PROCEEDINGS OF SPIE

SPIEDigitalLibrary.org/conference-proceedings-of-spie

Shack-Hartmann wavefront reconstruction by deep learning neural network for adaptive optics

Zareb Noel, Timothy Bukowski, Stanislav Gordeyev, R. Mark Rennie

Zareb A. Noel, Timothy J. Bukowski, Stanislav Gordeyev, R. Mark Rennie, "Shack-Hartmann wavefront reconstruction by deep learning neural network for adaptive optics," Proc. SPIE 12693, Unconventional Imaging, Sensing, and Adaptive Optics 2023, 126930G (3 October 2023); doi: 10.1117/12.2677670

SPIE.

Event: SPIE Optical Engineering + Applications, 2023, San Diego, California, United States

Shack-Hartmann Wavefront Reconstruction by deep learning neural network for adaptive optics

Zareb A. Noel, Timothy J. Bukowski, Stanislav Gordeyev, and R. Mark Rennie

Institute for Flow Physics and Control (FlowPAC)
Hessert Laboratory for Aerospace Research,
University of Notre Dame, Notre Dame, Indiana, 46556

ABSTRACT

Standard methods of Shack-Hartmann wavefront reconstruction rely on solving a system of linear equations, extracting wavefront estimates from measured wavefront slopes, which are calculated by retrieving centroids from a Shack-Hartmann Wavefront Sensor (SHWFS). As the dimensions of a micro-lens array in the SHWFS increase, the computational cost of processing wavefronts can become increasingly expensive. For applications that require rapid and accurate computations, such as closed-loop adaptive-optic systems, traditional centroiding and the least-squares reconstruction becomes the main bottleneck limiting performance. In this work, we apply a convolutional neural network (CNN) approach to directly reconstruct wavefronts from raw SHWFS measurements, circumventing both bottlenecks. The CNN model utilizes the ResU-Net framework to perform a zonal wavefront reconstruction, and a method for preprocessing the raw data was investigated with the prospect of enhancing the accuracy of this model specifically for the zonal approach to wavefront reconstruction.

Keywords: Shack-Hartmann, Deep Learning, Wavefront, Reconstruction, Zonal Reconstruction, CNN, Neural Network, Adaptive Optics, ResNet, UNet, Res-UNet SHWFS

1. INTRODUCTION

The Shack-Hartmann Wavefront Sensor (SHWFS) has proved a mainstay for wavefront measurements in aero-optical environments,[1] astronomy,[2] and adaptive-optic (AO) correction.[3] The SHWFS sensor works by discretizing a collimated beam using a lenslet array, focusing the light from each lenslet onto the camera's CCD array, creating a "dot pattern." The position of each focused spot on the detector array is then used to determine the deflection from the reference within that lenslet's area of interest (AOI). An example of a SHWFS image is shown in Fig. 1 with red boxes showing the AOIs around each dot. This is done by calculating the centroid of the focused dots,[4] where the displacements in both the x-direction and y-direction with respect to a reference image are then related to the local wavefront slope through the physical geometry of each lenslet. The wavefront slopes is then estimated the wavefront by using a least-squares surface fit, abiding by the constraint that the reconstructed wavefront gradient at each lenslet location must equal the measured slope.[5–7] While there are many ways of reconstructing a wavefront from its slope measurements, two choices of wavefront reconstruction algorithms seem to be the most common: modal[8] and zonal.[7] For the modal approach, the wavefront is reconstructed using a Zernike decomposition, whereas the zonal approach utilizes a least-squares reconstructor, such as the Southwell method.[7] For aero-optical applications, it has been found that the modal approach does not work well, as the Zernike polynomials used in the reconstruction are oriented on a polar grid with axial direction in line with the center of the signal beam aperture, whereas most aero-optical aberrations travel laterally across the aperture. Additionally, non-circular apertures cannot be used with the modal reconstruction method, and high-frequency details are difficult to recover.[9,10] This results in poor fitting, as many Zernike modes are needed to reconstruct the wavefront with minimal error.

Other author information:
Z.A. Noel (corresponding author): znoel@nd.edu
T.J. Bukowski: tbukowski@nd.edu
S. Gordeyev: sgordeye@nd.edu
R.M. Rennie: rrennie@nd.edu

Unconventional Imaging, Sensing, and Adaptive Optics 2023, edited by Jean J. Dolne, Mark F. Spencer, Santasri R. Bose-Pillai, Proc. of SPIE Vol. 12693, 126930G
 · 0277-786X · doi: 10.1117/12.2677670

