








Real-time machine-learning-driven control system of a deformable mirror for achieving aberration-free X-ray wavefronts

LUCA REBUFFI,^{1,*}  XIANBO SHI,¹  ZHI QIAO,¹  MATTHEW J. HIGHLAND,¹ MATTHEW G. FRITH,¹ ANTOINE WOJDYLA,²  KENNETH A. GOLDBERG,²  AND LAHSEN ASSOUFID¹

¹Argonne National Laboratory, 9700 S Cass Ave, Lemont, IL 60439, USA

²Lawrence Berkeley National Laboratory, 1 Cyclotron Rd, Berkeley, CA 94720, USA

*lrebuffi@anl.gov

Abstract: A neural-network machine learning model is developed to control a bimorph adaptive mirror to achieve and preserve aberration-free coherent X-ray wavefronts at synchrotron radiation and free electron laser beamlines. The controller is trained on a mirror actuator response directly measured at a beamline with a real-time single-shot wavefront sensor, which uses a coded mask and wavelet-transform analysis. The system has been successfully tested on a bimorph deformable mirror at the 28-ID IDEA beamline of the Advanced Photon Source at Argonne National Laboratory. It achieved a response time of a few seconds and maintained desired wavefront shapes (e.g., a spherical wavefront) with sub-wavelength accuracy at 20 keV of X-ray energy. This result is significantly better than what can be obtained using a linear model of the mirror's response. The developed system has not been tailored to a specific mirror and can be applied, in principle, to different kinds of bending mechanisms and actuators.

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

1. Introduction

Achieving and maintaining high-intensity focused X-ray beams with near-perfect wavefront quality and high stability is the main challenge and prerequisite for experiments at 4th-generation synchrotron radiation and free electron laser (FEL) beamlines. This represents the main challenge for X-ray optical elements, which will necessarily have much more demanding specifications than those for other applications because of the shorter wavelength and the ultra-small emittance of the radiation beams generated by these sources. When using coherent photons from diffraction-limited light sources, it is critical to maintain a well-controlled wavefront and suppress unwanted static distortions and dynamic disturbance [1]. Not only can the degradation of the wavefront be detrimental for phase-sensitive imaging techniques like tomography [2], but wavefront uniformity is of particular importance for coherent X-ray scattering experiments using X-ray Photon Correlation Spectroscopy, Coherent Surface Scattering Imaging, and Coherent X-ray Diffraction Imaging techniques. Wavefront distortions degrade the sample speckle contrast, which can hinder data interpretation [3] and lead to the failure of the sample phase retrieval and reconstruction. Therefore, the optical elements must i) be manufactured with a surface figure closely following an ideal mathematical shape, ii) automatically and repeatably align and focus the beam to match different sample and experiment requirements, and iii) provide real-time correction in response to wavefront deformations induced by thermal distortion of upstream optics, mechanical drifting, and other forms of photon-beam perturbations [4]. Adaptive Optics (AO) system is a viable solution to address the above problems.

Such a system can dynamically compensate for measured wavefront aberrations by using precise and highly controllable deformable mirrors [5]. However, up to now, the control of adaptive mirrors has relied on the linearity, dynamics, and repeatability of the optics response

to the actuators, which include mechanical bending, piezoelectric bimorph deformation, and thermal loading [6,7]. The use of piezoelectric actuators to induce deformation of relatively thick substrates is associated with nonlinearities, such as crosstalk between actuators, creep, and hysteresis [8–10]. The magnitude and rate of creep and hysteresis are affected by various factors, including the recent history of applied voltages [9]. These nanometer-scale effects are already relevant to the level of accuracy required for applications mentioned above.

The central challenge of controlling X-ray AO relying on piezoelectric actuators is the ability to compensate for their dynamic behavior. Without feedback and dynamic control, the mirrors step toward their targeted shapes and then drift slowly over minutes. Furthermore, the shapes they achieve are influenced by their recent shape history. Thus, repeated actuation to the same target shape can produce different results. Traditional creep compensation was carried out by applying a time-dependent compensating voltage profile immediately following the initial voltage application, which can give good results when the mathematical model is known, such as the creep radius of curvature is linear to the applied voltage [9].

With the help of a feedback system, iterative control methods based on the linear response model can be used [11,12], which can take tens of minutes to converge to the targeted solution and are, therefore, unsuitable for real-time AO control. Recent work demonstrated the possibility of controlling a bimorph mirror using machine learning (ML) based on a neural network (NN) trained on a collection of surface shapes measured by a Fizeau interferometer [13]. A closed-loop, fast-shaping control of a bimorph mirror was recently reported for the surface shape feedback system of an array of interferometric sensors [14].

In this work, we report the first demonstration of an NN-based controller of a bimorph mirror, achieving the desired wavefront shape directly at the beamline under working conditions. The bimorph mirror uses piezo actuators to provide global bending and local adjustment of the surface shape. Single-shot, absolute wavefront sensing was demonstrated using a coded-mask-based technique [15] and wavelet-transform-based analysis [16]. The NN controller was trained with measured, one-dimensional wavefront data at the Instrumentation Development, Evaluation, & Analysis (IDEA) beamline (28-ID-B) of the Advanced Photon Source (APS).

The NN-based controller achieved a response time of a few seconds and maintained target wavefront shapes (e.g., a spherical phase) with sub-wavelength accuracy at 20 keV. The system's speed, shape accuracy, and stability were significantly improved compared with a linear model of the actuation response functions. The results are a promising step toward the ultimate goal of achieving real-time 2D wavefront control with sub-second speed and picometer accuracy, which are crucial for next-generation beamlines. The developed ML model can also be applied to adaptive and active optics in broad applications, such as astronomical imaging.

2. Single-shot absolute wavefront sensing

The coded-mask-based X-ray wavefront sensing technique [15] is a newly developed variation of the near-field speckle-tracking method [17,18]. The main principle of the speckle-tracking method is to introduce a speckle generator, such as a piece of sandpaper or membrane filter, into the beam to induce wavefront phase modulation and, thus, a near-field speckle image on a downstream detector. Typically, the absolute beam wavefront can be extracted from two images recorded at two speckle generator-to-detector distances, D_1 and D_2 . Taking a one-dimensional case as an example, by tracking the local displacement of the speckle pattern, $\delta(x)$, the beam deflection angle can be determined as

$$\alpha(x) = \frac{\delta(x)}{D_2 - D_1}. \quad (1)$$

The differential wavefront phase (phase gradient) is then:

$$\frac{d\phi(x)}{dx} = \frac{2\pi\alpha(x)}{\lambda}, \quad (2)$$

where λ is the wavelength.

