



HAL
open science

A Study of Network-based Wavefront Estimation with Noise

Jeffrey Peter Smith, Jesse Cranney, Charles Gretton, Damien Gratadour

► **To cite this version:**

Jeffrey Peter Smith, Jesse Cranney, Charles Gretton, Damien Gratadour. A Study of Network-based Wavefront Estimation with Noise. Adaptive Optics for Extremely Large Telescopes 7th Edition, ONERA, Jun 2023, Avignon, France. 10.13009/AO4ELT7-2023-084 . hal-04414739

HAL Id: hal-04414739

<https://hal.science/hal-04414739v1>

Submitted on 31 Jan 2024

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

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



A Study of Network-based Wavefront Estimation with Noise

Jeffrey Smith^{a,b}, Jesse Cranney^b, Charles Gretton^a, and Damien Gratadour^c

^aSchool of Computing, Australian National University, Australia

^bAdvanced Instrumentation Technology Centre, Research School of Astronomy and Astrophysics, Australian National University, Australia

^cLESIA, Observatoire de Paris, University PSL, CNRS, Sorbonne Université, Université de Paris, 5 place Jules Janssen, 92195 Meudon, France

ABSTRACT

We provide an extended COMPASS simulation study of our existing Machine Learning based approaches to wavefront estimation using Convolutional Neural Networks (CNN). These proceed by training a *UNet* to estimate the wavefront phase from a wavefront sensor image. That network has application in PSF-R workflows and AO control; both considered here. We make several extensions and novel contributions. First, evaluating historical networks we study the spacial frequency response of network-based wavefront reconstruction for networks trained under distinct loss regimes; with and without an adversarial loss. Second, we study network-based estimation with low magnitude guide stars, a setting where photon noise and read-out noise can impact performance. We train new networks that use Deformable Mirror (DM) shape along with WFS data as input, as we find including the DM shape to be broadly advantageous. We evaluate the impact of a low signal to noise ratio under a learning regime that employs an adversarial loss. Previous work found that learning regime to be best when producing networks for PSF-R. We find that noise substantially impacts training performance in this setting, and also the utility of such networks for PSF-R. On the other hand, training networks without adversarial loss yields robust and performant networks in the presence of noise. We evaluate such networks in a detailed control case study.

Keywords: Machine Learning, Wavefront Estimation, Noise, Control, PSF-R, Modal Analysis, cGAN, UNet

1. INTRODUCTION AND BACKGROUND

Image translation using a Convolutional Neural Network (CNN) provides an effective technology for obtaining accurate wavefront estimates. Such estimates feature high spatial frequency information, which has a non-linear response in the classical measurement space (i.e., slopes) of the Shack-Hartmann (SH-WFS) and Pyramid Wavefront Sensors (P-WFS). Irrespective of the type of WFS being used, incumbent AO concepts are typically operated in a linear regime, ignoring most or all of the non-linear component of the WFS response. CNN based estimation that is fast and sensitive to the non-linearities offers a promising direction worthy of detailed exploration.

Further author information: (Send correspondence to Jeffrey Smith) jeffrey.smith@anu.edu.au

Use cases for detailed wavefront estimation include: (i) real-time control of AO systems, and (ii) Point Spread Function Reconstruction (PSF-R). We have explored both in previous studies [1, 2, 3, 4]. In the approach we developed a UNet trained to perform wavefront estimation. That network is either trained to minimise a *reconstruction loss*, or otherwise can be trained to minimise a weighted sum of reconstruction and adversarial loss terms following the cGAN concept [5]. Training a network without adversarial loss is an order of magnitude faster. From hereon we find it convenient to write UNet to refer to a network trained using only a reconstruction loss, and write cGAN to refer to a UNet trained specifically in the adversarial setting. Using the COMPASS simulation tool [6], we previously demonstrate that a cGAN provides estimates that enable the recovery of a highly detailed PSF. On the other hand, the UNet is not able to produce a detailed and accurate halo. We previously also explored the use of cGAN estimates for control, and in particular in a setting with a very high magnitude guide star.

The key remaining questions we now tackle relate to the behaviour of networks in the presence of noise. For that investigation, we find it informative to investigate the frequency response of networks, especially when it comes to detail in their outputs that is beyond spatial frequencies available to the WFS. In order to investigate CNNs systematically here we pursue a familiar statistical analysis method to investigate the spatial frequency response of networks. By studying in simulation using COMPASS, true information about the wavefront is also available for comparison purposes. By projecting wavefronts onto a KL modal basis [7], a vector of weights describing the impact of each mode on the total wavefront is produced. Our analysis proceeds by making comparisons of such modal weights. Specifically, by examining the variance of weights over a large number of samples, such as in a long exposure, we compare the ranges of weights produced by each network for each mode.

Our primary purpose here is to evaluate and train networks in noisy situations, and thereby get a sense of performance in situations with low photon flux. The COMPASS tool is able to simulate with photon noise and Read Out Noise (RON) to the simulated sensors. Using COMPASS we train and evaluate networks with a view to evaluating how noise affects performance in familiar control and PSF-R case studies.

2. MODAL ANALYSIS OF CGAN AND UNET

To investigate the spatial frequency response of the CNNs, we calculate the variance of the modal weights of the short exposure estimates over the duration of a long exposure period. We use 20,000 short exposure images in a continuous simulation and compare the statistics for the UNet and the cGAN with the COMPASS ground truth closed-loop residual wavefront. We also compare the ‘free turbulence’ wavefront—i.e., the open loop or uncorrected wavefront—which is included to validate our modal analysis. Plotted weights associated with free turbulence exhibit the expected smooth curve structure with smaller weights associated with larger mode indices.

The cGAN and UNet networks studied here are the existing trained networks from our previous paper [1]. Both networks were trained on the same simulated data. The UNet and cGAN have the same CNN structure, the difference between them being a consequence of training with (i.e., cGAN) and without (i.e., UNet) an adversarial loss term. The cGAN and the UNet are iteratively given the SH-WFS images from the simulated ground truth and used to create estimates from the same closed loop SH-WFS, collecting the variance statistics for each short exposure wavefront phase estimate over the 20,000 iterations. In addition the free-turbulence variance is also calculated for the same simulation by running in open loop, with the same pseudo-random seed and so with the same atmosphere for each iteration.