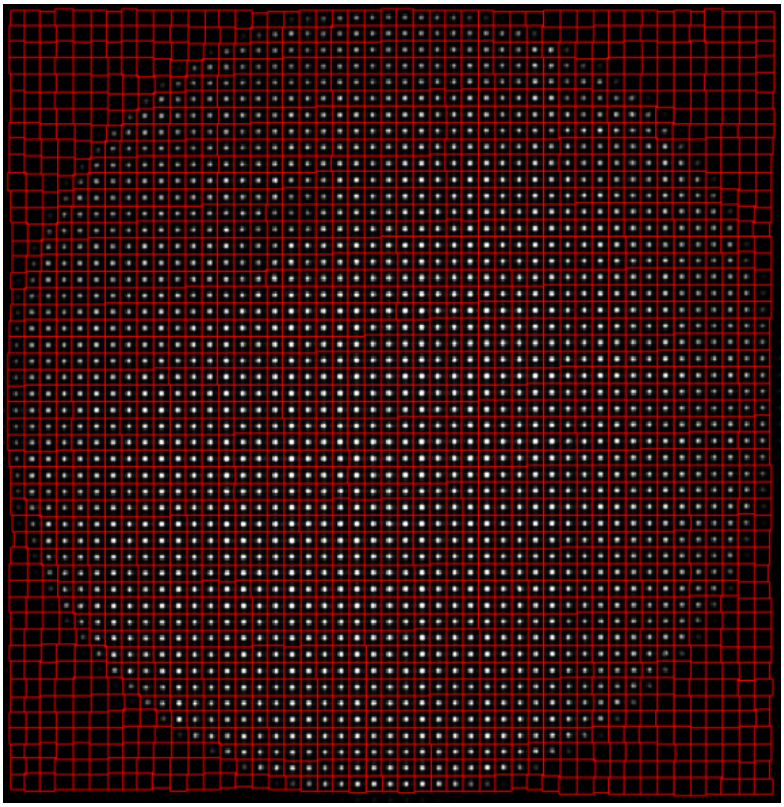

Figure 1. Raw SHWFS dot pattern with detected AOIs outlined in red. In the algorithm, only the AOIs that contain dots are used in centroiding and matrix inversion.

One of the main advantages of the SHWFS is its ease of setup and implementation of the wavefront reconstruction algorithm. The SHWFS does not need a strong illumination source, nor does it need the long coherence length of other wavefront sensing techniques such as digital holography.[11] Wavefront reconstruction can be performed by a single matrix inversion,[7] where only the slope measurements are needed. The ease of setup allows the SHWFS to be utilized in various wavefront-based imaging applications previously mentioned.

The primary focus of using deep learning in wavefront reconstruction is its potential application to adaptive optics (AO) systems, which aim to correct wavefront aberrations due to fluctuations in the refractive index caused by variable density turbulent flow. AO systems rely on wavefront measurements, which are used as the input for compensation hardware, such as deformable mirrors or fast-steering mirrors, to correct the wavefront aberrations. However, there exist considerable challenges for the implementation of a closed-loop adaptive optics system for turbulent flow, particularly that the compensation rate of AO systems is typically much slower than the evolution of the flow. Specifically, a major drawback of SHWFS in comparison to other wavefront sensing techniques is a lower phase-space resolution, which is limited by the number of lenslets, motivating an increase in lenslet array density. However, as the number of lenslets increases, so too does the computational demand for the slope and reconstruction calculations. Commonly used centroiding algorithms, which calculate the centroid of the focal spot of each lenslet, can especially be computationally slow, limiting the operating rates of adaptive optics systems.

Processing times for a standard reconstruction algorithm to process 10000 frames of SHWFS data are shown in Table. 2 as a function of number of lenslets. The time taken for centroiding is shown in blue, while the time spent performing the sparse matrix inversion for wavefront reconstruction is shown in orange. The remaining time was spent on miscellaneous operations such as loading images from directories or saving variables, and are shown in yellow. While the time spent performing miscellaneous operations remains fairly consistent as the number of AOIs increases, the cost of centroiding and wavefront reconstruction drastically increases. The associated

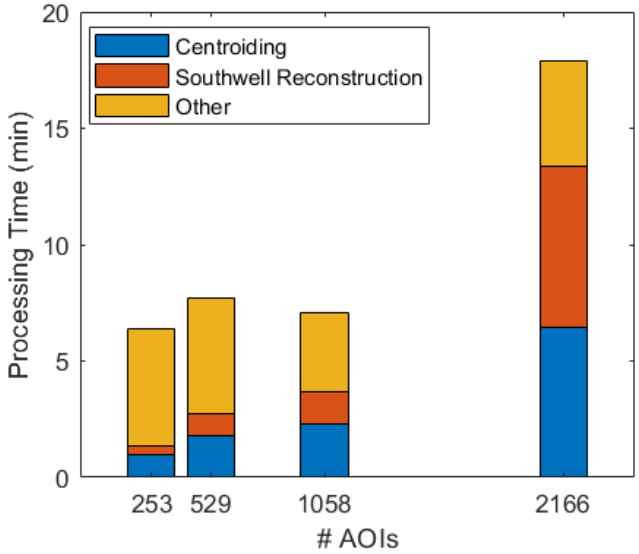

Figure 2. Processing times for varying lenslet array sizes. Processing times calculated from the average of 10000 frames using the standard least-squares Southwell[7] reconstruction algorithm.