The wavefront phase $\phi(x)$ can be obtained by integrating the phase gradient in the x direction. The local curvature of the wavefront $C(x)$ is the derivative of the phase gradient, given by

$$C(x) = \frac{1}{R(x)} = \frac{\lambda}{2\pi} \frac{d^2\phi(x)}{dx^2} = \frac{d\alpha(x)}{dx}, \quad (3)$$

where $R(x)$ is the local radius of curvature.

In this work, single-shot absolute wavefront sensing is realized using a coded phase mask as the speckle generator. The coded mask has a known design pattern that allows for the construction of a reference image by simulating a plane wave (constant phase) propagating through the mask to the detector at a mask-to-detector distance D . In this case, $\delta(x)$ can be obtained by comparing a single measured image with the simulated reference image. The local beam deflection angle, under the small angle approximation, becomes:

$$\alpha(x) = \frac{\delta(x)}{D}. \quad (4)$$

There are many algorithms for extracting $\delta(x)$, such as cross correlation [17], wavelet-transform [16], transport of intensity equation [19], maximum-likelihood optimization [15], and deep learning [20]. This work used the wavelet-transform-based speckle vector tracking (WSVT) method [16] for its controllable trade-off between resolution and speed and its enhanced noise robustness, which is superior to the direct, cross-correlation method. The full mathematical details of the wavelet-transform-based algorithm can be found in [16]. Here, we summarize the basic procedures for extracting the 1D displacement $\delta(x)$ from two 2D images (one measured image and one simulated reference). The one-dimensional analysis is appropriate as feedback for adaptive X-ray optics where the mirror shape actuation takes place strictly along the tangential direction of the mirror in one dimension.

For each pixel on the sample image [e.g., Fig. 1(a)], we construct a 1D signal array $I_x^s(t)$ from all the neighboring pixels within a template window size of N_w pixels, where $t \in [1, N_w]$. Similarly, $I_x^r(t)$ can be reconstructed for each pixel on the simulated reference image. In a traditional correlation-based method, each $I_x^s(t)$ will be compared to $I_x^r(t)$ for all the pixels. The displacement δx of a sample pixel is determined by finding the maximum cross-correlation coefficient between $I_x^s(t)$ and $I_{x+\delta x}^r(t)$. In WSVT, each $I_x^s(t)$ and $I_x^r(t)$ will be first decomposed into wavelet coefficients through discrete wavelet transform (DWT). Because DWT is an orthonormal transform, one can use the Euclidean distance (difference square of the wavelet coefficients) to compute the similarities between $I_x^s(t)$ and $I_x^r(t)$. This process is faster than the cross-correlation operation. In addition, the length of the wavelet coefficients can be truncated (by a factor of two in this work) to improve the speed even further without losing resolution.

Examples of the reconstructed 1D speckle displacement function δ and the corresponding phase ϕ and local radius of curvature R are shown in Figs. 1(b), 1(c), and 1(d), respectively. These profiles are candidate targets for the ML training (see Section 5.2 for comparison). The average time to obtain these profiles is only a few seconds, achieved by carrying a multi-thread data analysis in the background during the measurements. Such a high-speed analysis is essential to enable near-real-time data acquisition. Further improvement is under development by implementing deep learning [20].




Fig. 1. Examples of (a) sample image showing speckle pattern of radiation with photon energy of 20 keV from the coded mask, (b) reconstructed local displacement of speckle, (c) phase (in units of λ), and (d) local radius of curvature along the vertical direction (y coordinate).

3. Bimorph deformable mirror

The testing mirror is a PZT (lead zirconate titanate)-glued bimorph mirror manufactured by JTEC Corporation, as shown in Fig. 2. The Pt-coated, optically active area of the mirror has dimensions of 150 mm (length) \times 8 mm (width) on a silicon substrate with a dimension of 160

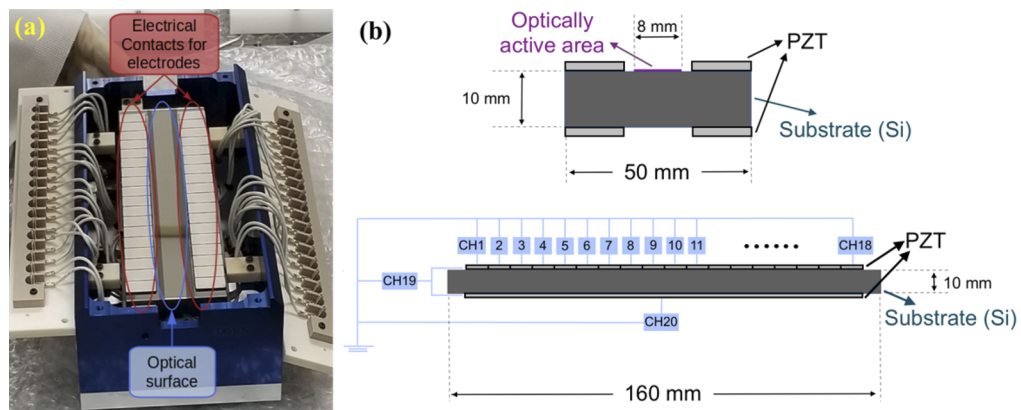


Fig. 2. Bimorph mirror: (a) photograph showing the 18 electrode pairs for the local shaping actuators and the optical surface sandwiched between the electrodes, (b) schematic for substrate dimensions and the piezo electrode channels.

mm (length) \times 50 mm (width) \times 10 mm (thickness). Two piezoelectric strips, with 18 separate electrodes, are glued symmetrically on each side of the optical surface. Each pair of electrodes, at the same position along the strips, forms one actuator capable of locally modifying the surface shape. There are 18 local surface actuators (ch1–ch18), each corresponding to one channel and one input voltage. Two long piezo strips are glued to the bottom surface of the mirror, forming a single global bender actuator (ch20) capable of bending the entire optical surface. A grounded channel (ch19) is on the backside of all piezo stripes.

The initial state of the mirror was set to 500 V on all electrodes. Then, the voltage applied to the bender actuator was tuned to 170 V to bend the mirror surface and form the focus in the desired position. In this study, as in [13], we limited the voltage change applied to the 18 local actuators to within a range of [−100 V, 100 V] for the training data collection.