Each short exposure wavefront image is projected onto the Karhunen–Loève (KL) modal basis to produce a vector of KL mode weights for each image, where the KL modal basis is generated in COMPASS. From the 20,000 samples, the variance can be calculated so the each mode can be compared. In addition to the variance of the modal weights, the variance of the difference can also be calculated for the cGAN and UNet inferred wavefront estimates where the difference is calculated as the difference in short exposure modal weights. This gives us the variance of the instantaneous or short exposure error that is necessary to test the accuracy of each inference, rather than the accuracy over the 20,000 frame long exposure.

2.1 Variance of KL modal weights in estimation methods

When looking at the accuracy of a stochastic process such as atmospheric variations it is useful to investigate the variance of the distribution of values over a sample data set. The set of modal weights for a long exposure wavefront estimate gives us a distribution of values for each mode, where the mode number is an analogue for the spatial frequency – i.e. the lower order modes describe lower spatial frequencies, and higher order modes describe higher spatial frequencies.

We know that there will be some spatial frequency cutoff where the SH-WFS contains diminishing degrees of higher spatial frequency information. Verifying that this is the case helps to explain what the CNNs are interpreting in the WFS. The second important analysis point is how well the distributions of the CNNs match that of the simulated ground truth. The variance for each modal weight from the cGAN, UNet, closed-loop simulated ground truth and open loop ‘free turbulence’ distributions are shown in Fig. 1.




Figure 1: Wavefront variance over 20,000 consecutive frames of the simulated AO loop at high photon flux.

Interpretation of Fig. 1 yields some very useful results for understanding the nature of the CNNs estimation. The first observation is that the distribution of the modal weights for the cGAN estimation is so well aligned with that of the simulated ground truth (COMPASS), it is difficult to see both series on the figure. This observation confirms what we have seen in previous analysis of the cGAN where the estimation has great accuracy, even at higher spatial frequencies (corresponding with the wider halo estimation) [1], in simulation of a long exposure PSF estimate.

A second observation is that the UNet modal weight variance deviates strongly from the other distributions after the first 400 – 500 modes, which is roughly where we would expect the SH-WFS low-pass cutoff to be. Our visual inspection of the UNet short exposure phase estimates indicate a loss of detail compared with the cGAN and simulated ground truth which aligns well with the result in Fig. 1.

Validation of the approach by comparing the variance of the modal weights of the simulated ground truth in closed (reference) and open-loop (free turbulence) shows the expected high variance of the modal weights in the

free turbulence to that of the closed-loop distribution where the free turbulence matches that of the closed-loop series (and the cGAN) once the low pass frequency limit is reached – again around 400 – 500 modes. This example is an interesting alternative to the typical visual comparison of the PSF when demonstrating the benefit of AO systems, where the variance of the wavefront distortions in closed-loop are significantly reduced.

2.2 Errors in inferred modal weights

We have seen in Fig. 1 that the UNet does not accurately estimate high spatial frequency information, and indeed appears to ignore that. On the other hand, the cGAN when evaluated in long exposure appears to accurately model such information. The question we now pursue, is whether this accuracy persists at short exposure. Continuing with modal analysis, we investigate how related the fine detail in the phase estimates of a cGAN is to a corresponding WFS image.




Figure 2: Wavefront inference error variance over 20000 consecutive frames of the simulated AO loop (with reference wavefront (green))

To investigate this, a second variance analysis of inference error between the short exposure estimates from both the cGAN, the UNet, and the ground truth. This will show the accuracy (or lack of it) of the modal weights in the short exposure, rather than over the long exposure. The variance of the inference error is calculated by subtracting the ground truth modal weight (COMPASS) from the respective modal weight for that of the UNet and cGAN for each short exposure estimation iteration and taking the variance of distribution of the error and comparing that to the variance of the reference ground truth modal weights (COMPASS).

The variance of the inference error is shown in Fig. 2. For the low order modes, the variance of the inference error in the modal weights for both the cGAN and the UNet show good accuracy as the variance of the difference is smaller than the variance of the ground truth (COMPASS). After the 400 – 500 mode mark we see the benefit of both CNNs disappear. The difference of the UNet to the ground truth becomes very small, however the cGAN becomes consistently worse at the higher modes. This indicates that while the UNet is not making any estimate contributions above the SH-WFS cut off, the cGAN is consistently adding incorrect information to its estimates.

This is a useful result for choosing a type of CNN. A cGAN makes excellent long exposure estimates, as seen in experimental results and now in the variance analysis from Fig. 1. Short exposure cGAN estimates at higher frequencies are consistently incorrect. The explanation for this is that the cGAN is able to interpret the low order information available in the SH-WFS, in the same way the UNet can, but at higher order it will generate the missing information from a learned distribution of the atmospheric statistics. These sampled statistics the cGAN uses to fill in the higher frequency information are purely generative, rather than translating from the SH-WFS. Indeed, such information is not available using that sensor. The excellent performance of the wavefront estimation for PSF-R owes its accuracy to the long exposure requirement of the PSF-R application. Generally we have found that cGANs fail to outperform similarly trained UNets on higher order deformable mirror control applications, specifically because the short exposure generative component of the cGAN is always going to be detrimental in the control setting where single frame accuracy is critical.

3. NOISE ANALYSIS OF CNN METHODS FOR WAVEFRONT ESTIMATION

There are two types of noise that can effect the SH-WFS in the AO system. One type of noise is photon noise, also known as Shot noise or Poisson noise, caused when there are limited numbers of photons incident on the CCD of the SH-WFS which adds a stochastic element to the measurement due to the discrete nature of light. The second noise effect is RON which refers to false measurements where a CCD pixel reads a spontaneous electron rather than the expected effect from an incident photon. The RON of an instrument is typically one or two photons per pixel.

Both of these noise effects are insignificant where there is a bright guide star and the photon flux is high, when the photon flux approaches the instruments lower limits (such as in the case of extreme adaptive optics) these noise effects become problematic. This section investigates the impact of noise on the CNN and characterises the effects. Table 1 shows the guide star magnitudes for our simulated instrument and the resulting photons per sub-aperture and photons per pixel in the SH-WFS.

