To be published in Optics Express:

Title:	Fast, efficient piston correction of deployable space telescopes using machine learning
Authors:	Daniel Martin,cyril Bourgenot,Andrew Reeves,Hubert P H Shum
Accepted:	21 January 26
Posted	21 January 26
DOI:	https://doi.org/10.1364/OE.582377

Fast, efficient piston correction of deployable space telescopes using machine learning

DANIEL MARTIN,[1,*] CYRIL BOURGENOT,[1], ANDREW REEVES [1] AND HUBERT. P. H. SHUM,[2]

[1] *Centre for Advanced Instrumentation, Durham University, UK,* [2] *Department of Computer Science, Durham University, UK*

[*] *daniel.martin2@durham.ac.uk*

Abstract: Alignment of segmented telescopes is essential in creating an effective monolithic primary mirror. A fast, two-step piston sensing technique has been developed to enable diffraction-limited imaging. Achieving such performance requires a Strehl ratio exceeding 0.8, which corresponds to a wavefront RMS error below 32 nm at $\lambda = 450$ nm.

A machine learning model has been implemented for a four-petal telescope to retrieve piston errors directly from PSF images and enable mirror correction. When tested on synthetic misalignments drawn from a uniform distribution within ±300 nm (or $\pm 2\lambda/3$), the model improved the mean Strehl ratio from a degraded state to 0.95 after one iteration, and to 0.99 after a second iteration. An SNR greater than 40 was found sufficient to achieve phasing corresponding to a Strehl ratio of at least 0.97.

Impact Statement

This study focuses on deployable space optics where satellites have a segmented primary mirror. It implements new machine learning architecture consisting of only 12,295 parameters to phase piston errors on each segment by analysing PSF images. The model is eight times smaller than recent developments [1] yet still predicts misalignments and corrects them to produce diffraction-limited PSFs. Being small and compact this procedure is applicable to cubesat and smallsat infrastructure, requiring little power or time resources which are valuable for such small payloads. This allows repeat alignment when needed, ensuring alignment to sub λ precision for high angular resolution. Also, by aligning in one motion, demands on servo actuators is reduced, extending their lifetime and the length of the functionality of the satellite as a whole. This innovative machine learning architecture supports the novel designs of deployable space optics, allowing more accessible and cost effective usage of satellites for uses in Earth observation, communications, and defence to name a few applications. Such novelty and impact reflects the key attributes of Optics Express.

1. Introduction

Deployable optics are increasingly driving the development of large, segmented space telescopes. Although the capability of launch vehicles continues to increase [2], the deployment of large monolithic mirrors remains impractical due to mass, volume, and structural constraints. Segmented primary mirrors consisting of smaller components provide a scalable solution to configure large scale mirrors in space. The motivation of increasing mirror size lies within the fundamental relationship between diameter and resolving power. In diffraction-limited systems, the angular resolution is inversely proportional to the pupil size, which is commonly defined by the diameter of the primary mirror. Consequently, larger aperture enable higher resolution imaging and improved overall performance. High resolution imaging is critical for applications ranging from exoplanet analysis, to Earth observation (EO), where spatial resolution directly determines the achievable scientific return. While deployable optics have traditionally been associated with flagship missions such as the James Webb Space Telescope (JWST) [3], there is

growing interest in extending these technologies to smaller-scale payloads. Notable applications are emerging with Cubesats [4–6], and compact satellites in the 1 – 2m class. To overcome stringent stowage and volume limitations of launcher fairings, current efforts [7–12] focus on the development of innovative deployable mechanisms. Concepts such as ADOT [7] and Supersharp [12] are designed to comply with launch vehicle constraints in their stowed configurations and subsequently deploy to achieve full aperture on orbit.

Exploiting the benefits of segmented mirrors requires the precise alignment of all segments to within sub-wavelength accuracy. Misalignments arising from mechanical tolerances, thermal drift, or launch-induced stress manifest as optical aberrations. Axial displacements and tilts introduce piston, tip and tilt (PTT) errors on each segment, resulting in a wavefront error (WFE). Coarse segments alignment enables an initial correction of the PTT terms by calibrating the telescope using through focus scanning and the overlap of each segment's point spread function (PSF) [7, 13]. Achieving diffraction-limited performance, however, requires extending this accuracy to sub-micron levels, which can be evaluated through wavefront sensing [14–16] and refined alignment via phase retrieval [17, 18]. Yet, on small satellites, dedicated wavefront sensors are often impractical due to size and power constraints. These hardware limitations have motivated the development of alternative, data-driven approaches, notably using neural networks [1, 7, 19, 20]. Training machine learning (ML) algorithms to analyse PSFs and infer the underlying aberrations [21] provides an effective means to reduce WFEs. Such models can be trained pre-flight, minimising the need for on-board computation, and offer faster, less resource-intensive alternatives to iterative correction methods - an advantage particularly valuable for small satellites where power and processing time are limited.

This study investigates the use of a CNN to analyse PSF images and estimate the corresponding aberrations. Previous work, both on ground-based infrastructures [22, 23] and general segmented systems [24], has demonstrated the potential of ML techniques for co-phasing segmented telescopes. These studies typically employ large, deep networks containing up to millions of parameters [25]. Even smaller-scale implementations with around 10^5 parameters [1] rely on complex architectures and controlled environmental conditions (e.g., limited noise).

In contrast, this work develops a compact model using Keras [26] and TensorFlow, departing from the deep network architectures commonly employed in similar studies [1, 19, 22]. While the overall computational power available for space applications is increasing, a "light", compact model frees up compute capacity for other required processes. This work also builds on the conclusions of [1] by investigating the use of both a defocus channel and two wavelength channels to break degeneracy (where similar PSFs are produced by different misalignments) and to correct larger aberrations. The main motivation in using these techniques is to enable correct phasing in one (or two) iterations thereby saving alignment time. By combining targeted data preprocessing with a compact, well-designed architecture, comparable—or even superior—correction accuracy can be achieved with a substantially smaller model. The resulting network, containing only $\sim$ 12,000 parameters, effectively restores diffraction-limited performance while remaining efficient enough for deployment in resource-constrained environments.

