))	240	298	21	0.97	0.96
Z_3^3	320	347	23	0.96	0.94
Z_4^0 (Fig. 5(c))	240	237	54	0.83	0.77
Z_4^0	160	198	19	0.98	0.9

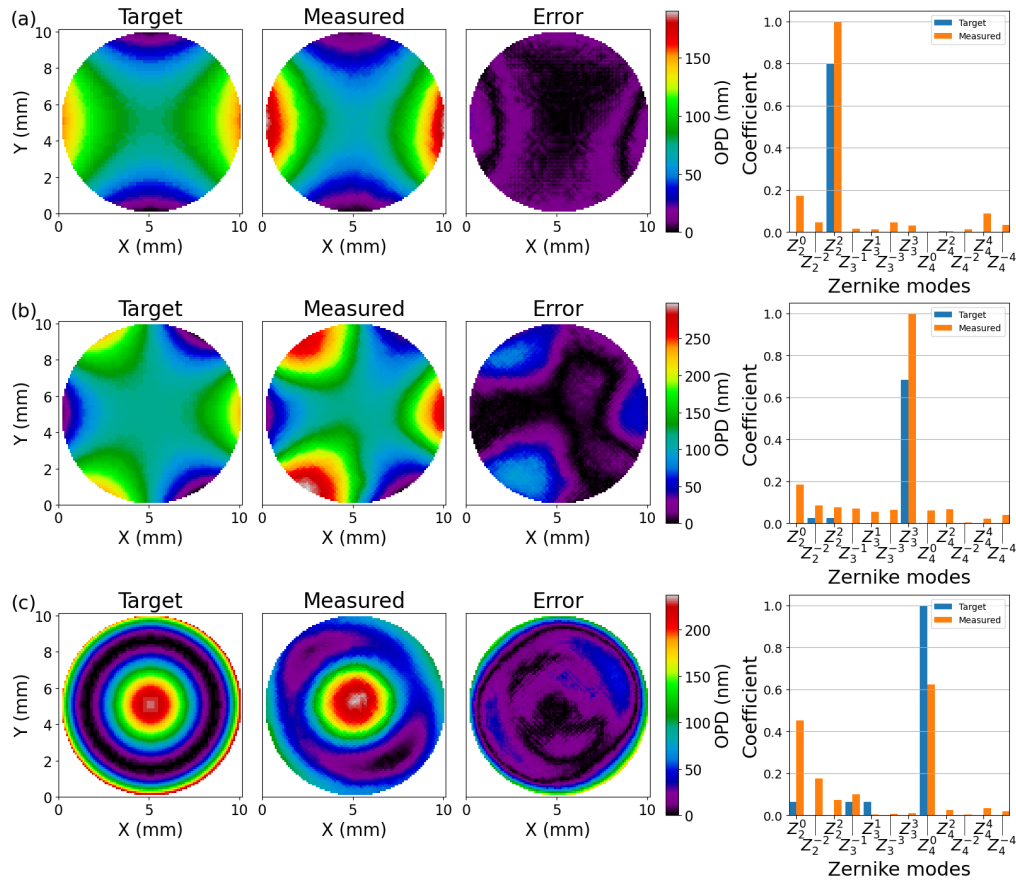


Fig. 5. Z_2^2 (a), Z_3^3 (b) and Z_4^0 (c): targeted wavefront, measured wavefront, residual phase difference and wavefronts projection of the Zernike modes.

5. Conclusion

To conclude, we have proposed the reconfigurable design of a novel thermo-optically addressed SLM (TOA-SLM) by a neural network. The phase modulation function relies on local heating of a thick birefringent and thermotropic LC layer, induced by local absorption of a shaped control beam in a metallic layer. While the physical system in itself is very promising, its actual use requires overcoming two difficulties. The first is to produce and tune a quantitatively accurate physical model of the experimental device. This model must include heat diffusion across a complex multilayer structure with poorly known boundary conditions (LC cell assembly dependent) and nonlinear thermo-optic effect. The second is to actually guess from that model what input is needed to achieve a prescribed phase profile. While the first task is difficult but doable, the second is analytically impossible and therefore would rely on a very time consuming and unreliable trial and error process based on numerical simulations of the model. In order to overcome these difficulties, we bypass physical modelling altogether and train a neural network to solve the inverse problem directly from experimental data.