Table 1. Photon counts at guide star magnitudes for SH-WFS with 16×16 sub-apertures and 8×8 pixels per sub-aperture

Guide Star Magnitude	Photons per sub-aperture	Photons per pixel
10	2500.00	39.06
11	995.27	15.55
12	396.22	6.19
13	157.74	2.46
14	62.80	0.98
15	25.00	0.39

While it is useful to understand the noise in typical *Signal to Noise Ratio* (SNR) terms, it is ambiguous in this setting. While the SNR can be calculated across all pixels in a sub-aperture of the SH-WFS, the true SNR will vary dependant on the turbulence effecting the concentration of the photons on a subset of pixels. For this reason, the typical measurement of the signal is the photon flux per sub-aperture, measured in photons as the number of photons per sub-aperture is not affected by the degree of turbulence. We adopt this convention for the purpose of noise analysis developed below. Table 2 shows the calculated SNR for Table 1, using Eq. 1 and the assumption of uniformly spread photons in the sub-aperture.

$$SNR = \frac{Photons}{\sqrt{Photons + RON^2}} \quad (1)$$

Table 2. Relative SNR to guide star magnitude for simulated geometry and Read-out Noise

Readout Noise	Guide Star Magnitude					
	10	11	12	13	14	15
0	6.25	3.94	2.49	1.57	0.99	0.63
1	6.17	3.82	2.31	1.32	0.70	0.33
2	5.95	3.52	1.94	0.97	0.44	0.19
3	5.63	3.14	1.59	0.73	0.31	0.13

3.1 UNet assisted open-loop control with noise

We report on a new network that we have trained called the ‘low-flux’ UNet. Our aim is to test the training feasibility and accuracy of inference. Training and test scenarios differ over multiple variables. In addition to seeking robustness to changes in Fried parameter, we introduce varying quantities of noise and also varying guide star magnitudes. Training and evaluating a network over a wide range of realistic guide star magnitudes assist noise analysis. Further, such a general network is motivated directly by control applications where guide star magnitudes will vary between observation targets. Otherwise our approach here follows our previous work with UNets and cGANs, where the simulated AO system and telescope parameters here are identical to those used in previous published work [2].

To examine the effect of noise where the photon flux does not easily dominate the noise effects, the simulated geometry was tested for the lowest guide star magnitude before the linear control of the Woofer DM shows significant instability. Instability was clear below a magnitude 13 guide star for this simulated system. From Table 1, we can see this has a photon flux of 157 photons per sub-aperture. This corresponds to a notional SNR of around 1 for a the typical noise level of one or two photons per pixel of RON expected for a SH-WFS – see Table 2. This new ‘low-flux’ UNet was trained successfully on the range of guide star magnitudes (10 to 13), noise (RON of 0, 1, 2 and 3 photons per pixel with photon noise on for positive values of RON) and r_0 (0.05, 0.10, 0.15, and 0.20m).

Analysis consists of the UNet controlling an open loop Tweeter DM that fine tunes the closed loop Woofer DM by estimating the residual wavefront after correction by the Woofer DM of the closed loop from the SH-WFS which we will refer to as UNet Assisted Open Loop Control (UNet AOLC). We use the analysis method from our previous work published in paper [2], where we compare the performance of the UNet AOLC with two baselines, one that replaces the UNet Tweeter control with linear Tweeter control (Linear AOLC) and the second baseline of a single DM in closed loop which we refer to as Closed Loop Control (CLC), where we eliminate the open loop Tweeter DM by holding the mirror flat.

We have found that exposing the DM shape to the network-based phase estimators universally improves performance, and for the sake of producing estimators for control, we expose the DM Woofer shape. The UNet was trained easily with the same parameters as previous UNets and the results of analysis for a Magnitude 10 guide star (highest photon flux trained) are shown in Tables 3 and 4 which follow the format of analysis results described in our previous study [2].

Using our ‘low-flux’ UNet for UNet AOLC with an additional 7 actuators in the open loop Tweeter mirror, the UNet AOLC gains more than 7.5 points of Strehl Ratio over the Linear AOLC baseline. This result is extended to 10 points of Strehl Ratio over the Linear AOLC when the Fried parameter is reduced to $r_0 = 0.06m$. The reader can note that the Strehl Ratio performance gains of our network-based approach relative to a linear baseline here exceed those we have reported on in a previous study [2]. From Table 4, the UNet AOLC consistently increases performance over the Linear AOLC as the Fried parameter is reduced — i.e., in poor viewing conditions due to the turbulence of the atmosphere. We can clearly see that the ‘low-flux’ UNet can be trained with noise and variable guide star magnitude with no degradation to noiseless inference. The use of the Woofer DM shape as an additional input channel for the UNet has extended performance, and we have found that UNets have outperformed cGANs for in all settings when considering the control use case. In the interest of space, here we only report on UNets when considering the robustness of network-approaches for control.

Table 3. UNet AOLC - Performance (in Strehl Ratio at 1650 nm) of each control scheme. Each controller is evaluated under each system configuration, with a fixed r_0 of 0.10m and Guide star Magnitude 10 and no noise (Δ WFE indicates the relative RMS wavefront error in nanometres for each method compared to the linear controller)

AO System	Linear Control	UNet Control		Oracle Control	
	LE SR	LE SR	Δ WFE	LE SR	Δ WFE
No Open Loop DM	62.20%	N/A	N/A	73.76%	108.42
Open-Loop +0 (17x17)	62.37%	69.00%	92.70	73.76%	107.55
Open-Loop +1 (18x18)	63.96%	71.90%	88.72	75.76%	108.06
Open-Loop +3 (20x20)	64.06%	73.02%	86.22	79.17%	120.85
Open-Loop +7 (24x24)	66.77%	75.41%	91.47	83.99%	125.79
Open-Loop +15 (32x32)	69.39%	77.28%	86.11	89.21%	131.63
Open-Loop +31 (48x48)	68.55%	77.61%	93.37	93.41%	146.08