processing rates are shown in Table 1. Note that the table only considers centroiding and reconstruction, as the computational time for miscellaneous operations are not totally necessary for adaptive optics. As shown, the processing rate decreases continuously as the lenslet resolution becomes coarser, with a maximum rate of 43 Hz for the smallest lenslet array of 253 lenslets, which gives substantially lower wavefront resolution than what would be required for adaptive optic applications, especially at the smallest turbulent length scales. The computations for Table. 2 were performed using an off-the-shelf laptop, but slow centroiding and reconstruction can certainly be mitigated using high-performance CPUs and GPUs. However, there are also limits to increasing CPU/GPU capacities, and simply scaling up these components is not feasible nor cost-effective for many applications. Recently, Kong et al.[12] employed the use of an FPGA-based system to improve centroiding and reconstruction speed and have shown success, but this requires the implementation of specialized hardware. Therefore, the development potential for faster and more efficient methods of SHWFS data analysis, that could be realized using deep learning-based approaches, could significantly improve the performance of adaptive optics systems in a cost effective way.

Table 1. Processing rates for varying lenslet array sizes.

Image Size (#AOIs)	Processing Rate(Hz)
2166	10
1058	27
528	32
253	43

Recently, concepts from deep learning have been used with a high degree of success in wavefront reconstruction from slope measurements. For example, Swanson et al.[10] reconstructed wavefronts from slope measurements using a U-Net convolutional neural network (CNN) model architecture. Similarly, Gomez et al.[13] used a convolutional neural network to predict the centroid displacements of SHWFS dot patterns, which were then used to reconstruct wavefronts from calculating the slopes. Hu et al. took this a step further, using the raw SHWFS dot patterns to directly reconstruct wavefronts without the need for slope measurements at all.[14] This was done using a ResUNet model architecture, which is the basis for the current research.

The purpose of this paper is to detail our progress in developing a CNN that performs wavefront reconstruction without the need to directly calculate centroids or to directly reconstruct wavefronts, allowing for easy integration

into the University of Notre Dame's AO System. Taking note of the results presented in Duffin et al.,[3] in order to achieve fully closed-loop AO correction; sampling, centroiding, and reconstruction, and correction occurs at frequencies at least two orders of magnitude higher than the correction rate desired. That is to say; if an AO system is targeting optical aberrations at 50Hz, then it must produce an estimation of that wavefront and send out correction signals at a rate of at least 5kHz.

In this paper, a ResUNet CNN, similar to Hu et al.,[14] is employed on a set of SHWFS images collected from wind tunnel experiments. The ResUNet architecture is a combination of the well-established ResNet[14] and U-Net,[15] which is an encoder-decoder architecture commonly used in image-to-image applications. The main concept behind ResNet is to improve model training by introducing residual connections throughout the network, which in this case bypass the subsequent convolutional layers, preserving information at specific resolutions while avoiding the problem of vanishing gradients.[10] U-Net, on the other hand, extracts feature maps from images at different resolution levels by downsampling using a combination of convolutional and maxpooling layers, and then upsampling the extracted feature maps back to the original resolution.[14] In this encoder-decoder architecture, information at each encoder level is preserved and transferred to its corresponding decoder level, creating dense feature representations without the loss of spatial information commonly associated with successive convolutions. By combining the two model architectures, ResUNet works by passing residual blocks from the encoder layers across to their equivalent decoder layers, where each residual block is composed of concatenated features extracted from the image by using differently-sized convolution kernels, as in.[14] The network's U-shape structure, shown in Fig3, reduces the spatial information with each consecutive block while increasing the number of features captured, allowing for detection and preservation of key features that are then expanded back to an image by the decoder.