A generic segmented telescope model is considered, with each segment controllable in piston, tip, and tilt. Section 2 describes the telescope model, PSF simulation process, and data generation pipeline. The CNN architecture is also introduced, highlighting key design choices and their benefits. The results of model testing are discussed in Section 3, followed by concluding remarks and potential directions for future work in Section 4.

2. Simulations and modelling

2.1. Telescope Model

Effective modelling and simulation of segmented telescopes are essential to assess the impact of segment misalignments.

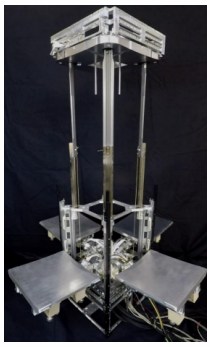

(a) Supersharp Swan prototype [12]

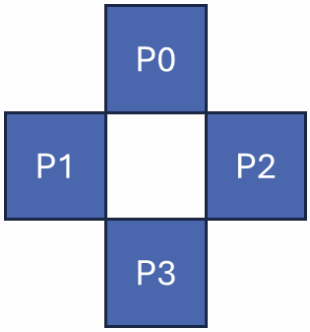

(b) Segment diagram

Fig. 1. (a) A deployable telescope cubeSat (Supersharp Swan prototype). Four deployable segments unfold with an extending boom for M2. The pupil function modelled from the prototype is shown in (b). Each active segment is labelled for easier reference.

A structure similar to the Swan prototype shown in Fig. 1a is considered, consisting of four deployable petals unfolding from the payload sides and an extruded telescopic secondary mirror. A detector is positioned in the focal plane between the petals. The petals are represented by an aperture defined by a two-dimensional pupil function, as illustrated in Fig. 1b. In the digital telescope model, incident light from a source at infinity is collected by the aperture and focused on-axis, onto the detector. The system is assumed perfectly stigmatic and modelled as a single mirror without a secondary. Field aberrations are neglected, and off-axis displacements are represented by a pure wavefront tilt. Under these far-field conditions, the digital telescope follows the Fraunhofer diffraction regime [27]:

$$\mathrm{PSF}(x', y') \propto |\mathcal{F}\{A(x, y)\, e^{i\phi(x,y)}\}|^2 \quad \text{where} \quad A(x, y) = A_0 P(x, y) \tag{1}$$

x', y' and x, y denote the coordinates in the focal and pupil planes, respectively. $\mathcal{F}$ represents a two-dimensional Fourier transform, ϕ the phase error across the aperture, and A the amplitude, which depends on the nominal amplitude A_0 and the pupil function P. Eq. 1 is used to obtain PSF images from the aperture function A with a constant pupil function P.

To investigate the effects of petal misalignments, piston, tip, and tilt aberrations are introduced as phase errors in the complex electric field, allowing the direct impact of phase on the PSF to be evaluated. From the simulated PSF, Strehl ratio is derived as a quantitative metric of optical quality, directly related to the near-diffraction-limited performance relevant to this application.

Each petal is modelled as 100 mm × 100 mm, giving the aperture a total diameter of 300 mm, which represents a realistic physical scale consistent with common CubeSat configurations. The PSF is computed for a system in which the ratio of the diffraction full width half maximum (FWHM) to detector pixel size is 2, i.e. Q = 2 corresponding to 2 pixels per FWHM. This is achieved by sampling the aperture (30 pixels) and padding (60 pixels) to obtain the desired sampling in the focal plane, which is twice the final detector sampling. Finally, the PSF is binned to match the correct detector pixel sampling.

Figure 2a illustrates the aperture structure under investigation. The nominal PSF, corresponding to a Strehl ratio of 1, is shown in Fig. 2b. The individual effects of each aberration are presented in Fig. 2c, where tip and tilt errors induce lateral translations of the sub-PSFs (i.e. the PSFs of individual segments) across the focal plane. These displacements can be corrected through

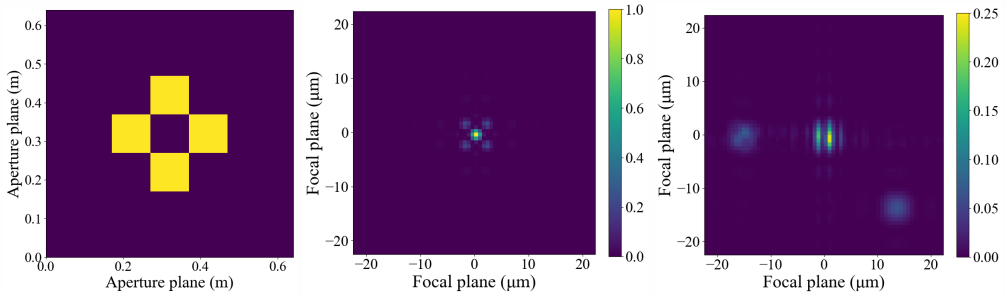

(a) Four segment aperture (b) Aligned optimal PSF (c) Misaligned PSF due to PTT

Fig. 2. Simulation of the four petal telescope as an aperture function. (a) shows the physical aperture plane of the segments. For no incident phase aberrations, the PSF remains optimal (b), with (c) showing the effects of misaligned segments through translation and interferences of sub-PSFs. Colour bars indicate maximum intensity, normalised to the nominal PSF in (b). $\lambda = 450\ nm$.

independent segment motions to realign and overlap the sub-PSFs, thereby merging the four sub-PSFs into a single one. Relative piston errors produce the most significant degradation, distorting the image and generating fringe patterns due to coherent interference between segments. Proper co-phasing of all segments requires precise overlap of the sub-PSFs and complete elimination of piston errors.

Standard alignment procedures are typically classified as 'coarse' or 'fine', with coarse alignment addressing piston errors on the order of $\sim \lambda/2$ RMS or larger, and fine alignment targeting residual piston errors in the range of $\lambda/10$–$\lambda/20$ RMS. In this study, a pre-alignment step is assumed, during which severely degraded sub-PSFs are first brought to focus through piston adjustment and subsequently overlapped by nulling tip and tilt aberrations. From this stage, only piston errors are corrected using a two-step procedure: an initial coarse correction that reduces the WFE to better than $\sim \pm\lambda/10$ RMS, followed by a fine correction phase that achieves sub-diffraction-limited performance.