Table 4. UNet Assisted Open loop with additional DM shape inputs - Robustness to atmospheric turbulent conditions (r_0) of optimal setting, compared to linear control baselines (both with +7 actuators on the tweeter DM, in Strehl Ratio at 1650 nm and Guide star Magnitude 10 and no noise. Δ WFE indicates the relative RMS wavefront error in nanometres compared to the linear controller)

Fried parameter r_0 (m)	Linear Control +7 actu (24x24)	UNet +7 actu (24x24)	
	LE SR	LE SR	Δ WFE
0.04	9.87%	17.08%	234.98
0.05	25.22%	35.01%	174.47
0.06	38.72%	48.58%	142.06
0.08	56.43%	65.71%	107.82
0.10	66.77%	75.41%	91.47
0.12	73.34%	81.41%	85.19
0.14	77.77%	85.35%	80.98
0.16	80.94%	88.09%	77.51

3.2 Effects of noise on UNet Assisted Open-loop Control Performance

With our new ‘low-flux’ UNet showing great performance with a magnitude 10 guide star and no noise, we can now investigate performance of the same network with the introduction of noise, down to the limit of the photon flux that the linear controller can tolerate. Here we analyse the addition of photon and RON noise to the analysis and also lower the SNR by limiting the photon flux with decreasing guide star magnitude.

As with previous performance measurements, we test both zero extra actuators in the Tweeter mirror, and then with seven extra actuators. The results for zero extra actuators are shown in Fig. 3 for each of the different guide star magnitudes and associated photon flux. At Magnitude 10, the photon flux is 2,500 photons per sub aperture which is relatively high and corresponds to an estimated SNR of around 6 (see Table 2). Note here that for the zero noise case, the photon noise is also turned off in the simulation tools, and is ‘on’ for all positive values of RON. As photon noise is proportional to the inverse square root of the photon flux, the impact of photon noise on SNR will increase as the photon flux becomes small.

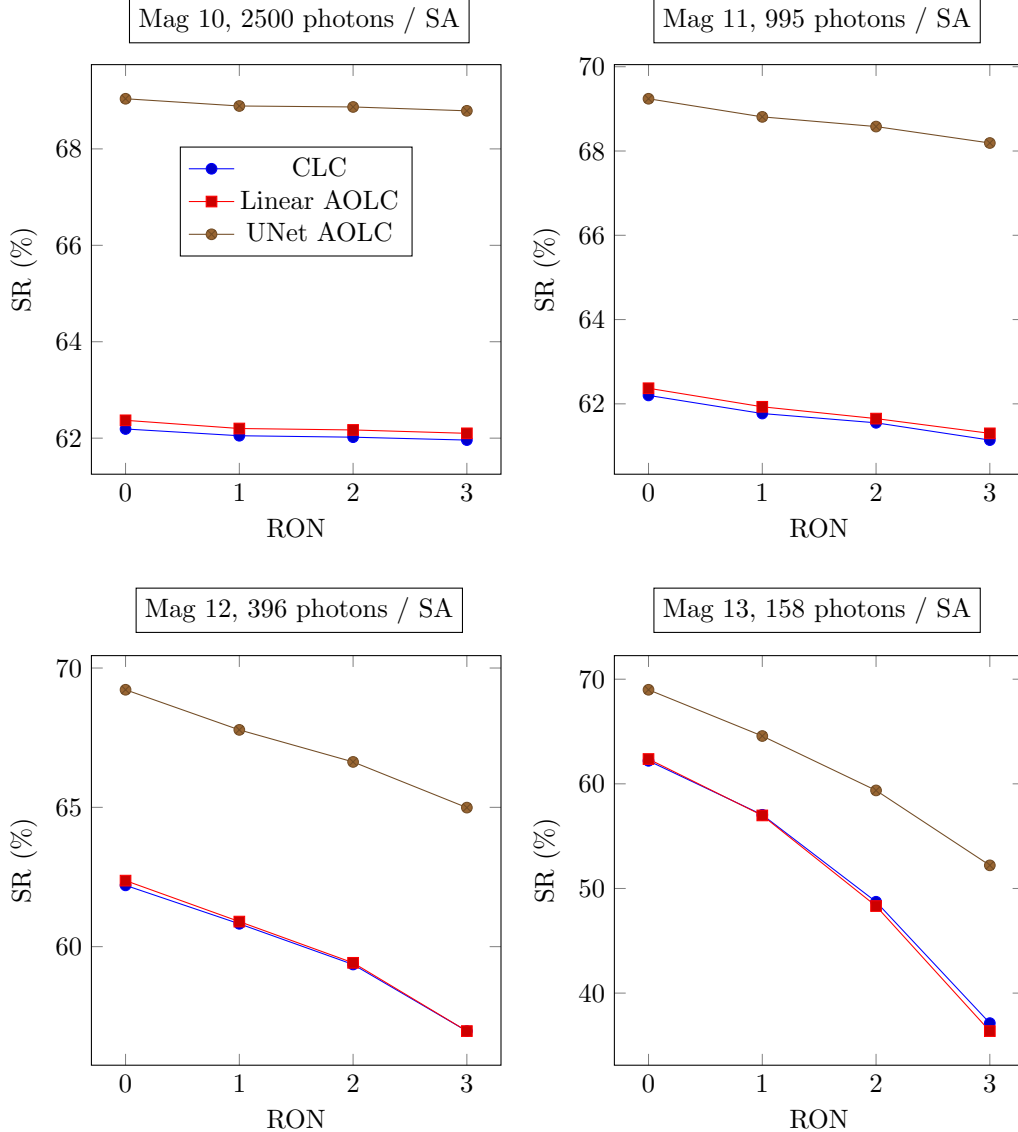


Figure 3: Expected SR achieved by control regimes are reported for a range of signal to noise ratios encountered by varying the guide star magnitude and RON. Zero extra actuators in Tweeter mirror and $r_0 = 0.10$ m.

The performance measured in estimated SR shows that for zero noise, we observe the expected 7% SR improvement over the Linear AOLC benchmark consistent with Table 3. As we turn on the photon noise and increase the RON we can see a degradation in all three methods but this is slight. As the guide star is sufficiently bright, the signal is almost unaffected by noise. The main observation here is the presence of photon noise has a greater impact on performance than the RON, which can be seen between RON values of 0 and 1. While the photon noise effect on performance is small, it is clearly a significant effect relative to the RON.