2. MODEL ARCHITECTURE

While other methods use CNNs to implement modal-based wavefront reconstruction,[16] this does not work particularly well for SHWFS measurements of aero-optical environments as the distortions due to turbulent fluctuations travel laterally and are not oriented on a polar grid. Therefore, the current work aims to implement training for a zonal-based reconstruction of aero-optical wavefronts, utilizing a similar architecture to that of Wang et al.[16] To improve the training of the network, it is important to assess what exactly the model is training to learn. In modal reconstruction methods, as stated in Swanson et al.,[10] the network attempts to extract features that reproduce wavefronts as a combination of the Zernike modes. The CNN is training to learn the feature base composed of different "groupings" of dot positions that resemble Zernike modes. On the other hand, a zonal-based method works by solving a linear system of equations that relate the dot deflection value in each AOI to a wavefront value whose gradient produces that value, see.[7] As dot deflections are measured relative to some reference position (sometimes an average position), representing the relative motion from the reference positions within the input layer should produce more accurate results. Take the representation of this concept in Fig. 4; when looking at a single lenslet (AOI), the focused dot position within the AOI does not immediately detail the value of the slopes at that measurement location. However, by adding the time-average dot position, the data can now better represent deviations from the average position. Preparing the data to the input layer as a set of time-average-added SHWFS dot patterns, each AOI in each frame represents a unique deviation from an average position, which carries with it a unique value for slope vector. Put another way; the representation of dot deflections in the input layer is a feature that the network picks up on and trains for. An example of this combination is shown in Fig. 5.

For the average-added data to be properly utilized, additional changes to the model architecture are required. Primarily, the current model opts for a convolutional layer with a stride of 4 instead of a max-pooling layer as in Ref.[16] This layer down samples the image resolution by a factor of 4 without sacrificing spatial information, as is the case with max pooling layers. Average pooling layers were also considered at this stage of the model to "blend" the reference and aberrated dots, but this implementation was saved for future work. Moreover, in contrast to Wang et al.,[16] the residual blocks in the current model had branches with convolutions of varying kernel sizes to capture features of different spatial scales. Adapting from Hu et al.,[14] kernel sizes of 3x3, 5x5, 7x7, and 9x9 are used to extract and represent those features. These features are then concatenated with the input and convoluted with a 1x1 filter with no stride. Batch normalization layers were included to make the distribution

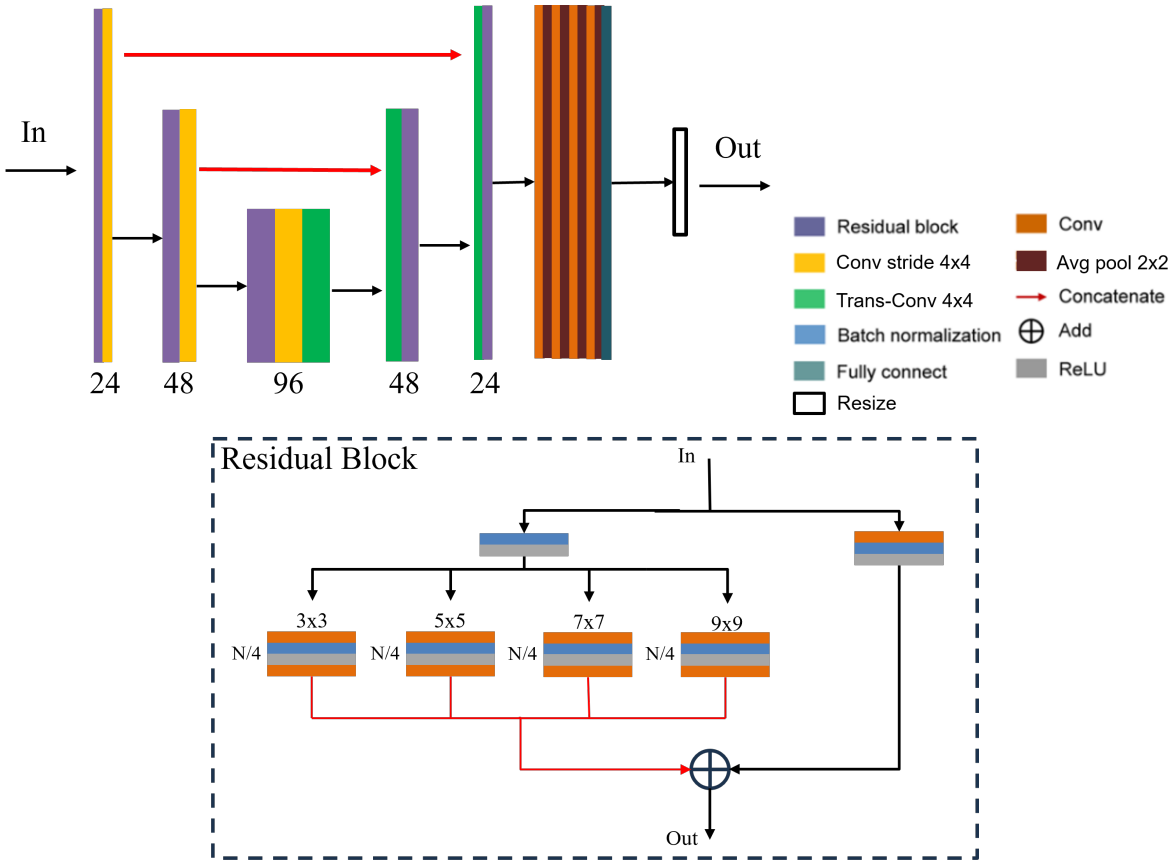

Figure 3. Current Res-UNet model architecture. Yellow blocks denote down-sampling with a 3 x 3 convolutional layer with a stride of 4. Green blocks represent up-sampling performed by transposed convolutional layers. Red arrows denote concatenation between the encoder and corresponding decoder layers. 'N' in 'N/4' denote the number of desired features in the residual block layer. Orange and red blocks indicate 3 x 3 convolutional down-sampling and average pooling layers, respectively, with a stride of 1. The final layer before the resizing layer is a fully-connected dense layer with a hyperbolic tangent activation function.