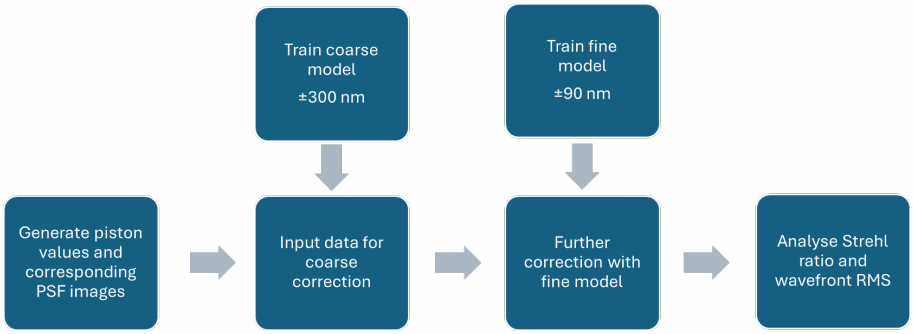

Fig. 3. Schematic flowchart of model implementation. A coarse and fine model will be separately trained, then used on randomly generated piston and PSF data. Resulting metrics will be analysed for performance checks.

The flowchart in Fig. 3 illustrates the procedure followed in this study. The training of coarse and fine models is integrated into a two-step correction framework. Random piston aberrations and their corresponding PSF images are generated as input data for the two ML models, which are trained sequentially, first on larger piston values of ±300 nm, followed by finer corrections

within ±90 nm. The 300 nm, corresponding to approximately two-thirds of the investigated wavelength, spans, when investigated on positive and negative values, over one full wave and enables the analysis of large piston errors. The fine alignment range allows examination of more subtle residual errors.

The performance of the ML models is evaluated based on the quality of the PSF correction. Once piston values are predicted, their deviations from the true values are computed. Two quality metrics are employed to assess these residuals and their effect on the PSF. The Strehl ratio, $\mathcal{S}$, provides a direct measure of performance relative to the nominal PSF, with a quality threshold of $\mathcal{S} > 0.8$ ensuring near-diffraction-limited performance according to the Maréchal criterion [28]. In addition, the RMS WFE is used as a complementary metric to quantify alignment accuracy.

Noise is further introduced into the system to enhance data realism and bridge the domain gap between synthetic and experimental data, thereby improving generalisation to unseen real-world conditions. This study considers both photon noise and readout noise, where the signal-to-noise ratio (SNR) is defined as:

$$\mathrm{SNR} = \frac{f}{\sqrt{f + \sigma_{ron}}} \tag{2}$$

The SNR depends on the photon flux, f, and the readout noise, σ_{ron}. The readout noise, taken from [1], has an average value of 2.1 e^-, which is typical of a CMOS detector. Noisy PSF data are generated following the same procedure as before, using Eq. 1. Photon and readout noise are then applied using Poisson and Gaussian distributions, respectively. Assuming constant readout noise, Eq. 2 is used to compute the corresponding photon flux f, where approximately 100 photons yield an SNR of 10, and 10,000 photons correspond to an SNR of 100.

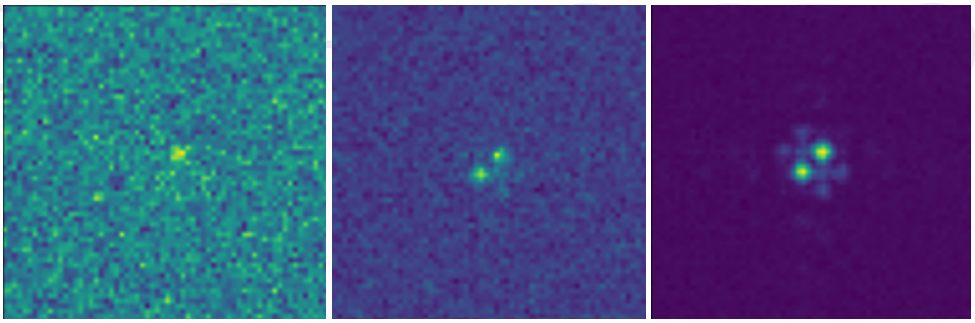

(a) SNR = 10 (b) SNR = 30 (c) SNR = 100

Fig. 4. Generated PSF images for a random configuration of piston values. Going from (a) - (c) the SNR increases from 10 – 100.

The effects of noise on the PSFs are shown in Fig. 4, illustrating SNRs of 10, 30, and 100. For each configuration of piston values, different levels of noise were applied. The SNR magnitude was varied by adjusting the photon flux. At low SNR values, important features are obscured, making it difficult for the ML model to distinguish them. As the SNR increases, distinctive features become more apparent, leading to higher model recognition accuracy. High SNR values, such as in Fig. 4c, are comparable to noiseless PSFs and retain the key global structures necessary for model analysis.

2.2. Degeneracies

For effective training of the ML model, large sets of input data are required. Careful data evaluation and selection are necessary to remove sources of degeneracy. Firstly, due to the

inherent shape and symmetry of the four-petal design, off-axis pairs can produce identical or nearly indistinguishable PSFs. For example, piston values applied to $P0$ and $P1$ (as shown in Fig. 1b) generate the same PSF as if the same pistons were applied to $P2$ and $P3$. Degeneracy also arises from pistons applied on opposite segments: a piston x (sub-λ) on one segment produces the same PSF as $-x$ on the opposite segment. ML models rely on observing distinct patterns to learn the mapping between input and output; a direct and unique (bijective) relationship is therefore essential. When similar PSF images can be generated by multiple piston combinations, a direct 1-to-1 mapping of piston values to images becomes impossible. Additionally, the λ ambiguity imposed by the exponential term in Eq. 1 introduces further degeneracy, as phase cycles of multiples of $\pm 2\pi$ (or $\pm\lambda$) produce identical monochromatic PSFs.