Progressively decreasing the photon flux by moving to fainter guide stars, the effects of noise become more pronounced. A clear result is that the benefits of the UNet AOLC are not lost with increased noise at low photon flux, nor does the decrease in photon flux affect the UNet AOLC more than the benchmark methods. In fact, at very low photon flux and high RON, we see significant out-performance by the UNet Assisted Open Loop where the extreme cases are well over 10 % SR. This is clearly in a regime where the linear controller is becoming unstable and so, not only does the UNet perform comparatively well, the network is clearly compensating for some of the linear controller instabilities.




Figure 4: Expected SR achieved by control regimes are reported for a range of signal to noise ratios encountered by varying the guide star magnitude and RON. Seven extra actuators in tweeter mirror and $r_0 = 0.10$ m.

By adding the extra 7 actuators to the Tweeter mirror we can see the effect of increasing the fidelity and allowing the benefit of the high detail estimates from the UNet for controlling the Tweeter mirror. In Fig. 4 we can now see the performance of the two baseline configurations diverge as the Linear AOLC also makes use of the additional actuators in the Tweeter mirror. The UNet AOLC extends the performance benefit over the Linear AOLC benchmark to over 8.5% SR improvement, and more than 13% SR over the CLC baseline performance.

As the photon flux decreases and the noise is increased, again we see the effect of photon noise is significant. The UNet AOLC and CLC baseline follow similar trends to the zero extra actuator case. However, the Linear AOLC benchmark suffers greatly with low photon flux and higher noise. While the UNet AOLC is still showing strong robustness to noise, as the CLC method struggles to interpret the noisy SH-WFS images, the higher order estimation required for Linear AOLC is strongly impacted. This shows that while our preferred benchmark of the Linear AOLC is a good comparison in high photon flux, in the more extreme cases of low photon flux and high noise, that benchmark fails to compete with the CLC baseline, and is outperformed by the UNet AOLC by over 20 % SR in more extreme noise and low photon flux.

3.3 Modal analysis of UNet inference for noisy estimation

With the performance analysis showing great results for robustness to noise and photon flux, we can now use the modal analysis to investigate the effect of noise on the variance of the modal weights of the inferred wavefront estimates. This will show which modes are effected by noise in contrast to the variance of the reference ground truth modal weights. As with previous modal analysis we take 20,000 short exposure wavefront images and calculate the variance statistics of the modal weights.



Figure 5: Variance (log) between KL modal weights of the inferred wavefront estimates and simulated ground truth

The chart in Fig. 5 plots the log of the variance of the modal weights for each mode for the reference linear controller for both noise Off (yellow) and noise On (green) for a magnitude 12 guide star where we know the photon flux is low and will provide significant noise for analysis. The variance statistics of the inferred wavefront estimates from the UNet are calculated in the same way for noise Off (blue) and On (orange). The noise value of 1 represents a RON of 1 and the switching on of photon noise.

We can see from the reference series that there is only a small difference in the variances of the modal weights (the green series is barely visible) and the only slight visible differences are in the first 200 modes with the lowest order modes more affected with a small increase in the variances. In contrast, the UNet tends to underestimate the variance slightly with the increase in noise and it affects all modes. We also see the UNet follows a similar profile to that of the high photon flux analysis in Section 2, where there is a slightly smaller variance in the estimates to that of the reference ground truth image from the simulator. These noise observations show slight variations however and may not be indicative of any significant affects.

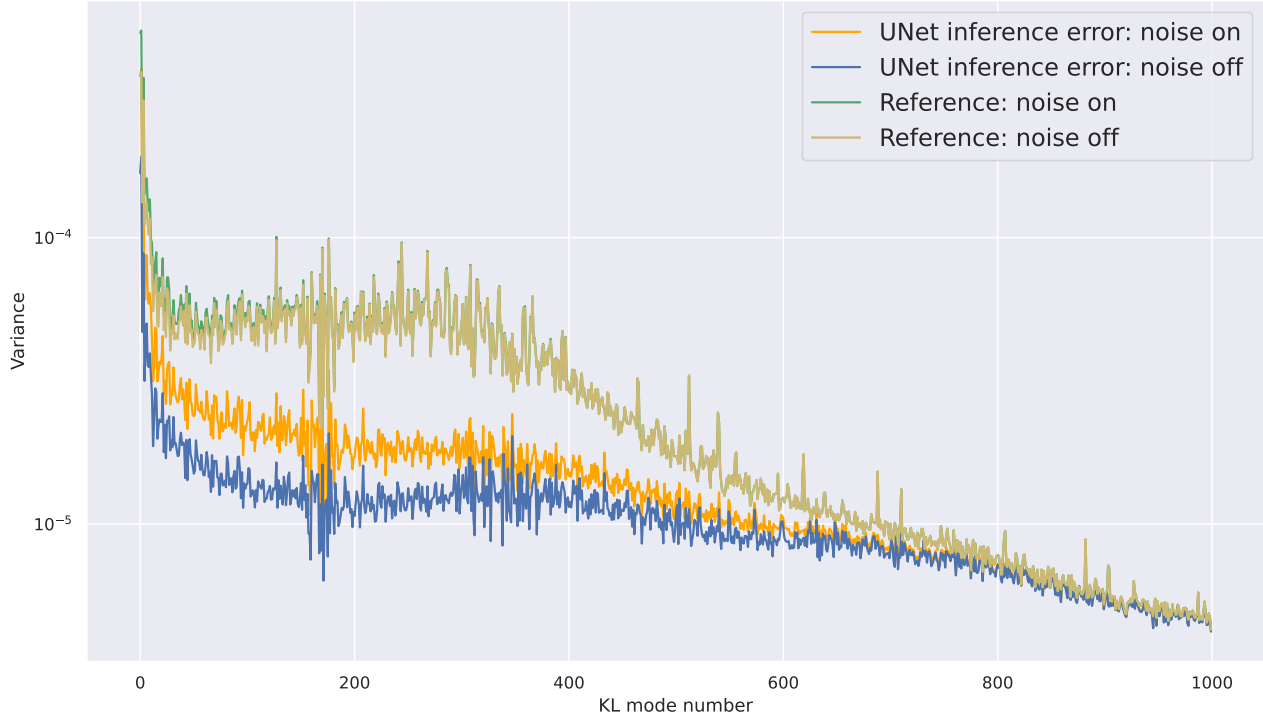


Figure 6: Variance (log) of the difference between KL modal weights of the inferred wavefront estimates and simulated ground truth

Moving to the variance of the inference error analysis, we can see the effect of noise on the inferred wavefront estimates. Now that we are measuring the variance of the error in modal weights between the inferred estimates and the simulated ground truth, there is a jump in the variance for the noisy data which is clearly significant and effecting the first 600 modes where the lowest order modes are more affected. The impact of noise on the UNet in low SNR regimes is that the wavefront estimation error increases, albeit remaining much smaller than the reference variance over the range of modes important for AO control.