We have shown that a relatively small convolutional network (about 2.4×10^5 weights) can be trained very well on the moderate amount of existing experimental data. After training, the network is capable of accurately predicting (in less than 0.1 s) a control beam profile leading to many kinds of prescribed phase profiles, as we have illustrated here on convergent and divergent

mirrors and Zernike polynomials up to the fourth order, with a stroke several hundreds of nm. While these results (including a generation efficiency above 90%) fully validate the overall approach, several avenues for improvement exist. First, designing an experimental measurement capable of providing an accurate piston value would allow for much better prediction capabilities: the model would then consider the physically relevant absolute values of the phase instead of only gradients. Second, progressing on a quantitatively accurate physical model would open the way to transfer learning for the inverse problem after training a network on numerical data (perhaps including a tandem network approach [20,35]). The neural network approach could also be used to learn the thermal conditions at the cell boundaries.

Even at the current stage, we have demonstrated that a TOA-SLM device can be accurately operated thanks to a neural network-based statistical model. This step, together with the ultra-broadband acceptance of the device and its ability to introduce continuous and deep phase modulation over a large aperture, opens the way for ultrafast laser aberration compensation using this new technology.

Funding. Agence Nationale de la Recherche (ANR-19-CE30-0006-01, ANR-21-PRRD-0001-01); Association Nationale de la Recherche et de la Technologie (2019/0660).

Disclosures. The authors declare no conflicts of interest.

Data availability. Data underlying the results presented in this paper are not publicly available at this time but may be obtained from the authors upon reasonable request.

References

1. C. Rosales-Guzman and A. Forbes, *How to Shape Light with Spatial Light Modulators* (SPIE. Spotlight, 2017).
2. N. Collings, T. Davey, J. Christmas, D. Chu, and B. Crossland, "The applications and technology of phase-only liquid crystal on silicon devices," *J. Display Technol.* **7**(3), 112–119 (2011).
3. P.-A. Blanche, "Holography, and the future of 3d display," *Light: Adv. Manuf.* **2**, 28 (2021).
4. S. Residori, U. Bortolozzo, and J. P. Huignard, "Liquid crystal light valves as optically addressed liquid crystal spatial light modulators: optical wave mixing and sensing applications," *Liq. Cryst. Rev.* **6**(1), 1–16 (2018).
5. A. S. Ostrovsky, C. Rickenstorff-Parrao, and V. Arrizón, "Generation of the perfect optical vortex using a liquid-crystal spatial light modulator," *Opt. Lett.* **38**(4), 534–536 (2013).
6. A. M. Weiner, "Femtosecond pulse shaping using spatial light modulators," *Rev. Sci. Instrum.* **71**(5), 1929–1960 (2000).
7. F. Snik, O. Absil, and P. Baudoz, *et al.*, "Review of high-contrast imaging systems for current and future ground-based and space-based telescopes III: technology opportunities and pathways," in *Advances in Optical and Mechanical Technologies for Telescopes and Instrumentation III*, vol. 10706 R. Navarro and R. Geyl, eds., International Society for Optics and Photonics (SPIE, 2018), p. 107062L.
8. O. Albert, L. Sherman, G. Mourou, T. B. Norris, and G. Vdovin, "The smart microscope : an adaptive optics learning system for aberration correction in multiphoton confocal microscopy," *Opt. Lett.* **25**(1), 52–54 (2000).
9. M. Kasprzak, B. Canuel, F. Cavalier, R. Day, E. Genin, J. Marque, D. Sentenac, and G. Vajente, "Performance of a thermally deformable mirror for correction of low-order aberrations in laser beams," *Appl. Opt.* **52**(12), 2909–2916 (2013).
10. S. Bonora, J. Pilar, A. Lucianetti, and T. Mocek, "Design of deformable mirrors for high power lasers," *High Power Laser Sci. Eng.* **4**, e16–e20 (2016).
11. P. S. Salter and M. J. Booth, "Adaptive optics in laser processing," *Light: Sci. Appl.* **8**(1), 110 (2019).
12. V. M. DiPietro, S. Bux, N. Forget, and A. Jullien, "Phase-only pulse shaper for multi-octave light sources," *Opt. Lett.* **45**(2), 543–546 (2020).
13. J. Cannon, "Some numerical results for the solution of the heat equation backwards in time," in *Numerical Solutions of Nonlinear Differential Equations*, (J. Wiley New York, 1966), pp. 21–54.
14. C.-S. Liu and C.-W. Chang, "Analytic series solutions of 2d forward and backward heat conduction problems in rectangles and a new regularization," *Inverse Probl. Sci. Eng.* **28**(10), 1384–1406 (2020).
15. G. Genty, L. Salmela, J. M. Dudley, D. Brunner, A. Kokhanovskiy, S. Kobtsev, and S. K. Turitsyn, "Machine learning and applications in ultrafast photonics," *Nat. Photonics* **15**(2), 91–101 (2022).
16. D. Faccio, A. Velten, and G. Wetzstein, "Non-line-of-sight imaging," *Nat. Rev. Phys.* **2**(6), 318–327 (2020).
17. S. Barland and F. Gustave, "Convolutional neural network for self-mixing interferometric displacement sensing," *Opt. Express* **29**(8), 11433–11444 (2021).
18. M. Del Hougne, S. Gigan, and P. Del Hougne, "Deeply subwavelength localization with reverberation-coded aperture," *Phys. Rev. Lett.* **127**(4), 043903 (2021).
19. L. Pilozzi, F. A. Farrelly, G. Marcucci, and C. Conti, "Machine learning inverse problem for topological photonics," *Commun. Phys.* **1**(1), 57 (2018).