Consequently, several methods are implemented to resolve these issues and ensure a direct piston-to-PSF mapping. Previous works have addressed these challenges by enforcing asymmetry, using an odd number of segments, cropping petals, or limiting aberration ranges [1, 29, 30]. However, in the present study, alternative solutions are investigated to accommodate symmetrical structures:

- **Diagonal symmetry** - To remove ambiguity caused by the symmetrical segment structure, a reference segment is selected. By setting $P0$ as the reference segment in the training and only correcting the remaining three petals, diagonal pairs are eliminated. Since only relative piston differences are required for phase correction, aligning to this new reference preserves wavefront information while simplifying the learning task.
- **Opposite petals** - Even with a reference segment set as $P0$, the left and right petals can still produce degenerate PSFs. Introducing an additional channel by adding a known defocus to the incoming wavefront helps distinguish the signs. This can be achieved by moving the secondary mirror or detector and recording a second PSF image.
- 2π **ambiguity** - Adding a second wavelength removes this limitation, ensuring that at least one wavelength is out of phase with the other unless perfect phasing is achieved. For clear feature differentiation and easier generalisation, blue (450 nm) and red (650 nm) are chosen as two colour channels. Piston movements then correspond to different phase shifts at each wavelength.

Fig. 5 presents examples from the dataset, showing differences in the blue and red PSF images. The first column displays two blue PSFs at different $P1$ values, with the other petals held constant at values specified in the figure caption. In the upper row, $P1$ is set to -270 nm, while in the lower row it is 180 nm. Varying by the blue wavelength of 450 nm, which corresponds to 2π, these images are identical. Only through the addition of the second wavelength can these cases be distinguished, thereby resolving the 2π ambiguity. For each piston combination, a defocus of 1.2 radians was applied to the incident wavefront to eliminate sign ambiguities, as illustrated in the upper-right panel for the blue PSF. While this introduces unique images, additional training data provide the model with even more information, enabling more effective learning.

One dataset consists of four PSF images: an in-focus and a defocus image for each of the two wavelengths considered. This method applies novel diversity through defocus and wavelength channels to correct large aberrations. This approach aligns with standard imaging and calibration methods that the satellite is expected to support. Defocus can be introduced through movement of M2 or the detector, or using multiple detectors, removing mechanical complexity and reducing image acquisition time. Broadband detectors and filters can be used to capture both colour channels required. A total of 20,000 sets are generated for each model, using piston values within $[-300, 300]$ nm for the coarse model and $[-90, 90]$ nm for the fine model. These ranges correspond to approximately $2\lambda/3$ and $\lambda/5$, where λ is chosen as the smallest wavelength investigated (blue, 450 nm), ensuring that the correction procedure also accommodates the larger

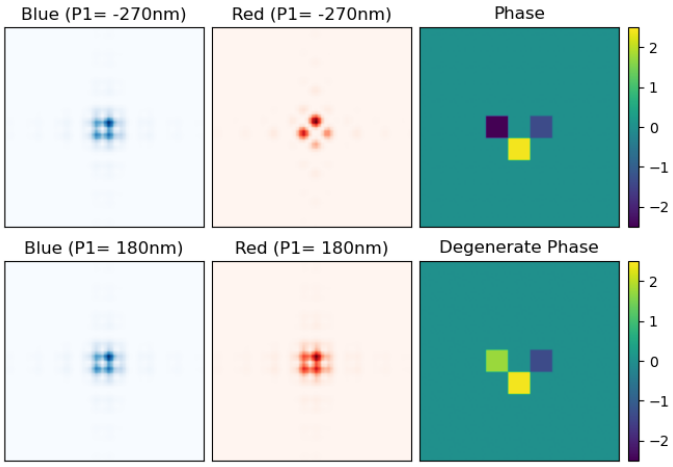

Fig. 5. Plots of PSF images in training data for Q = 2. Pistons are set as [0,-270,-140,250] nm, generated in blue (450 nm) and red (650 nm) wavelengths. Bottom row plots are generated from pistons of [0,180,-140,250] nm, i.e. 2π shift in P1 at 450 nm. The third column shows the phases induced by the different piston values.

wavelength. Tip and tilt values are set to zero. The defocus magnitude was selected to produce a noticeable difference between the images, making them distinguishable for the model. A value of $\lambda/5$ (corresponding to 1.2 radians at λ = 450 nm) was chosen, providing sufficient variation to break the 2π degeneracy. Fifteen percent of the training set was reserved for validation, and an additional 1,000 sets were generated at this stage for post-training testing.

2.3. Machine Learning Model

In this study, a model is trained using PSF images and their corresponding piston values, which serve as the inputs and outputs, respectively. The four images (in-focus and defocus images at two wavelengths) are provided as four 2-dimensional input channels, while the three piston values result in a final layer of 3 outputs.

For image analysis, convolutional layers are employed. Here, a 3×3 kernel is applied to each pixel in the input images, capturing relationships between neighbouring pixels and encoding patterns into the learned coefficients. The width of each convolutional layer, defined by the number of filters they contain, starts at 4 and increases sequentially to 8 and 16. This progressive increase in width enables hierarchical learning, where different levels of feature complexity can be captured. The first layer extracts local features and global symmetry, while subsequent layers identify higher-order features and more complex patterns. After feature extraction, the data passes through two fully connected layers before reaching the output layer.

Additional components are incorporated throughout the network layers to enhance performance. Rectified Linear Unit (ReLU) functions [31] are applied in each layer, introducing non-linearity and mitigating the vanishing gradient problem, thereby enabling deeper network architectures. Batch normalisation [32], applied after each convolution, normalises intermediate activations, improving training stability and acting as a form of regularisation to reduce overfitting. A max pooling layer [33], inserted between the convolutional and dense layers, reduces spatial dimensions while preserving the most salient features. These integrated elements collectively improve performance and contribute to a more compact and effective model.