4. EFFECT OF NOISE ON PSF-R USING NETWORKS TRAINED WITH ADVERSARIAL LOSS

While the best performance we have measured in real-time control uses a UNet, the UNet was demonstrably inadequate for reconstruction the PSF halo as shown in our previous paper [1]. Without the adversarial loss term used to train the cGAN, the wavefront estimates do not contain the high spatial frequency information required for accurate halo reconstruction for the purposes of PSF-R. Given the spatial frequency response of deformable mirrors, this limited wavefront reconstruction is enough for AO control.

The cGAN learns more than just how to interpret the available information in the SH-WFS. Specifically, due to the Markov discriminator component accessing the statistics of the Kolmogorov turbulence model related to the simulated wavefront, the cGAN is able to generate a random sample of the missing spatial frequencies above the limitations of the SH-WFS. While the short exposure inference for the generated higher spatial frequencies from the cGAN network are inaccurate due to the stochastic generative nature of the inference, it is very accurate over a long exposure estimate that is the average of 20,000 samples. The cGAN is also more sensitive to the presence of noise in the SH-WFS and we examine this in the following section.

4.1 Noise effect on training cGAN networks

Several attempts to train a ‘low-flux’ cGAN were made for on the same low magnitude guide star and noise data we used to train the ‘low-flux’ UNet. However, the training of the cGAN networks did not have the same success as the UNet. Training of the cGAN collapsed early in the training process. We have not been successful in training with noisy data and an adversarial loss. Problematic training of GANs is well known, and there is an opportunity here for future work with more advanced cGAN methods that are designed to mitigate these training problems.

4.2 Effect of noise on cGAN trained without noise

Since we were unable to train a cGAN with the same training dataset as the UNet, a new ‘low-flux’ cGAN network was trained on a dataset that contained only a single guide star magnitude of 12 and with no noise in the training dataset and the typical range of r_0 values (0.05, 0.10, 0.15, 0.20 m). This did effectively train and produces wavefront estimates and long exposure PSFs that exhibit quality metrics in the case of a low magnitude guide that are analogous to those presented in our earlier published works [1]. By using a guide star of magnitude 12 the photon flux is reduced to where photon and read-out noise are significant with respect to the signal, and so we expect to see a strong effect when turning on the noise for analysis.

As the ‘low-flux’ cGAN was not able to be trained with simulated noisy data, as we did for the UNet analysis, we can only add noise to the analysis data and observe the effects on the wavefront estimates when compared with noiseless analysis results, the ground truth and our reference benchmark [8] with the same methods described in [1]. The ‘low-flux’ cGAN estimates of the wavefront are significantly impacted by the addition of photon noise and RON, even at small values of RON. The effect of turning on the noise is seen in Fig. 7 where the x and y split profiles show a dramatic change in the overall accuracy of the long exposure PSF built from the inferred wavefront estimates. The primary fault is the significant error in the central peak of the PSF, which corresponds to the estimate SR missing its target by nearly 7 points of SR, much worse than the noiseless result. The top down view shows significant impact as well in Fig. 8 where the circular average has jumped in the inferred result, though still retains significant improvement over the benchmark in the near halo, retaining much of the benefit of the data driven accuracy.





Figure 7: Long Exposure PSF-R analysis with RON of 1 and photon noise - x and y splits of the long exposure PSF for magnitude 12 guide star, near the limit of closed loop linear control stability. - 396 photons per sub-aperture, $r_0 = 0.10\text{m}$