of the output features fall into a range with a larger activation function gradient to improve convergence speed. The encoder-decoder branch results in an image of the same size as the input image 489x489 pixels, however, the target wavefront is 43x43 so another set of 4 alternating convolution layers with average pooling layers was added to downsample the image. The intention of adding average pooling layers after the convolutional layers is to average out adjacent pixels at the finer resolution. Finally, a fully-connected dense layer was used to connect the output of the ResU-Net to the target wavefront accompanied by a resizing layer to coerce the output to the desired 50 x 50 size. As the resulting wavefront can take on values in the negative range, a hyperbolic tangent activation function was used for the fully connected dense layer to improve performance.

The model was implemented using Tensorflow 2.10 in Python 3.9.16, with the Adaptive Moment Estimation (ADAM) optimizer set at a learning rate of 0.0001. Hyperparameter optimization was carried out using the Keras-tuner library's Hyperband class, which randomly varied parameters such as the number of features in each block of the encoder/decoder, the number of downsampling layers, the stride size for downsampling, and the learning rate. The optimizer then ran each set of parameters for 2 epochs to determine the optimal configuration. The results showed that the optimal number of downsampling layers was 3, which was also the maximum value tested, with a stride of 4. The number of features to pull from the first convolution was 24 and this was doubled with each downsampling. The L2 Norm of the error between the target wavefront and the network output was used to assess training performance.

$$R(h|D) = \frac{1}{N^2} \sum_{k=1}^{N} \sum_{j=1}^{N} |h(x_{kj}) - W_{k,j}|^2 \tag{1}$$

where N is the size of one dimension the wavefront array, h(x) is the model output, W is the target wavefront, and k and j denote the spatial points in the image.

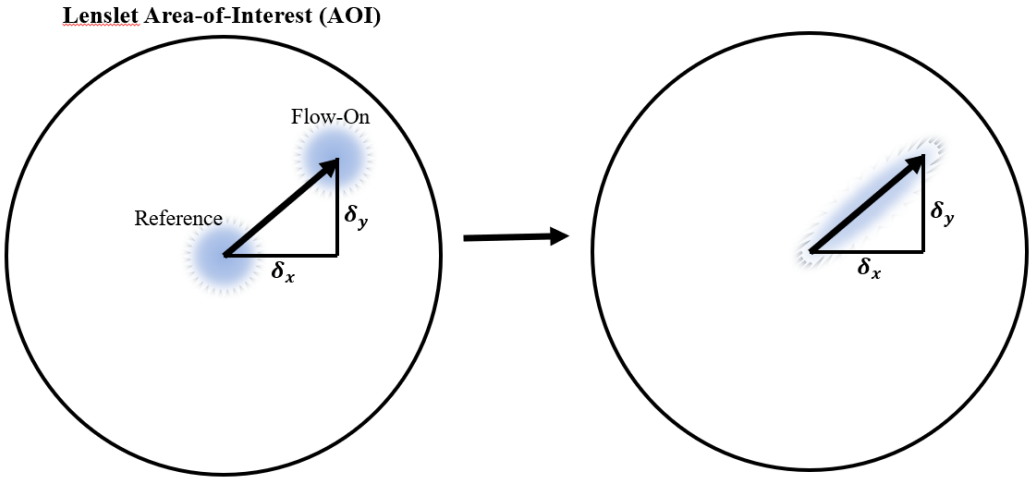

Figure 4. Example AOI dot with added reference (left). With convolutional layers and downsampling, the two dots blend together to form a feature (right) that uniquely describes the measured deflection.