4. ML-based controller of the bimorph mirror

The controller is built by extending the system described in [13]. The system is a two-part framework: (i) an ML-based predictive model of the non-linear dynamics of the adaptive optics using a feedforward neural network and (ii) a controller to drive the mirror to obtain a desired spatial distribution of a chosen wavefront property using non-linear quadratic cost regulation over a finite horizon. The system can use one of the wavefront properties that can be measured at the beamline: local speckle displacement δ (directly related to the differential phase), integrated absolute phase ϕ , or local radius of curvature, R , of the wavefront.

4.1. NN-based predictive model

The predictive model is a discrete-time (Δt) model of the non-linear dynamics of the bimorph mirror. For any time t , the predicted wavefront property wp_{t+1} , at $t + \Delta t$, is expressed as a function of a history of n discrete-time set of voltages $v_{t-k} = [v_i^{actuator}]_{t-k}^{i=1 \dots N_a}$ applied to the chosen N_a actuators at time $t - k \cdot \Delta t$ and corresponding obtained wavefront property wp_{t-k+1} at time $t - (k - 1) \cdot \Delta t$, with $k = 0, 1 \dots n$ ($0 =$ current time), given by

$$wp_{t+1} = f(wp_t, wp_{t-1}, \dots, wp_{t-n+1}, v_t, v_{t-1}, \dots, v_{t-n}). \quad (5)$$

The meaning of wp_{t+1} is that the model predicts the wavefront property spatial distribution one time step in the future given a finite history of wavefront property shapes and voltages, starting with its current shape, and applied set of voltages, going back n time steps. In this study, we kept $\Delta t = 2$ s, as used in [13], for comparison purposes and to safely have time to change the input voltages and collect the wavefront measurement. The chosen wavefront property is detected as a 2D spatial distribution and mapped to 1D profiles along vertical and horizontal coordinates. However, since the bimorph mirror only focuses in one direction, we used the corresponding 1D vertical profile obtained from the selected portion of the detected image (see Fig. 1).

The machine learning system to determine the function f is a feedforward neural network [21], with five fully connected layers and exponential linear unit activation, and an additional skip connection introduced due to their effectiveness in modeling an identity mapping [21,22]. The size of the input layer is determined by the dimension of our wavefront property spatial distribution, the number of actuators being controlled, and the number of previous time steps incorporated in making a prediction.

In this experiment, the beam size at the 28-ID beamline illuminated an area of the bimorph mirror surface corresponding to the central 10 actuators (ch5–ch14). The 1D section of the selected wavefront property is limited to 700 points (pixels), corresponding to a central portion of the whole detected image, where the signal is unaffected by the effects of upstream beamline optics that blur the boundaries of the image (e.g., diffraction from slits.) To reduce the computational time, we represented the wavefront property with a spline of 50 points with an added quality

check by setting a limit on the difference between the original distribution and spline to avoid loss of accuracy.

To train the NN model, we performed 15-fold cross-validation on the training dataset with 3,000 iterations of the Adam stochastic optimization [23] per fold. The PyTorch library [24] was used as the ML framework. Predictive models with different n time step histories ($n = 4, 5,$ and 10) were trained for comparison.

4.2. NN-based controller

The mirror controller uses the trained NN model to determine a sequence of voltage inputs to be applied from the current state to a final state in a finite number of m control time steps (with the same Δt used for the training dataset for the predictive model). The controller algorithm solves a quadratic cost function with an expression similar to the iterative linear quadratic regulator [25] that penalizes the prediction error at each intermediate step. As described in [13], the controller provides the succession of voltages to be applied to the actuators by directly minimizing the non-convex objective function using Adam [23]. We explicitly chose to maintain the conditions of our previous study for comparison purposes, even though it does not theoretically guarantee the optimality of the converged solution.

The cost function is given by

$$\{v\}^* = \arg \min_{\{v\}} \sum_{t=1}^m c_t \cdot \|wp_t - wp^*\|_2^2, \quad (6)$$

with initial conditions

$$wp_t = wp_0 \quad \forall t \leq t_0, \quad v_t = v_{0-1} \quad \forall t \leq t_0, \quad (7)$$

where wp^* is the desired wavefront property spatial distribution, and c_t coefficients are used to apply a weighted penalty to the prediction error at each time step. We assumed that the first set of voltages, v_0 , is applied to the actuators at the time step 0, when the mirror is at rest with some arbitrary shape. Additionally, we used constant weights c_t that have been optimized empirically [13]. In this work, we have chosen an $m = 10$ step controller for all experiments.

5. Experimental results at the IDEA beamline

5.1. Experimental setup

The 28-ID-B beamline at APS is designed to characterize the performance of state-of-the-art X-ray optics and devices planned for the Advanced Photon Source Upgrade (APS-U). The scientific goal of the IDEA beamline is to obtain performance metrics for proposed beamline optics and components for the APS-U by simulating the expected upgraded operating conditions for the multi-bend achromat undulator source and to provide data to validate, optimize, or re-engineer the best possible performance of all the planned new and enhanced beamlines. In particular, the testing program focuses on performance under high heat load and coherence preservation by beamline optics and components.

The IDEA beamline source is described in Table 1. The beamline has a horizontal double multilayer monochromator (HDMM) to provide large-bandwidth ($1 \sim 1.5\%$) monochromatic X-ray beams. The optical components of the beamline are listed in Table 2. The reported distances are relative to the center of the undulator.

In experimental hutch B, we built a Kirkpatrick–Baez (KB) mirror focusing system composed of an outward-reflecting, horizontally focusing bendable mirror using a flexure bender [26] and the upward-reflecting bimorph mirror for vertically focusing. Both mirrors were operated at a grazing angle of 3 mrad. The KB mirror pair creates a focused beam at the focal plane, upstream of the wavefront sensor (coded mask and detector), as shown in Fig. 3, to create a $\sim 1 \text{ mm}^2$ square

Table 1. Characteristics of undulator U33 at APS.

Period (mm)	Number of Periods	K	E_{1st} (keV)	Electron Beam					
				E (GeV)	I (mA)	σ_x (μm)	σ_y (μm)	σ'_x (μrad)	σ'_y (μrad)
33	71	2.247	4.0	7	100	280.5	10.2	11.8	3.4

Table 2. Optical components of the IDEA beamline. Distances are relative to the center of the undulator.