Fig. 6 shows the full network architecture, with Table 1 providing detailed properties and parameters, totalling 12,295. Untrainable parameters correspond to constant weights, calculated

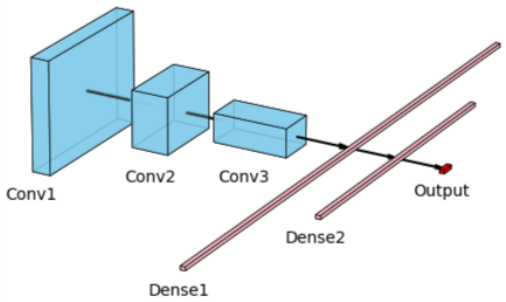

Fig. 6. Diagram of the ML model and process of training. Four PSFs are set as input, passing through three convolution layers (light blue) of increasing depth then two fully connected layers (pink), with a final dense output of N piston values (red). Parameter sum to a total of 12,295 of which 56 are untrainable and are part of the model infrastructure.

Layer	Size	Parameters
Input	(64, 64, **4**)	0
Convolution 1	(32, 32, **4**)	164
Convolution 2	(16, 16, **8**)	328
Convolution 3	(8, 8, **16**)	1232
Global Pooling	(**16**)	0
Dense 1	(**128**)	2176
Dense 2	(**64**)	8256
Output	(**3**)	195

Table 1. Summarised properties of each model layer. Their label and function is given, alongside the output shape of data passing through, and number of parameters associated. Quantities in bold represent the number of neurons - width of each layer.

from mean and variance rather than being updated via backpropagation.

Both the coarse and fine models were trained using the same training rate and batch size, resulting in a training time of approximately 254 seconds on a typical office laptop (see figure caption). Fig. 7 shows model accuracy versus training epoch for each model, with the corresponding validation (val) data also displayed. In Fig. 7a, the mean piston RMS is calculated over the entire sample after each epoch. Over time, the RMS decreases, plateauing after an initial period of rapid improvement.

For the fine model (green), smaller input residuals result in an inherently lower RMS throughout training. The fine model reaches a plateau even before epoch 10; an early stopping callback further reduces training time by halting learning once the validation loss ceases to deviate above a specified threshold. Model behaviour is further illustrated in Fig. 7b, which shows the increase in accuracy as training progresses. Notably, the validation loss and RMS are smaller than those for the training set, and correspondingly, validation accuracy is higher. This discrepancy may arise from several factors, most likely the regularisation effect of batch normalisation. During training, batch normalisation uses the mean and variance of each batch, introducing additional noise into the network's forward pass. In validation, a rolling mean computed from the training phase is used, avoiding this noise, causing the observed variation.

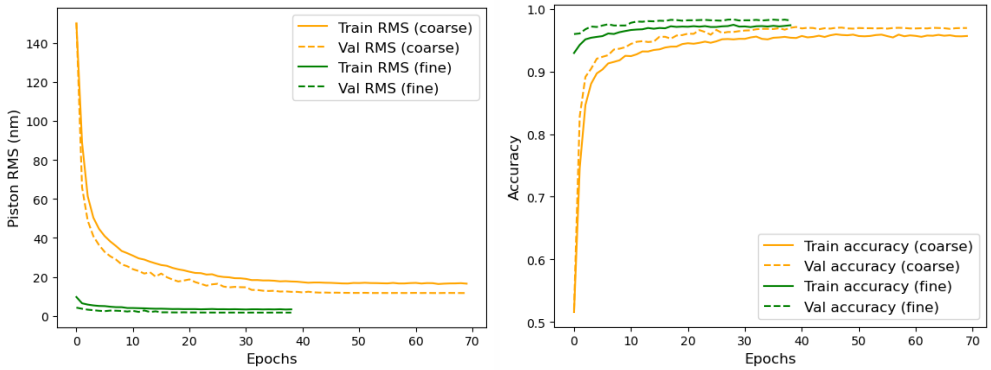

(a) Mean piston RMS during training

(b) Model accuracy during training

Fig. 7. Analysis of training with respect to piston RMS and accuracy. Both the training and validation metrics are plotted for the coarse (orange) and fine (green) models. (a) Plot of mean piston RMS of the entire data set against epochs. Accuracy (ratio of correct values) is shown in (b). For each model, 100 epochs are used with a batch size of 32, taking 253.74 s to train. Early stopping capped epochs to 70 and 44 for the coarse and fine respectively. Training was carried out on a laptop with an AMD Ryzen 5 PRO 7535U processor (6 cores, 12 threads, 29-4.55 Ghz) and 16 GB of DDR5 RAM.

3. Results

3.1. Noiseless Data

Fig. 8 summarises key performance metrics of the two-step phasing procedure. For 1,000 testing sets, model performance in terms of the Strehl ratio is shown in Fig. 8a and Fig. 8b, for the initially generated PSF images and those reconstructed after each step using the predicted pistons. The Strehl ratio of the blue channel is reported here, as pistons have the greatest effect on shorter wavelengths. Initially, the Strehl ratio ranges from 0.36 to 0.95, with only 8.7% of cases above the intended threshold of 0.8. The coarse model improves this to 95.2% above 0.8, increasing the mean Strehl ratio from 0.6 to 0.92. Similar improvements are observed in the red channel, increasing from 0.64 to 0.95. The reduction in standard deviation indicates higher precision throughout the correction process.

Introducing the residual piston error as input to the fine model further improves the Strehl ratio, with 99.9% of cases above the threshold and a mean ratio of 0.99 for both colour channels.

Fig. 8c shows the residual piston error for each petal in every test set, plotted against the input piston error of that petal. Results are shown after applying the coarse model alone and after both models are applied. The accuracy of the coarse model decreases at large piston values, particularly near the training boundary of 300 nm, where the model fails to correctly distinguish piston values. This is expected, as such errors were not encountered during training. Interestingly, the skew of residuals changes sign across the range, with negative pistons being over-predicted and positive pistons under-predicted. The scarcity of training data near these boundaries contributes to this reduced accuracy. Elsewhere, the spread of residuals remains fairly constant, with a standard deviation of 33.5 nm. Over the full range of pistons, the subsequent fine model reduces residuals to within ±10 nm, with a standard deviation of 3.7 nm. Some deviations reflect gaps in data and further training may improve accuracy further.