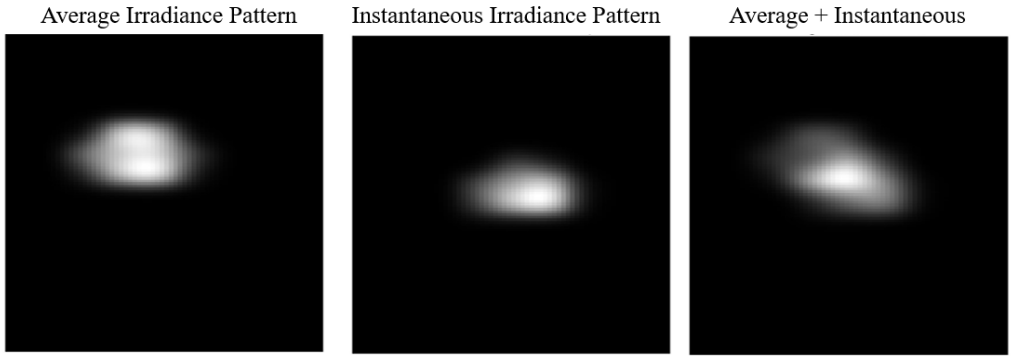

Figure 5. Sample SHWFS irradiance pattern depicting the time-series averaged dot (right) added to the aberrated dot (center) within an AOI to produce a feature (left) representative of local wavefront aberration.

3. DATASET, PREPROCESSING, AND HARDWARE

The model was trained using experimental data collected at the University of Notre Dame's White Field Mach 0.6 wind tunnel, where a collimated beam was propagated through the wake behind a flat plate and collected on a SHWFS camera. The freestream flow had a Mach number of 0.4. The camera sampled at 40 kHz with an exposure of 0.38 μs. The lenslet array had a pitch of 0.3 mm and a focal length of .0382 m. The SHWFS images were processed using the standard least-squares reconstruction processing algorithm.[7] The input consisted originally of images with a resolution of 489x489, whereas the target consisted of wavefronts with a size of 50x50. The images were then normalized to be on the order of 0 to 1 and were converted to half-precision floating-point values or float16. The wavefronts were converted from units of micrometers to scale between -1 and 1 and were also

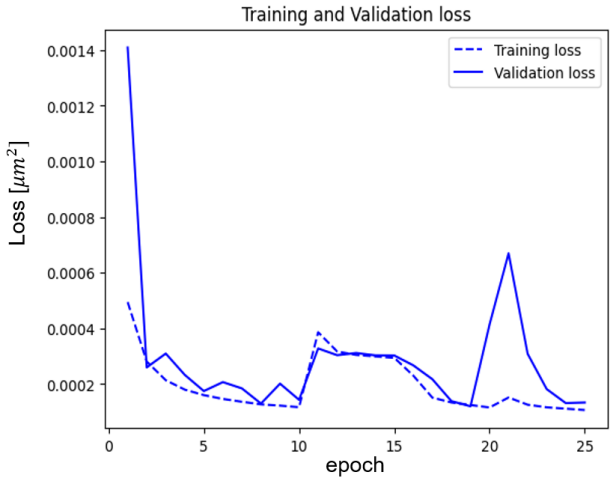

Figure 6. Validation and training loss vs epoch. Loss defined as mean square error between target and predicted wavefronts.

converted to float16. The image-wavefront pairs were then randomized. A dataset of 20000 images was trained using a 75-25 training–validation split where 1000 frames were reserved for testing from the training-validation. Training was performed in Google Colaboratory, a cloud-based system that has 12.7 GB of CPU RAM and 15 GB of GPU RAM. Training the network with the current configuration space took approximately 5 hours and future work aims to increase processing speed.

4. RESULTS

The model was trained for a total of 25 epochs, and the training and validation curves are shown in Fig. 6. After model training and validations, 1000 frames were used to test the accuracy of the model on data that had not been used during training. An example result is shown in Fig. 7. Here the wavefront reconstructed using the standard reconstruction algorithm is shown on the left. While the prediction visually looks similar the the target, the RMS of the absolute error was slightly less than half of the RMS of the wavefront, taking on these values near the boundary of the aperture. The extreme value discrepancy may be caused by the lenslets at the edge of the aperture, which can be very sensitive to noise, leading to erroneous predictions. While the RMS of the sample prediction shown in Fig. 7 deviated approximately 11% from the target, the RMS of the absolute error, averaged over the entire testing dataset of 1000 frames deviated from the frame-averaged RMS of the target by 43%. Furthermore, the peak-to-valley of the absolute error was approximately 44% of the peak-to-valley of the target wavefront. Other works[15,16] found the average RMS error between the target and network output was around 1% of the RMS of the wavefront. Particularly, Wang et al.[16] also uses the peak-to-valley distance as another metric and reported a peak-to-valley error that was less than 3% of the peak-to-valley of the target wavefront. The goal of this research is to get comparable results, however, from preliminary testing, the model produces significant errors. One potential reason for this is that the data used to train the model is experimental data, whereas other works[10,14,16] used simulated data. The effects of noise, misalignment between SHWFS dot patterns and reconstructed wavefronts, and various other parameters all have the potential to significantly decrease the effectiveness of training. Another source of the low accuracy most certainly stems from an under-training of the network; only 46000 sample data points were used, whereas other works use many more sample points.[14] Additionally, due to memory constraints on the CPU and GPU, training was performed in chunks of 5000 frames. The result of this training is evident in behavior of the the training-validation loss curves shown in Fig. 6, where the loss abruptly increases when a new chunk of the dataset is trained.