Distance (m)	Component	Dimension (mm^2)	Description/Comments)
28.3	Slits	0.06×1.0	White beam slit
30.0	Mirror (M1)	30×500	Horizontal, outward reflecting, bendable mirror
30.7	Mirror (M2)	30×500	Horizontal, inward reflecting, bendable mirror
32.7	HDMM	30×160	Horizontal double multilayer monochromator
63.6	Slits	0.42×0.85	Monochromatic beam slit
64.97	H-KB		Horizontally focusing bender mirror
66.1	V-KB		Vertically focusing bimorph mirror
68.6	C-M		Coded mask (5- μm pattern period, 2- μm -thick Au electroplated on silicon nitride substrate)
69.1	Detector		20- μm thick LuAG(Ce) scintillator, 10 \times objective lens, and Zyla 5.5 sCMOS camera. Effective pixel size: 0.65 μm

image on the detector (see Fig. 1). Details of the coded mask and detector system can also be found in Table 2. The monochromator was tuned to a photon energy of 20 keV ($\lambda = 0.062$ nm) using the 5th harmonic of the undulator.

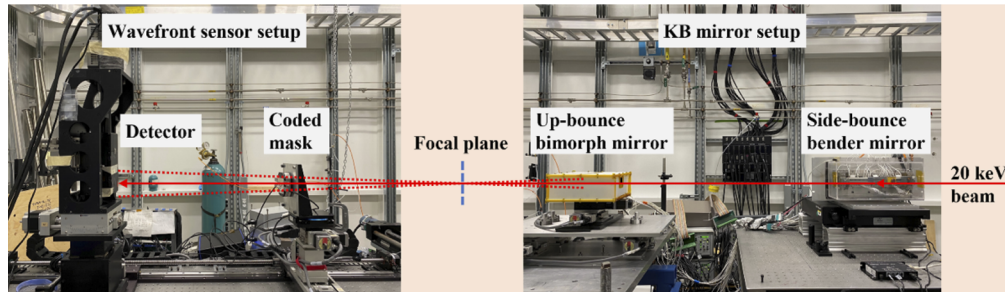


Fig. 3. KB mirror setup and wavefront sensor system in the experimental hutch B at the APS IDEA beamline.

5.2. Data collection and model training

To keep the mirror conditions close to the conditions used in [13], for comparison, the global bender channel (ch20) was excluded from machine learning in the first set of experiments. It was set to a fixed voltage of 170 V to keep the overall shape of the mirror and the focal position fixed. The voltages of ch1–ch4 and ch15–ch18 were fixed at 500 V and excluded from the machine learning since the beam footprint did not illuminate the areas modified by their relative actuators. The voltage inputs of each used surface actuator (ch5–ch14) were limited to a range of $[-100 \text{ V}, +100 \text{ V}]$ around 500 V to avoid possible damage to the mirror. Finally, the timestep was set to $\Delta t = 2$ s, as used in [13]; this was achieved by starting an image collection at 1.7 s at each step with an acquisition time of 0.3 s.

To train the NN, we acquired 24,500 wavefront measurements corresponding to random applied voltage inputs in the following datasets:

- 10,000 wavefronts corresponding to a random set of voltages, changing every step.
- 5,000 wavefronts corresponding to a random set of voltages, changing every 5 steps.
- 5,000 wavefronts corresponding to a random set of voltages, changing every 10 steps.
- 4,500 wavefronts corresponding to a random set of voltages, changing every 15 steps.

These training datasets cover conditions on large and immediate wavefront changes caused by adjustments of the voltages applied to all the actuators, as well as by small and slow changes due to the non-linear dynamic of the piezoelectric motors when the voltages are held constant. An additional 1,000 datasets were collected but not used in the model training. Instead, they served as a random shape database to be selected from as target shapes guaranteed to be achievable.

Firstly, the model was trained by using the three wavefront properties (see section 2) separately: speckle displacement, the local radius of curvature, and the absolute phase. The $m = 10$ -step control experiments were carried out using the $n = 4$ predictive model trained on the three wavefront properties.

For each control experiment to achieve one particular wavefront profile, the wavefront was first measured at the time t_0 at the current mirror condition. Then, the NN-based controller changed the voltages of the mirror actuators (ch5–ch14) in $m = 10$ consequential steps (every $\Delta t = 2$ s). The wavefront was measured during the last 0.3 s of each step before changing the voltages. Then, there were 10 measured wavefront profiles at t_1 to t_{10} . In addition, a final wavefront at t_{fin} was measured after another 40 s (at 60 s from the first voltage change) without changing the actuator voltages. The 11 measured wavefront profiles were then compared to the desired wavefront target for error analysis. To improve the measurement, 25 experiments to achieve 25 different but random wavefront profiles were carried out using NN-based models trained on the three wavefront properties, with results shown in Fig. 4.




Fig. 4. Comparison between the prediction error of NN-based models trained on different wavefront properties: (a) speckle displacement, (b) local radius of curvature, and (c) absolute phase. The target wavefronts are random profiles. Average rms error values are reported in the insets.

We used the rms phase error (in units of λ) between the desired and measured wavefront profiles as the criteria. Figure 4 shows the rms phase error distribution of the 25 experiments. The average rms phase error and its standard deviation are listed in the figure legends for time steps at t_1 , t_{fin} , and all the 11 time steps. It is evident that the model trained on the local displacement [Fig. 4(a)] provides the best result with the lowest average phase error and narrowest distribution (smallest $\pm\sigma$ value in the figure). This is reasonable since the speckle displacement is the most

direct measurement, while phase and radius of curvature are both secondary properties derived with additional integration or differentiation processes. Throughout the rest of this work, unless otherwise stated, all results were obtained using models trained on the local displacement.

It is worth noting that the capability of distinguishing one wavefront from another (detecting wavefront aberrations) depends on the resolution and sensitivity of the wavefront sensor. The coded-mask-based technique in this work provides the highest available phase sensitivity among all single-shot methods, which is at the $\lambda/100$ level [14]. This is critical for monitoring and controlling wavefronts at low-emittance light sources with super-precision adaptive optics. Another limiting factor is the stability of all beamline components. A wavefront stability measurement over a similar period (~ 75 minutes for 25 experiments) shows an rms phase stability of 0.02λ .

Next, control experiments using predictive models trained with $n = 4, 5,$ and 10 discrete-time steps are compared in Fig. 5. No significant differences were observed, suggesting that the original model with $n = 4$ for mirror shape prediction [13] is sufficient for the wavefront prediction. A time history of four steps is enough to account for the time-dependent dynamics of the mirror and thus accurately predict and control the wavefront. In addition, the phase error at t_1 being close to t_{fin} indicates that the ML controller can drive the mirror to the correct shape within the first step and keep it stable for all later steps. These fast-converging and stable features are essential for real-time feedback control of any adaptive mirror system.