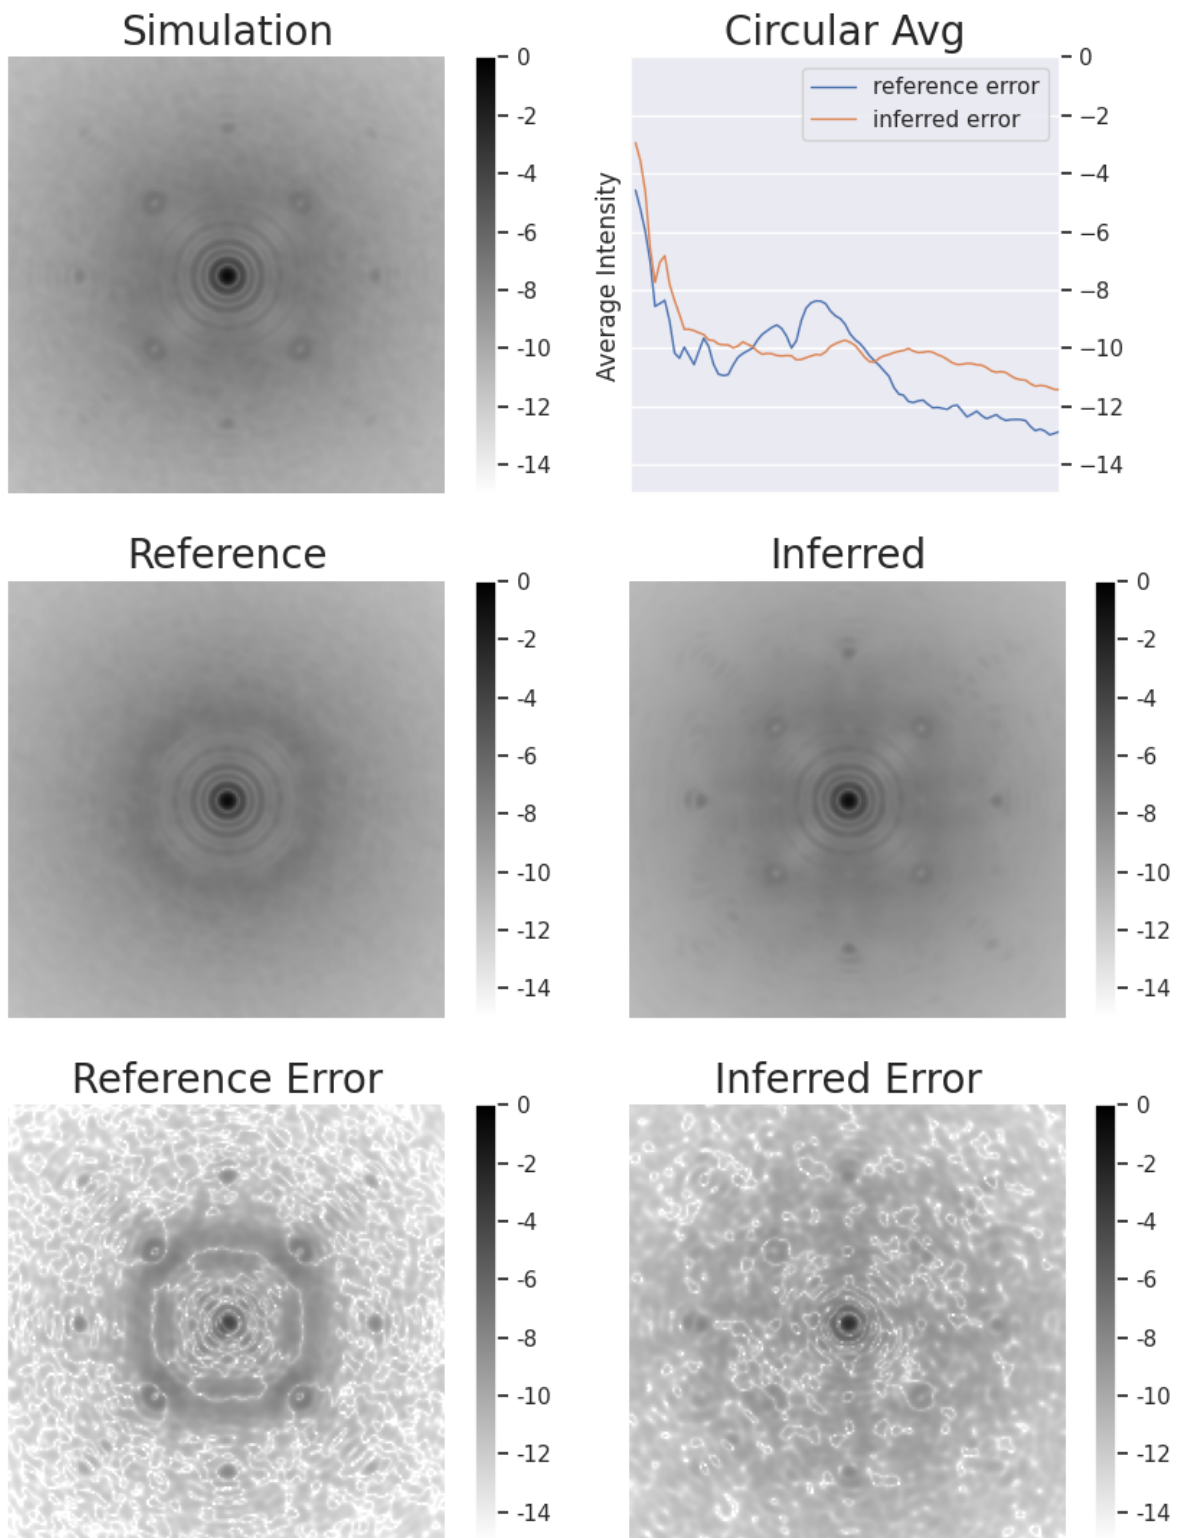


Figure 8: Long Exposure PSF-R analysis with RON of 1 and photon noise - top down view of the long exposure PSF for magnitude 12 guide star, near the limit of closed loop linear control stability. - 396 photons per sub-aperture




Figure 9: Wavefront inference error for cGAN with (orange) and without (blue) photon noise for magnitude 12 guide star, near the limit of closed loop linear control stability. 396 photons per sub-aperture, $r_0 = 0.10$ m and $\text{RON} = 0$

Testing with photon noise turned on and a RON of 1 photon per pixel for analysis data, at a photon flux of 396 photons per sub-aperture the cGAN estimates show strong deviation from the noiseless estimates, where the linear controlled wavefronts from simulation were barely effected. Ablation testing with the RON set to zero and only photon noise added to the SH-WFS in simulation, the analysis shown in Fig. 9 is clearly indicating photon noise is the dominant noise effect. Given that the photon noise is purely controlled by the number of photons available to each SH-WFS measurement, the photon noise will be a limiting factor for the cGAN at low photon flux for PSF-R.

5. DISCUSSION AND FUTURE WORK

We developed a modal analysis technique to study the frequency response and accuracy of network-based methods for wavefront estimation. We were able to investigate the accuracy and composition of wavefront estimates. We suppose our approach shall be useful in future investigations of wavefront estimation methods. It has been particularly useful in this project to further validate our choice of networks for PSF-R and AO control tasks, and also to demonstrate limits in WFS and DM frequency response for the low pass filter behavior that these devices create intrinsically by their design. The use of a UNet for Assisted Open Loop control is well justified, with comparative results in this paper well exceeding our previous methods using the cGAN for control. With the addition of the DM shape as an AO loop telemetry input to the UNet along with the SH-WFS image, the UNet can achieve up to 10 points of SR improvement over linear control methods. One other interesting feature of our results is that as the atmospheric conditions become more difficult (ie. r_0 is reduced), the UNet improves performance in terms of SR and RMS error. This is expected to be due to the form of the sub-aperture PSF producing more diverse and so more easily interpreted high frequency information for the UNet to interpret.