While the current model closely follows other models,[16] the goal was to make this model simpler so that it would require less training and have faster computation in a closed-loop AO application. The objective of this work was to increase the speed of computation compared to the standard centroiding and wavefront reconstruction method. The current status of Notre Dame's AO system has just the centroiding process occuring on the order of 100Hz. With wavefront reconstruction, the throughput of the current system clocks in at just over 10Hz.

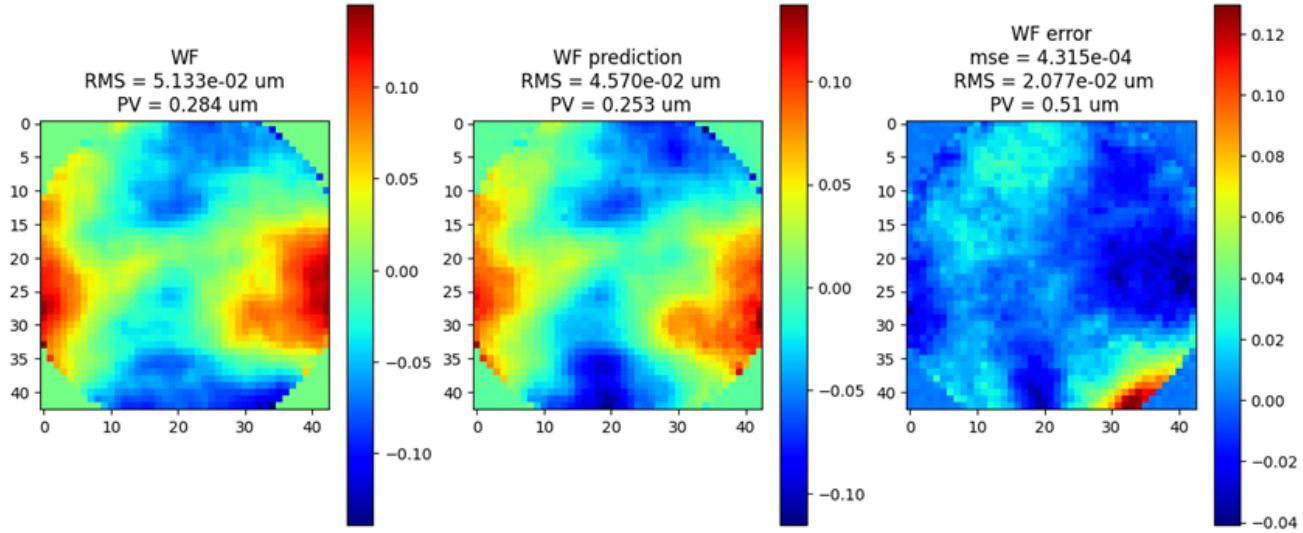

Figure 7. Sample frame of least-squares reconstructed wavefront(left), network predicted wavefront (center), and absolute error between left and center plots (right). PV indicates peak-to-valley. Colorbars show wavefront value in microns.

The network presented here was able to process data at about 120 Hz, which, while much slower than other implementations,[16] is still and order of magnitude faster than the current system.

These preliminary results are promising, as the model was able to predict wavefronts that resembled the target wavefront from experimental data. However much work is needed to improve accuracy. Future work entails improving the robustness of the model in order to more easily extract features from the SHWFS dot patterns and use those to perform the zonal reconstruction method. Additionally, training on a larger and much more diverse set of data from different experimental conditions and test setups will be conducted to improve generalizability and model accuracy. Once an acceptable accuracy is reached, steps will be taken to simplify the model to reduce the computational cost and increase the throughput speed while retaining accuracy. Finally, the model will be implemented into the existing AO system and its performance evaluated in an experimental capacity.

5. CONCLUSION

In this study a Res-UNet architecture CNN was developed and trained with the goal of implementing a faster method for wavefront reconstruction from SHWFS measurements. The model was trained using experimental data collected by measuring the aero-optical environment of a flat-plate wake in Mach 0.4 flow. This model was found to process wavefronts at speeds around 10x that of the University of Notre Dame's AO system centroiding-based approach. The specialized approach to image preprocessing for zonal reconstruction is also evaluated and shown to produce results that closely match the qualitative behavior of the least-squares reconstructed wavefront. However, while these predicted wavefronts qualitatively look promising, the accuracy requires improvement with the absolute error between the target and predicted wavefronts have a mean spatial RMS error of 2.7 μm, which was on the same order of magnitude as the spatial RMS of the target wavefront. Future work will aim to reduce the size of the model in order to maximize processing speed, while improving accuracy as well as study the limits of generalizability of this neural network-based approach.