The consistent accuracy across the full range of trained piston values suggests that the model prioritises correcting the Strehl ratio over precisely predicting individual pistons. Sign ambiguities between pistons on $P1$ and $P2$ may still persist for piston values $< \lambda/14$, leading the model

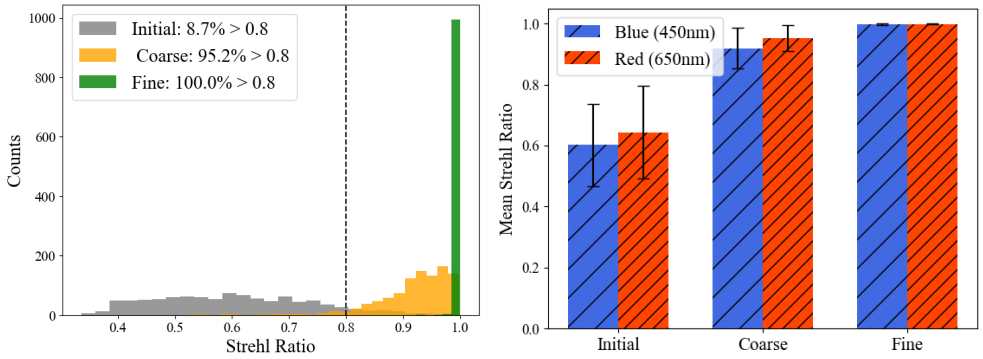

(a) PSF Strehl ratios

(b) Mean Strehl ratios

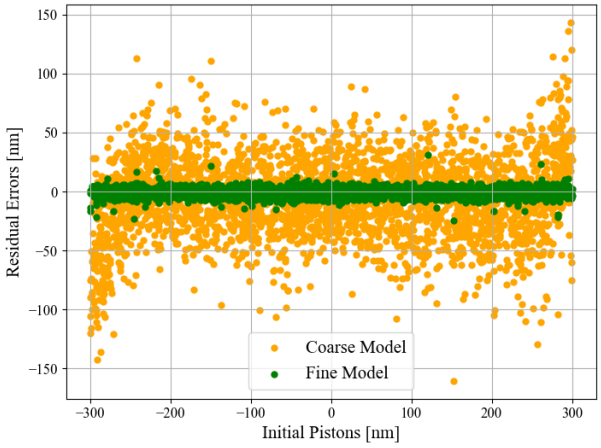

(c) Residual Piston Values

Fig. 8. Analysis of the two-step correction procedure through testing of 1,000 sets. (a) Histogram of Strehl ratios for blue (450 nm) before correction and after each step, noting the percentage above 0.8. This threshold is marked with a dashed line. The mean Strehl value for each step, for each wavelength, is shown in (b). Blue (/) and red (//) bars represent the respective colour channels and error bars show the standard deviation error on the mean. Physical residual piston errors are plotted in (c), with both coarse (orange) and fine (green) values against the initial piston error per mirror petal.

to predict $\pm x$, effectively averaging the relative pistons to zero. Alternatively, the model may struggle to distinguish pistons of such magnitude to high precision as aberrations $< \lambda/14$ have little effect upon the PSF images. This limited performance can be seen in the fine model accuracy of Fig. 7b.

3.2. Effect of Noise

Testing the original models with noisy data demonstrates that they are incapable of efficiently analysing such images. Even at an SNR of 100, the lowest mean RMS WFE remained approximately 150 nm after correction. Two new models were trained following the same approach as described in Section 3.1, with a coarse 300 nm model and a fine 90 nm model. 10,000 piston combinations were used to generate PSF images, to which photon and readout noise were applied prior to training. Varying the photon flux (f in Eq. 2) generated an individual training set of PSFs at each SNR level. The resulting models were then tested over an SNR range

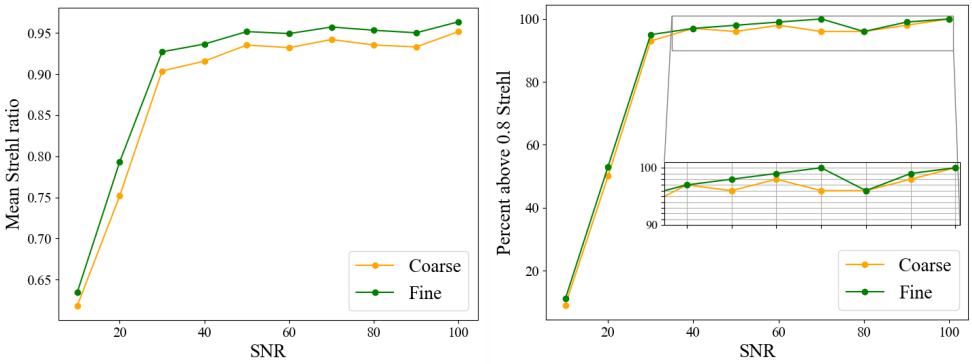

(a) Mean Strehl ratio

(b) Percentage of PSFs with > 0.8 Strehl

Fig. 9. Results from testing on noisy PSF images. (a) Plot of mean Strehl ratio for each step in the correction process as a function of SNR. Percentage of diffraction limited PSFs are shown in (b). The coarse (orange) and subsequent fine (green) model are compared for initial piston values ±300 nm, and fine phasing of ±90 nm. Inset highlights differences at higher percentages.

of 10–100 for 100 cases at each level.

Fig. 9 presents the results of applying the procedure to noisy datasets, including both photon and readout noise. The resulting mean Strehl ratios are shown in Fig. 9a, evaluated after coarse alignment and again after fine alignment. Performance is poor at low SNR (= 10), but improves with increasing signal, achieving diffraction-limited performance for both models at SNR ≥ 30. Notably, the coarse model plateaus shortly after reaching a Strehl ratio of 0.9, with a maximum of 0.95 at SNR = 100. This level of correction alone is sufficient for effective phasing. The fine model provides slightly higher accuracy, reaching a maximum mean Strehl ratio of 0.97. Thus, the fine model is not strictly necessary to achieve high Strehl ratios but can further refine the system with only one additional iteration, requiring minimal additional resources and time.