For PSF-R, we can see from our efforts at training and limited noise analysis, that photon noise is a problem for the cGAN. Clearly where the photon flux is high enough to overcome the photon noise, the cGAN is an excellent

estimation tool for long exposure PSF-R. At low photon flux configurations, there may be a few approaches that can assist a workflow that uses a cGAN. A network cannot learn to filter noise unless it is exposed to noise in training, and from our experience with the cGAN this motivates future work examining architectures and loss regimes that will be relatively robust to noise at both training and inference time. We expect future investigations into the MMD-GAN [9, 10] and Wasserstein GAN [11] to be potentially fruitful. As another approach, noise could be treated at inference time, through the use of denoising methods such as filtering the SH-WFS images through an autoencoder to remove the noise effects prior to wavefront estimation [12]. Another possible method could be to use the UNet for robust estimation of some modes, and in a hybrid configuration use numerical methods to complete the halo for PSF-R for offline workflows. This has the benefit of out-of-the-box robustness to noise, with the accuracy of the numerical methods at higher frequencies.

6. CONCLUSION

We have developed a method of using a KL modal basis to uncover statistical details of the cGAN and UNet CNNs. We have used this to understand the strengths and weaknesses of network-based image to image translation for wavefront estimation. The analysis informs network choice and characterises the relative strengths and weaknesses of networks. Specifically we show that the UNet translates spatial frequency information from the WFS, which gives a corresponding wavefront estimate which is missing the higher order spatial frequency information; also missing from the SH-WFS. In contrast, the cGAN is adding a generative component that is learned from the simulated Kolmogrov turbulence, and is adding a sample of this to the wavefront estimate which is not accurate in the single frame estimate, however which statistical convergence is extremely accurate over a long exposure.

We demonstrate an improved method of UNet Assisted Open Loop control with up to 10 points of SR benefit over the Linear Assisted Open Loop controlled baseline where the DM mirror shape is added to the SH-WFS input to the CNN from AO loop telemetry data. Given our results we show that the UNet is ideal for AO control and even excels in low photon flux, high noise and over a range of magnitudes, Fried parameters and degrees of noise, where the cGAN is ideal for PSF-R tasks provided there is insignificant photon noise.

ACKNOWLEDGMENTS

Many thanks to Florian Ferreira for donating his time and knowledge assisting with COMPASS. Thanks also to Bartomeu Pou Mulet and Hao Zang for numerous discussions that enhanced knowledge in this field of research. This work was supported in part by Oracle Cloud credits and related resources provided by the Oracle for Research program. This research was undertaken with the assistance of resources from the National Computational Infrastructure (NCI Australia), an NCRIS enabled capability supported by the Australian Government.

References

- [1] Jeffrey Smith et al. “Image-to-image translation for wavefront and point spread function estimation”. In: *Journal of Astronomical Telescopes, Instruments, and Systems* 9.1 (2023), p. 019001. DOI: [10.1117/1.JATIS.9.1.019001](https://doi.org/10.1117/1.JATIS.9.1.019001). URL: <https://doi.org/10.1117/1.JATIS.9.1.019001>.
- [2] Jeffrey Smith et al. “Enhanced adaptive optics control with image to image translation”. In: *Uncertainty in Artificial Intelligence, Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence, UAI 2022, 1-5 August 2022, Eindhoven, The Netherlands*. Ed. by James Cussens and Kun Zhang 0001. Vol. 180. Proceedings of Machine Learning Research. PMLR, 2022, pp. 1846–1856. URL: <https://proceedings.mlr.press/v180/smith22a.html>.
- [3] Jeffrey Smith et al. “Image-to-image translation for wavefront and PSF estimation”. In: *Adaptive Optics Systems VIII*. Ed. by Laura Schreiber, Dirk Schmidt, and Elise Vernet. Vol. 12185. International Society for Optics and Photonics. SPIE, 2022, p. 121852L. DOI: [10.1117/12.2629638](https://doi.org/10.1117/12.2629638). URL: <https://doi.org/10.1117/12.2629638>.
- [4] B. Pou et al. “Model-free reinforcement learning with a non-linear reconstructor for closed-loop adaptive optics control with a pyramid wavefront sensor”. In: *Adaptive Optics Systems VIII*. Ed. by Laura Schreiber, Dirk Schmidt, and Elise Vernet. Vol. 12185. International Society for Optics and Photonics. SPIE, 2022, 121852U. DOI: [10.1117/12.2627849](https://doi.org/10.1117/12.2627849). URL: <https://doi.org/10.1117/12.2627849>.

- [5] Phillip Isola et al. “Image-to-image translation with conditional adversarial networks”. In: *Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017*. Vol. 2017-January. 2017.
- [6] Florian Ferreira et al. “COMPASS: An efficient GPU-based simulation software for adaptive optics systems”. In: *Proceedings - 2018 International Conference on High Performance Computing and Simulation, HPCS 2018*. 2018.
- [7] E. Gendron and P. Léna. “Astronomical adaptive optics. II. Experimental results of an optimized modal control”. In: *Astronomy and Astrophysics Supplement Series* 111 (Apr. 1995), p. 153.
- [8] F Ferreira et al. “Numerical estimation of wavefront error breakdown in adaptive optics”. In: *Astronomy and astrophysics (Berlin)* 616 (2018), A102.
- [9] Mikołaj Binkowski et al. “Demystifying MMD GANs”. In: *6th International Conference on Learning Representations, ICLR 2018 - Conference Track Proceedings*. 2018.
- [10] Chun Liang Li et al. “MMD GAN: Towards deeper understanding of moment matching network”. In: *Advances in Neural Information Processing Systems*. Vol. 2017-December. 2017.
- [11] Martin Arjovsky, Soumith Chintala, and Léon Bottou. *Wasserstein GAN*. 2017.
- [12] B. Pou et al. “Denoising wavefront sensor images with deep neural networks”. In: *Adaptive Optics Systems VII*. Ed. by Laura Schreiber, Dirk Schmidt, and Elise Vernet. Vol. 11448. International Society for Optics and Photonics. SPIE, 2020, 114484J. DOI: [10.1117/12.2576242](https://doi.org/10.1117/12.2576242). URL: <https://doi.org/10.1117/12.2576242>.