ACKNOWLEDGMENTS

The authors would like to acknowledge the contribution of Luke Butler, whose experimental data was used for training the model in this work.

REFERENCES

[1] Jumper, E. J. and Gordeyev, S., "Physics and measurement of aero-optical effects: Past and present," *Annual Review of Fluid Mechanics* **49**, 419–441 (Jan. 2017).

[2] Wizinowich, P. L., Mignant, D. L., Bouchez, A. H., Campbell, R. D., Chin, J. C. Y., Contos, A. R., van Dam, M. A., Hartman, S. K., Johansson, E. M., Lafon, R. E., Lewis, H., Stomski, P. J., Summers, D. M., Brown, C. G., Danforth, P. M., Max, C. E., and Pennington, D. M., "The w. m. keck observatory laser guide star adaptive optics system: Overview," *Publications of the Astronomical Society of the Pacific* **118**, 297–309 (Feb. 2006).

[3] Duffin, D. A. and Jumper, E. J., "Feedforward adaptive-optic correction of aero-optical aberrations caused by a two-dimensional heated jet," *AIAA Journal* **49**(6), 1283–1291 (2011).

[4] Nightengale, A. M. and Gordeyev, S., "Shack-hartmann wavefront sensor image analysis: a comparison of centroiding methods and image-processing," *Optical Engineering* **52(7)**, 071413 (2013).

[5] Fried, D., "Least-square fitting a wave-front distortion estimate to an array of phase-difference measurements," *J. Opt. Soc. Am.* **67**, 370–375 (1977).

[6] Hudgin, R., "Wave-front reconstruction for compensated imaging," *J. Opt. Soc. Am.* **67**, 375–378 (1977).

[7] Southwell, W., "Wave-front estimation from wave-front slope measurements," *J. Opt. Soc. Am.* **70**, 998–1006 (Aug 1980).

[8] Mochi, I. and Goldberg, K. A., "Modal wavefront reconstruction from its gradient," *Appl. Opt.* **54**, 3780–3785 (Apr 2015).

[9] Genberg, V., Michels, G., Doyle, K., Bury, M., and Sebastian, T., "Computational methods to compute wavefront error due to aero-optic effects," in [*Optical Modeling and Performance Predictions VI*], Kahan, M. A. and Levine, M. B., eds., **8840**, 88400A, International Society for Optics and Photonics, SPIE (2013).

[10] Swanson, R., Lamb, M., Correia, C., Sivanandam, S., and Kutulakos, K., "Wavefront reconstruction and prediction with convolutional neural networks," in [*Adaptive Optics Systems VI*], Close, L. M., Schreiber, L., and Schmidt, D., eds., **10703**, 107031F, International Society for Optics and Photonics, SPIE (2018).

[11] Wilcox, C. C., Radosevich, C. J., Healy, K. P., Tuffli, A. L., Agena, B. D., Spencer, M. F., and III, D. J. W., "Digital holography wavefront sensing with a supersonic wind tunnel ," in [*Holography: Advances and Modern Trends VI*], Fimia, A., Hrabovský, M., and Sheridan, J. T., eds., **11030**, 110300L, International Society for Optics and Photonics, SPIE (2019).

[12] Kong, F., Cegarra Polo, M., and Lambert, A., "Fpga implementation of shack-hartmann wavefront sensing using stream-based center of gravity method for centroid estimation," *Electronics* **12**(7) (2023).

[13] Suárez Gómez, S. L., González-Gutiérrez, C., Díez Alonso, E., Santos Rodríguez, J. D., Sánchez Rodríguez, M. L., Carballido Landeira, J., Basden, A., and Osborn, J., "Improving adaptive optics reconstructions with a deep learning approach," in [*Hybrid Artificial Intelligent Systems*], de Cos Juez, F. J., Villar, J. R., de la Cal, E. A., Herrero, Á., Quintián, H., Sáez, J. A., and Corchado, E., eds., 74–83, Springer International Publishing, Cham (2018).

[14] Gu, H., Zhao, Z., Zhang, Z., Cao, S., Wu, J., and Hu, L., "High-precision wavefront reconstruction from shack-hartmann wavefront sensor data by a deep convolutional neural network," *Measurement Science and Technology* **32**, 085101 (May 2021).

[15] Kaji, S. and Kida, S., "Overview of image-to-image translation by use of deep neural networks: denoising, super-resolution, modality conversion, and reconstruction in medical imaging," (2019).

[16] Wang, K., Zhang, M., Tang, J., Wang, L., Hu, L., Wu, X., Li, W., Di, J., Liu, G., and Zhao, J., "Deep learning wavefront sensing and aberration correction in atmospheric turbulence," *PhotoniX* **2** (June 2021).