20. D. Liu, Y. Tan, E. Khoram, and Z. Yu, "Training deep neural networks for the inverse design of nanophotonic structures," *ACS Photonics* **5**(4), 1365–1369 (2018).
21. P. R. Wiecha, A. Arbouet, C. Girard, and O. L. Muskens, "Deep learning in nano-photonics: inverse design and beyond," *Photonics Res.* **9**(5), B182–B200 (2021).
22. J. R. P. Angel, P. Wizinowich, M. Lloyd-Hart, and D. Sandler, "Adaptive optics for array telescopes using neural-network techniques," *Nature* **348**(6298), 221–224 (1990).
23. D. Sandler, T. Barrett, D. Palmer, R. Fugate, and W. Wild, "Use of a neural network to control an adaptive optics system for an astronomical telescope," *Nature* **351**(6324), 300–302 (1991).
24. Y. Guo, L. Zhong, L. Min, J. Wang, Y. Wu, K. Chen, K. Wei, and C. Rao, "Adaptive optics based on machine learning: a review," *Opto-Electron. Adv.* **5**(7), 200082 (2022).
25. I. Goodfellow, Y. Bengio, and A. Courville, *Deep learning* (MIT, 2016).
26. J. Li, S. Cauza, and S.-T. Wu, "Temperature effect on liquid crystal refractive indices," *J. Appl. Phys.* **96**(1), 19–24 (2004).
27. V. M. DiPietro, A. Jullien, U. Bortolozzo, N. Forget, and S. Residori, "Thermally-induced nonlinear spatial shaping of infrared femtosecond pulses in nematic liquid crystals," *Laser Phys. Lett.* **16**(1), 015301 (2019).
28. V. M. DiPietro and A. Jullien, "Broadband spectral domain interferometry for optical characterization of nematic liquid crystals," *Appl. Sci.* **10**(14), 4701 (2020).
29. P. Bon, G. Maucort, B. Wattellier, and S. Monneret, "Quadriwave lateral shearing interferometry for quantitative phase microscopy of living cells," *Opt. Express* **17**(15), 13080–13094 (2009).
30. F. Chollet, "keras," (2015).
31. O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*, (Springer, 2015), pp. 234–241.
32. M. D. Zeiler, D. Krishnan, G. W. Taylor, and R. Fergus, "Deconvolutional networks," in *2010 IEEE Computer Society Conference on computer vision and pattern recognition*, (IEEE, 2010), pp. 2528–2535.
33. F. Chollet, *Deep learning with Python* (Simon and Schuster, 2021).
34. P. Gavrikov, "visualkeras," <https://github.com/paulgavrikov/visualkeras> (2020).
35. A. Estrada-Real, A. Khairah-Walieh, B. Urbaszek, and P. R. Wiecha, "Inverse design with flexible design targets via deep learning: Tailoring of electric and magnetic multipole scattering from nano-spheres," *Photonics Nanostructures-Fundamentals Appl.* **52**, 101066 (2022).
36. M. Born and E. Wolf, *Principles of Optics* (Pergamon, 1980).
37. J. Herrmann, "Phase variance and strehl ratio in adaptive optics," *J. Opt. Soc. Am. A* **9**(12), 2257–2258 (1992).