Fig. 9b illustrates the effective phasing by plotting the percentage of PSF images phased to the diffraction limit (i.e., > 0.8). For SNR ≥ 30, both models successfully phase 96–100% of cases to the diffraction limit. Both models accurately predict piston values with sufficient precision; the fine model is not strictly required but provides additional accuracy across all SNR levels.

4. Discussion

The results demonstrate that a two-cycle, CNN-based machine learning alignment model can achieve excellent piston-correction accuracy. Using the two-step process, a mean Strehl ratio of 0.96 can be achieved at SNR ≥ 100, with a failure rate of less than 0.01%. These results show the effectiveness of adding defocus and colour channels into the CNN models, developing the work of [**dumiont**].

For space telescope alignment applications, an ideal ML model must reach diffraction-limited correction while remaining compact enough for on-board storage. The memory and computational resources of space-qualified electronics are limited, making a model with fewer parameters advantageous. This is achieved with a model requiring significantly fewer internal parameters than previous approaches, including recent ResNet developments in [1] and standard ResNet architecture. The reduced architecture, with only 12, 295 parameters, enhances performance while minimizing memory and power requirements. Together, the coarse and fine models occupy only 436 KB.

Accurate alignment of deployable telescopes in space enables larger effective apertures,

directly improving angular resolution. High-resolution Earth observation, imaging, and space communications benefit from sub-λ aligned mirrors, which enable more detailed mapping, diffraction-limited performance, and higher efficiency. Applying this ML architecture and phasing strategy to segmented telescopes corrects misalignments to excellent accuracy.

While the coarse model alone achieves efficient alignment (RMS 10–20 nm, with 95.2% of cases phased to the diffraction limit), the fine model further refines performance to $\sim$ 1 nm RMS with minimal additional resource cost. Compared to alternative iterative techniques, requiring only two cycles relaxes the demands on servo actuators for segment movement. Fast convergence reduces mechanical wear and alignment time. With only a defocus image per colour channel, piston values can be determined and corrected in less than one second, conserving both power and time—critical resources aboard small satellites. This method can be repeated regularly when in contact with a ground beacon or other alignment source, ensuring precise alignment throughout the satellite's operational lifetime. This performance strengthens small deployable satellites as a scalable, accessible, and affordable technology. Pre-flight training time on Earth is also reduced due to the compact model size, allowing rapid re-training to improve performance or adapt the model for different payloads.

For deployment on production space telescopes, additional considerations are required, particularly regarding the fidelity of the digital twin relative to the actual telescope, which may include small manufacturing or polishing errors. Each payload should be considered individually, and a model trained for that specific geometry. Surface errors of each segment affect the PSF and thus the performance of the models; these errors should be measured and incorporated into training to ensure optimal results. While some errors can be simulated, unknown thermal deformations and material property variations in space present additional challenges. So far, the model incorporates a reference segment (segment with no alignment capability) to remove degeneracies caused by the symmetry of the 4-petal system. In a real system, each segment has its own residual surface figure error, leading to a unique sub-PSF. This naturally breaks the symmetry when the model is trained to account for these measured errors. In addition, incorporating the full telescope design, including the secondary mirror, will be necessary to capture the increased complexity of aberration propagation between optical surfaces. The method indicates that fast alignment is feasible, enabling adaptive and continuous correction. However, to achieve optimal performance, well-calibrated piston control for each segment is essential; otherwise, a closed-loop system is required. While demonstrated for piston values up to 300 nm, aberration magnitudes are payload-specific. Resulting piston values from simulations may be physically unattainable with existing hardware due to actuator resolution or thermal deformation, and thus the model's achievable accuracy may vary for each application.

While showing excellent accuracy in piston correction, this work contributes to the broader field of segmented-mirror alignment. In future work, it will be beneficial to investigate how using additional PSF channels is affected by the inclusion of residual tips and tilts. The CNN will be analysed and adapted accordingly to account for these additional errors, changing outputs and parameters as required. It is expected that the inclusion of further aberrations will only slightly increase the computational cost, as it relies on the same overall CNN architecture. Additionally, an interesting avenue for future study is the sampling of the PSF and an assessment of the model's limitations when correcting images from systems with different detector pixel scales.

5. Conclusion

A convolutional neural network has been demonstrated to effectively analyse and determine piston corrections for PSF images from a segmented telescope. PSF images are generated from four segments with relative pistons and processed in two colour channels, including the addition of a defocus image. Training of both coarse and fine alignment models demonstrated the ability to correct pistons to achieve a Strehl ratio > 0.8. The coarse alignment model, trained with

pistons in the range ±300 nm, phased PSFs to a mean Strehl ratio of 0.92. The subsequent fine alignment model further improved the correction, achieving a mean Strehl ratio of 0.99.

The effect of photon and readout noise was also investigated, providing insight into the impact of SNR on model performance when trained on noisy datasets. An SNR > 40 was sufficient for diffraction-limited phasing, with > 97% of cases reaching a mean Strehl ratio of greater than 0.93 after the two-step process. An SNR of 100 phases 99.9% of cases to a mean Strehl of 0.96.

Further improvements in data processing and model architecture could enhance the performance of the ML approach. Introducing small residual tip/tilt errors would better represent realistic operating conditions. Additionally, including surface errors across each petal could improve the fidelity of PSF simulations and reduce degeneracies in the training data, supporting more effective learning. Another avenue for investigation is PSF sampling and an assessment of the limitations of the model when correcting images from systems with smaller diffraction ratio Q.

Funding

Acknowledgement

The authors would like to thank Steven Knox, Surrey Satellite Technology Limited, for his valuable input of deployable optical systems.

Disclosures

The authors declare no conflicts of interest.

Data Availability Statement

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. Dumont, C. M. Correia, J.-F. Sauvage, N. Schwartz, M. Gray, and J. Cardoso, "Phasing segmented telescopes via deep learning methods: Application to a deployable cubesat," *J. Opt. Soc. Am. A*, vol. 41, pp. 489–499, 2024.

[2] E. Seedhouse, *SpaceX: Making Commercial Spaceflight a Reality*. Springer Science & Business Media, 2013.