Fig. 5. Comparison between the prediction error of models trained with different step histories: (a) $n = 4$, (b) $n = 5$, and (c) $n = 10$ to achieve randomly generated arbitrary wavefront profiles. Average rms error values are reported in the insets.

5.3. Aberration-free wavefront control

The ideal wavefront of a focusing beam is a spherical phase profile with a constant local radius of curvature equal to the optics-to-focus distance. Based on Eqs. (2) to (4), the differential phase profile $d\phi/dy$ and the measurable speckle displacement δ are both linear functions of y for a spherical wavefront. In practice, targeting such ideal spherical wavefronts is the most common and essential case at a beamline. In other words, the adaptive optics controller aims to remove any wavefront aberration from the desired profile.

While collecting the training datasets, we observed that the mean values of the local radius of curvature in all datasets fell between 1.43 m and 1.67 m, a range achievable within the set voltage limit of $[-100 \text{ V}, 100 \text{ V}]$ for electrodes ch5–ch14. For this reason, the control experiments were carried out to achieve a selected set of constant radii within the $[1.43 \text{ m}, 1.67 \text{ m}]$ range covered by the NN-based model.

Results from control experiments using the three n -steps trained models (Fig. 5) are summarized in Fig. 6. A similar fast-converging and stable performance is observed. The overall phase errors are larger than the random wavefront cases (Fig. 5) since the NN-based model was trained with

random wavefront outputs. Nevertheless, the wavefront phase error is small, for example, 0.265λ for the $n = 4$ model, which corresponds to an rms height error of 1.8 nm on the mirror surface obtained from back projection, which is close to the direct mirror shape control results (~ 1.4 nm) in [13]. This is a particularly satisfying result considering that in Ref. [13], the NN-based model was trained on a directly measured surface shape with the bender in its rest position, while in this study, the mirror was operated at an X-ray beamline, bent into a focusing shape. Also note that the contribution from the beamline instability (in particular, a drift of the vertical mirror pitch) over the time (~ 6 h) performing 125 experiments is about 0.06λ , which could explain the slight difference.




Fig. 6. Comparison between the prediction error of models trained with different step histories: (a) $n = 4$, (b) $n = 5$, and (c) $n = 10$ to achieve wavefront profiles with random but constant local radii of curvature. Average rms error values are reported in the insets.

As an example, Fig. 7 shows a comparison between the desired and measured wavefront phase and speckle displacement for one selected control experiment. The displacement control error is less than one pixel, also thanks to the high sensitivity provided by the wavefront sensor.

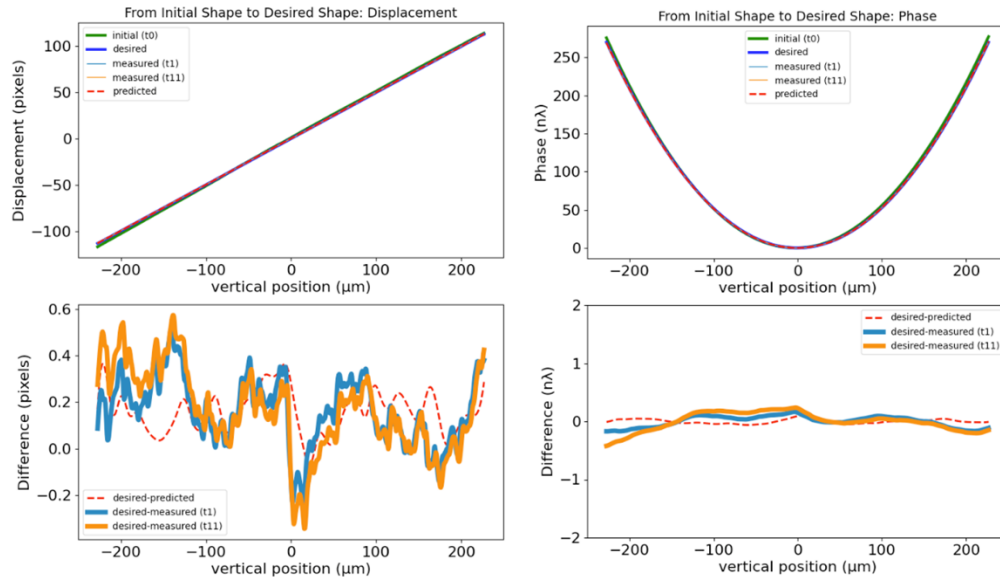


Fig. 7. Comparison of speckle displacement (left) and phase (right) between a spherical wavefront and the measured one achieved with the NN-based controller. Top: absolute values. Bottom: differences.

5.4. Comparison with linear controller

To understand the capability and the effectiveness of the NN-based controller, experiments were also carried out with a traditional linear controller based on the measured wavefront response functions (RF) of the actuators [10]. In this work, the piezo RF of each electrode was obtained by measuring the relative wavefront change when increasing each electrode voltage individually from 500 V to 600 V. The measured response functions for the phase and speckle displacement are shown in Fig. 8. A similar study was conducted to compare using different wavefront properties (phase, speckle displacement, and local radius of curvature) as the target. Like the results observed for the NN-based model in Fig. 4, the use of the speckle displacement provided the best performance and was thus chosen for all the subsequent studies.

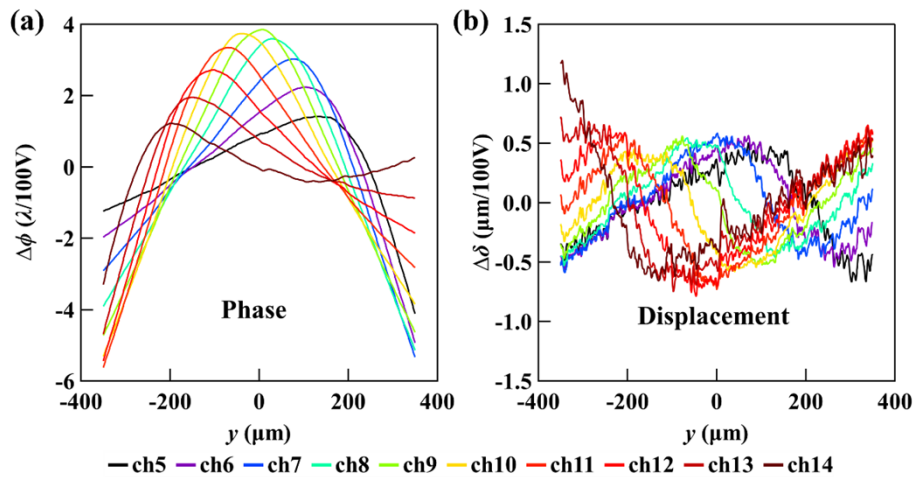


Fig. 8. Measured response functions in terms of (a) wavefront phase and (b) speckle displacement changes per 100 V changes on each piezo electrode, to be used by the linear controller.