[3] P. A. Lightsey, C. B. Atkinson, M. C. Clampin, and L. D. Feinberg, "James webb space telescope: Large deployable cryogenic telescope in space," *Opt. Eng.*, vol. 51, no. 1, p. 011 003, 2012.

[4] R. P. Welle, "Overview of cubesat technology," *Handb. Small Satell. Technol. Des. Manuf. Appl. Econ. Regul.*, pp. 1–17, 2020.

[5] D. Selva and D. Krejci, "A survey and assessment of the capabilities of cubesats for earth observation," *Acta Astronaut.*, vol. 74, no. 1, pp. 50–68, 2012.

[6] National Academies of Sciences, Engineering, and Medicine, *Achieving Science with CubeSats: Thinking Inside the Box*. Washington, DC: National Academies Press, 2016.

[7] N. Schwartz et al., "Phasing a small deployable optic space telescope using focal-plane wavefront sensing," in *Space Telescopes and Instrumentation 2024: Optical, Infrared, and Millimeter Wave*, 2024.

[8] D. Dolkins and H. Kuiper, "Design and end-to-end modelling of a deployable telescope," *Front. Astron. Space Sci.*, vol. 4, p. 13, 2017.

[9] N. Schwartz et al., "Active deployable primary mirrors on cubesat," in *Proceedings of 13th IAA Symposium on Small Satellites for Earth Observation*, 2021.

[10] I. Parry et al., "Unfolding, self-aligning thermal space telescopes for high resolution earth observations," in *International Workshop on High-Resolution Thermal EO*, 2023.

[11] D. Gooding et al., "A novel deployable telescope to facilitate a low-cost< 1m gsd video rapid-revisit small satellite constellation," in *International Conference on Space Optics—ICSO 2018*, SPIE, vol. 11180, 2019, pp. 102–110.

[12] I. Parry et al., "Innovative technologies for very-high-resolution MWIR and LWIR Earth observations," in *Proceedings of the Small Satellites Systems and Services Symposium (4S 2024)*, Conference Paper, 2025.

[13] M. D. Perrin et al., "Preparing for jwst wavefront sensing and control operations," in *Space Telescopes and Instrumentation 2016: Optical, Infrared, and Millimeter Wave*, vol. 9904, International Society for Optics and Photonics, 2016, 99040O.

[14] I. Surdej, N. Yaitskova, and F. Gonte, "On-sky performance of the zernike phase contrast sensor for the phasing of segmented telescopes," *Appl. optics*, vol. 49, no. 21, 2010.

[15] R. A. Gonsalves and R. Chidlaw, "Wavefront sensing by phase retrieval," in *Applications of Digital Image Processing III*, SPIE, vol. 207, 1979.

[16] M. R. Bolcar and J. R. Fienup, "Sub-aperture piston phase diversity for segmented and multi-aperture systems," *Appl. Opt.*, vol. 48, no. 1, A5–A12, 2009.

[17] S. Ragland and L. Gers, "A phase retrieval technique to measure and correct residual segment piston errors of large aperture optical telescopes," in *Ground-based and Airborne Telescopes IX*, International Society for Optics and Photonics, vol. 12182, SPIE, 2022.

[18] R. A. Gonsalves, "Phase Retrieval And Diversity In Adaptive Optics," *Opt. Eng.*, vol. 21, no. 5, 1982.

[19] Y. Wang et al., "Deep learning wavefront sensing for fine phasing of segmented mirrors," *Opt. Express*, vol. 29, no. 6, pp. 8745–8756, 2021.

[20] L. Sauniere, W. Gillard, and J. Zoubian, "Decoding optical aberrations of low-resolution instruments from psfs: Machine learning and zernike polynomials perspectives," in *Space Telescopes and Instrumentation 2024: Optical, Infrared, and Millimeter Wave*, SPIE, vol. 13092, 2024, pp. 1262–1274.

[21] D. Li, S. Xu, D. Wang, and D. Yan, "Large-scale piston error detection technology for segmented optical mirrors via convolutional neural networks," *Opt. Lett.*, vol. 44, 2019.

[22] F. Rossi, C. Plantet, A.-L. Cheffot, G. Agapito, E. Pinna, and S. Esposito, "Machine learning techniques for piston sensing," in *Proceedings of SPIE Vol. 12185, Adaptive Optics Systems VIII*, vol. 12185, 2022.

[23] J. R. P. Angel, P. Wizinowich, M. Lloyd-Hart, and D. Sandler, "Adaptive optics for array telescopes using neural-network techniques," *Nature*, vol. 348, 1990.

[24] J. Fang and D. Savransky, "Automated alignment of a reconfigurable optical system using focal-plane sensing and kalman filtering," *Appl. Opt.*, vol. 59, no. 31, pp. 9649–9659, 2020.

[25] K. He, X. Zhang, S. Ren, and J. Sun, *Deep residual learning for image recognition*, 2015.

[26] F. Chollet et al., *Keras*, https://keras.io, 2015.

[27] C. S. Adams and I. G. Hughes, *Optics f2f: From Fourier to Fresnel*. Oxford University Press, 2018.

[28] A. Maréchal, *Étude des effets combinés de la diffraction et des aberrations géométriques sur l'image d'un point lumineux*. Ed. de la Rev. d'Optique théoret. et instrumentale, 1948. [Online]. Available: https://books.google.co.uk/books?id=5tSmzwEACAAJ.

[29] B. Pope, N. Cvetojevic, A. Cheetham, F. Martinache, B. Norris, and P. Tuthill, "A demonstration of wavefront sensing and mirror phasing from the image domain," *Mon. Not. R. Astron. Soc.*, vol. 440, pp. 125–131, 2014.

[30] F. Martinache, "The asymmetric pupil fourier wavefront sensor," *Publ. Astron. Soc. Pac.*, vol. 125, no. 926, pp. 422–430, 2013.

[31] V. Nair and G. E. Hinton, "Rectified linear units improve restricted boltzmann machines," in *Proceedings of the 27th International Conference on Machine Learning (ICML)*, 2010, pp. 807–814.

[32] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proceedings of the International Conference on Machine Learning (ICML)*, 2015, pp. 448–456.

[33] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems*, vol. 25, 2012, pp. 1097–1105.