The linear controller deals with the relative wavefront changes, as opposed to the NN-based model, which works on the absolute wavefront. The initial wavefront at t_0 was first measured, and the difference between the measured wavefront property (in terms of speckle displacement) and the desired spherical wavefront was taken as the target for tuning the electrode voltages. The process was carried out using a procedure described in [6] based on linear matrix operations. The output is an array of relative voltage changes for the 10 used electrode channels. Once the voltages were applied, the new wavefront was then measured at the same t_1 to t_{10} and t_{fin} as in the NN-based model studies.

The results using the NN-based (including all results from $n = 4, 5,$ and 10 models in Fig. 6) and linear controllers are compared in Fig. 9 for t_1 and t_{fin} . The NN-based controller provides a significant improvement in the average phase error and the sigma value (\pm values in the figure) of the error distribution. The NN-based model gives a control error of 0.262λ at t_1 , which is a factor of 2.8 better than the linear model, as shown in Fig. 9(a). The rms error at t_{fin} is 0.303λ for the NN-based model, a factor of 2.1 better than the linear model, as shown in Fig. 9(b). The total error over all time steps favors the NN-based model by a factor of 2.4. The NN-based model is generally more accurate and achieves the target wavefront in fewer steps. Also, the rms phase errors of all experiments fall into a narrow distribution (small $\pm\sigma$ values in the figure). This indicates that the NN-based model is more reliable in providing correct solutions over a broader range of wavefronts with different radii of curvatures. On the other hand, the linear model is only accurate when the required changes are small [13]. Moreover, since the response function of the

linear model was acquired at static conditions (waited one minute after each voltage change for the mirror shape to settle), it failed to account for the short-time dynamics of the mirror response.




Fig. 9. RMS phase error distribution obtained by the NN-based controller ($n = 4, 5,$ and 10) compared to a linear model (LM) based on response functions at (a) t_1 and (b) t_{fin} . The total error was defined as the average value of the 11 time steps (t_1 to t_{10} , and t_{fin}). Average rms error values are reported in the insets.

5.5. Adding the bender to the ML model

In all previous experiments, the global bender actuator (ch20) was not changed but imposed a fixed baseline wavefront phase with a fixed radius of curvature. Here, ch20 is included in the NN-based model, with limited movements around the initial position (170 V). We aimed to reduce the total excursion of the individual actuators (ch5–ch14) to achieve the desired shape. A new set of training data was collected with the following pattern:

- 10,000 wavefronts corresponding to a random set of voltages, changing every step.
- 10,000 wavefronts corresponding to a random set of voltages, changing every 5 steps.

While the surface actuators (ch5–ch14) were allowed to fluctuate within a range of $[-100 \text{ V}, 100 \text{ V}]$, the range for the bender actuator (ch20) was limited to $[-10 \text{ V}, 10 \text{ V}]$, a fluctuation compatible with the desired range of radius of curvature in the previous experiments. The model performance with the bender actuator is shown in Fig. 10. There is a slight improvement in the overall prediction error (0.232λ) as compared to that (0.265λ) of the model without the bender actuator [Fig. 6(a)]. The improvement ($\sim 12\%$) is not significant enough to conclude that adding the bender to the NN-based model is beneficial, but enough to promote it as a direction worth exploring in the future.

It is also important to underline the preliminary, albeit important, characteristic of this study. The bender actuator is different from the surface actuators, with different dynamic responses and weights in modifying the wavefront. The model training must be carried out with care to consider the optimal and limiting range of voltage fluctuations for the bender. Nevertheless, the results suggest the flexibility and possibility of the NN-based model to accommodate actuators with different forms and dynamics, making it suitable for more complicated adaptive optical systems.

5.6. Robustness and Repeatability

To test the robustness of the system, the experiments were repeated in two additional experimental sessions without retraining the model but using the same model trained in the first experimental




Fig. 10. RMS phase error distribution obtained by the NN-based controller ($n = 4$) with the bender actuator (ch20). The remaining parameters used in the NN-based model are the same as those used in Fig. 6(a). Average rms error values are reported in the insets.

session (described in section 5.2). In the two subsequent experimental sessions, separated by 1 and 3 months from the first, the entire optical system was rebuilt from scratch, including tuning the undulator source, white beam mirrors, slits, and the monochromator to the same condition, repositioning the KB mirrors, motor stages, supporting tables, and the wavefront sensor, and realigning the KB mirrors to provide the same focal position. As shown in Fig. 11, the NN-based model trained in the first experimental session provided progressively deteriorated but overall stable and consistent results. Considering the extreme conditions of the test, the repeatability and robustness of the controller model are more than satisfactory for our application without model retraining. This is significant for real-world operation, knowing that the model training is long-lasting even for beamlines with versatile configurations and applications.

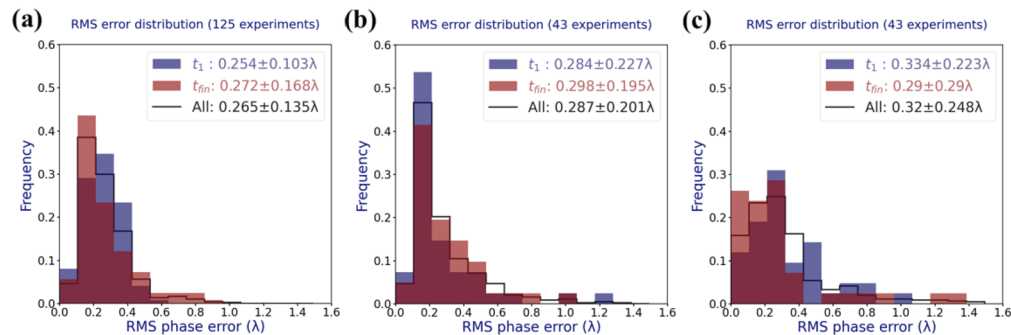


Fig. 11. Comparison between the prediction error of the NN-based model trained with data from the first experimental session but applied to (a) the first experimental session (a replot of Fig. 6(a) with a different scale for comparison), (b) the second experimental session after one month, and (c) the third experimental session after three months. Average rms error values are reported in the insets.

6. Conclusions

We reported the application of an NN-based prediction model and related controller for a bimorph mirror to control the X-ray beam wavefront. The goal is to achieve a high-quality wavefront control and to remove aberrations introduced by the photon transport system of the beamline. The system has been realized by combining two key developments: (i) a high-speed, reference-free,

single-shot wavefront sensor to directly measure the wavefront properties and feedback to the controller and (ii) an NN-based control system to drive the shape of the mirror surface and thus the output wavefront to the desired shape. The IDEA beamline, dedicated to optics characterization and testing, is also essential for such extensive studies over multiple beamtimes.

The best performance was achieved using the NN-based model trained on the speckle displacement (equivalent to the differential phase), compared with models trained on the absolute phase and local radius of curvature. This is a natural outcome of using the coded-mask-based wavefront sensing technique since it directly measures the speckle displacement. The speed of the wavefront sensor and data analysis algorithm also enabled the collection of large training datasets directly at the beamline. Further optimizing the NN-based model suggested that the $n = 4$ history step is sufficient for realizing the time-dependent dynamics of the system and thus predicting and controlling the wavefront to a sub-wavelength level.

The NN-based model achieved the best results in shaping the wavefront to a desired spatial distribution, especially a spherical wavefront (constant local radius of curvature), showing more than a factor of two improvements compared to the linear control model. Moreover, the NN-based model captured the non-linear dynamic properties of the adaptive mirror, which is essential for achieving real-time feedback and wavefront control for advanced applications.

The NN-based model corrected the wavefront phase to a sub-wavelength error level, equivalent to an rms height error of 1.8 nm on the mirror surface. This is within the accuracy required to achieve diffraction-limited performance in many applications of next-generation X-ray sources. The NN-based model also gave narrow error distributions, indicating reliability, when tested with arbitrary random wavefront shapes. Its merits of fast converging (seconds) and stability over time make it suitable for real-time feedback control of adaptive optics.

The implementation of a mixed bending actuator and small-scale surface actuators was tested in the training model. Further development will be needed to add boundaries to actuator movements in the training process to avoid potentially unsafe or unrealistic shape conditions and improve the training data.

A fundamental aspect of ML-driven systems is the stability and repeatability of the conditions under which the training data will be collected, and the optics operated, something that could not be taken for granted on a synchrotron beamline. We can speculate that an auto-alignment system based on artificial intelligence could be developed to monitor the beam properties (like position and shape) and restore desired initial conditions. Such a beamline auto-alignment system running in the background during the experiments could potentially be fundamental in keeping the stability and repeatability of the beam conditions from the ML training data collection to the operation of the NN-based controller, ensuring that the system is consistently capable of delivering optimal performance. Further improvements may be achieved by adding fine adjustments to the NN-based model. The model training speed may be accelerated by using higher-power supercomputers.

In addition to the NN model improvement, further developments in wavefront sensing accuracy and data analysis speed are still needed for real-time wavefront control with sub-second speed and picometer accuracy. The absolute wavefront sensing accuracy is sensitive to the aberration in the detector and the coded mask, which can be improved by calibration with a known wavefront shape, such as a spherical wavefront generated by a small pinhole. The recently developed deep-learning-based model [20] shows a promising first step toward high-resolution sub-second analysis.

It is important to underline that the NN-based model has been developed without being tailored and bonded to a specific motion mechanism or to a specific kind of optical element. The system correlates wavefront shapes to corresponding hardware setups represented by a list of numbers. It can be applied to optical elements with any deformation mechanism provided by an arbitrary number and type of hardware devices.

Moreover, the developed NN-based system is not limited to correcting wavefront aberration but could also provide a “next-generation” wavefront engineering solution by driving the optics to precisely manipulate the wavefront phase towards the experimental needs. For example, it can introduce tailored spatial interactions between the illumination and the sample in imaging and scattering measurements, enabling new contrasts and higher sensitivity, typically accessible only through holographic techniques. Moreover, this development paves the way for the advanced control of more sophisticated adaptive optics, such as grating optics for multicolor wavefront engineering.

Funding. Department of Energy, Office of Science, Office of Basic Energy Sciences (DE-AC02-05CH11231, DE-AC02-06CH11357).

Acknowledgments. This work was supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under contract DE-AC02-06CH11357 and DE-AC02-05CH11231. The work was also supported by the Laboratory Directed Research and Development (LDRD) funding (LDRD 2023-0004) from Argonne National Laboratory, provided by the Director, Office of Science, of the US Department of Energy under Contract No. DE-AC02-06CH11357. Work by Matthew J. Highland on experiment design and data collection was supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, Materials Science and Engineering Division. We warmly thank Dr. Gautam Gunjala for many insightful discussions related to this research work.

Disclosures. The authors declare no conflicts of interest in this paper.

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. M. Yabashi, K. Tono, H. Mimura, S. Matsuyama, K. Yamauchi, T. Tanaka, H. Tanaka, K. Tamasaku, H. Ohashi, S. Goto, and T. Ishikawa, “Optics for coherent X-ray applications,” *J. Synchrotron Radiat.* **21**(5), 976–985 (2014).
2. A. Rack, T. Weitkamp, L. Assoufid, T. Rack, I. Zanette, C. Morawe, R. Kluender, and C. David, “Protocol to study wavefront preservation capabilities of reflective X-ray optics with coherent synchrotron light,” *Nucl. Instrum. Methods Phys. Res., Sect. A* **710**, 101–105 (2013).
3. A. Zozulya, A. Shabalin, H. Schulte-Schrepping, J. Heuer, M. Spiwek, I. Sergueev, I. Besedin, I. Vartanians, and M. Sprung, “Wavefront preserving channel-cut optics for coherent x-ray scattering experiments at the P10 beamline at PETRAIII,” *J. Phys.: Conf. Ser.* **499**, 012003 (2014).
4. D. Cocco, G. Cutler, M. Sanchez del Rio, L. Rebuffi, X. Shi, and K. Yamauchi, “Wavefront preserving X-ray optics for Synchrotron and Free Electron Laser photon beam transport systems,” *Phys. Rep.* **974**, 1–40 (2022).
5. I.-T. Nistea, S. G. Alcock, V. Badami, R. Signorato, and K. Sawhney, “Controlling an active bimorph deformable mirror with sub-nanometre resolution,” *Proc. SPIE* **11109**, 111090E (2019).
6. X. Shi, Z. Qiao, S. Mashrafi, R. Harder, D. Shu, M. Wyman, J. Anton, S. Kearney, L. Rebuffi, T. Mooney, J. Qian, B. Shi, S. Matsuyama, K. Yamauchi, and L. Assoufid, “Prototype design and experimental tests of a zoom mirror system for the APS upgrade,” *Proc. SPIE* **11491**, 1149110 (2020).
7. S. Mashrafi, R. Harder, X. Shi, D. Shu, Z. Qiao, M. Wyman, T. Mooney, J. Anton, S. Kearney, L. Rebuffi, J. Qian, B. Shi, and L. Assoufid, “Machine learning control of an elliptically bent hard X-ray mirror,” *Proc. SPIE* **11491**, 114910T (2020).
8. S. G. Alcock, I. Nistea, J. P. Sutter, K. Sawhney, J. J. Fermé, C. Thellier, and L. Peverini, “Characterization of a Next-Generation Piezo Bimorph X-Ray Mirror for Synchrotron Beamlines,” *J. Synchrotron Radiat.* **22**(1), 10–15 (2015).
9. S. G. Alcock, I. T. Nistea, R. Signorato, and K. Sawhney, “Dynamic adaptive X-ray optics. Part I. Time-resolved optical metrology investigation of the bending behaviour of piezoelectric bimorph deformable X-ray mirrors,” *J. Synchrotron Radiat.* **26**(1), 36–44 (2019).
10. S. G. Alcock, I. T. Nistea, R. Signorato, R. L. Owen, D. Axford, J. P. Sutter, A. Foster, and K. Sawhney, “Dynamic adaptive X-ray optics. Part II. High-speed piezoelectric bimorph deformable Kirkpatrick–Baez mirrors for rapid variation of the 2D size and shape of X-ray beams,” *J. Synchrotron Radiat.* **26**(1), 45–51 (2019).
11. D. Cocco, C. Hardin, D. Morton, L. Lee, M. L. Ng, L. Zhang, L. Assoufid, W. Grizolli, X. Shi, D. A. Walko, G. Cutler, K. A. Goldberg, A. Wojdyla, M. Idir, L. Huang, and G. Dovillaire, “Adaptive shape control of wavefront-preserving X-ray mirrors with active cooling and heating,” *Opt. Express* **28**(13), 19242–19254 (2020).
12. H. Wang, K. Sawhney, S. Berujon, J. Sutter, S. G. Alcock, U. Wagner, and C. Rau, “Fast optimization of a bimorph mirror using x-ray grating interferometry,” *Opt. Lett.* **39**(8), 2518–2521 (2014).
13. G. Gunjala, A. Wojdyla, K. A. Goldberg, Z. Qiao, X. Shi, L. Assoufid, and L. Waller, “Data-driven modeling and control of an x-ray bimorph adaptive mirror,” *J. Synchrotron Radiat.* **30**(1), 57–64 (2023).
14. S. G. Alcock, I.-T. Nistea, V. G. Badami, R. Signorato, M. Fusco, L. Hu, H. Wang, and K. Sawhney, “Fast shaping control of X-ray beams using a closed-loop adaptive bimorph deformable mirror,” *Optica* **10**(2), 172 (2023).

15. Z. Qiao, X. Shi, M. J. Wojcik, L. Rebuffi, and L. Assoufid, "Single-shot x-ray phase-contrast and dark-field imaging based on coded binary phase mask," *Appl. Phys. Lett.* **119**(1), 011105 (2021).
16. Z. Qiao, X. Shi, R. Celestre, and L. Assoufid, "Wavelet-transform-based speckle vector tracking method for X-ray phase imaging," *Opt. Express* **28**(22), 33053 (2020).
17. S. Berujon, R. Cojocaru, P. Piau, R. Celestre, T. Roth, R. Barrett, and E. Ziegler, "X-Ray Optics and Beam Characterization Using Random Modulation: Theory," *J. Synchrotron Radiat.* **27**(2), 284–292 (2020).
18. S. Berujon, R. Cojocaru, P. Piau, R. Celestre, T. Roth, R. Barrett, and E. Ziegler, "X-Ray Optics and Beam Characterization Using Random Modulation: Experiments," *J. Synchrotron Radiat.* **27**(2), 293–304 (2020).
19. D. M. Paganin, H. Labriet, E. Brun, and S. Berujon, "Single-Image Geometric-Flow x-Ray Speckle Tracking," *Phys. Rev. A* **98**(5), 053813 (2018).
20. Z. Qiao, X. Shi, Y. Yao, M. J. Wojcik, L. Rebuffi, M. J. Cherukara, and L. Assoufid, "Real-Time X-Ray Phase-Contrast Imaging Using SPINNet—a Speckle-Based Phase-Contrast Imaging Neural Network," *Optica* **9**(4), 391 (2022).
21. C. M. Bishop, "Pattern Recognition and Machine Learning (Information Science and Statistics)," Springer-Verlag: Berlin, Heidelberg (2006).
22. K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *arXiv*, arXiv:1512.03385 (2015).
23. D. P. Kingma and J. Ba, "Adam: a method for stochastic optimization," *arXiv*, arXiv:1412.6980 (2014).
24. A. Paszke, S. Gross, and F. Massa, *et al.*, "PyTorch: An Imperative Style, High-Performance Deep Learning Library," *Advances in Neural Information Processing Systems* **32** (2019), Curran Associates, Inc., pp. 8024–8035. Available at: <http://papers.nips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
25. W. Li and E. Todorov, "Iterative Linear Quadratic Regulator Design for Non-linear Biological Movement Systems," (2004). Available at: <https://homes.cs.washington.edu/~todorov/papers/LiCINCO04.pdf>.
26. D. Shu, V. De Andrade, J. Anton, S. Kearney, K. Fezzaa, S. Bean, A. Deriy, W. Liu, J. Maser, B. Lai, J. Z. Tischler, and F. De Carlo, "Mechanical design of a flexural nanopositioning stage system for hard x-ray nanofocusing at the Advanced Photon Source 32-ID-C station," *Proc. SPIE* **11112**, 111120N (2019).