Optimizing branch-point detection using machine learning

Thomas E. DeFoor*[a], Matthew Kalensky[b], and Stanislav Gordeyev[a]

[a]Department of Aerospace and Mechanical Engineering, University of Notre Dame, Notre Dame, Indiana, USA
[b]Integrated Engagement Systems Department, Naval Surface Warfare Center Dahlgren Division, Dahlgren, Virginia, USA

ABSTRACT

In this paper, we leverage wave-optics simulations to explore a machine learning approach for identifying the locations of branch points in Shack–Hartmann wavefront sensor (SHWFS) data. We train a convolutional neural network to receive a single SHWFS sub-aperture irradiance pattern image and estimate whether the pupil of that area of interest (AOI) contains at least one branch point. We test the model alongside a branch-point detection approach that relies on the beam spread ($D_{4\sigma}$) of individual SHWFS irradiance patterns to determine whether or not a branch point is present within the AOI. We explore the performance of both the machine-learning-based and $D_{4\sigma}$-based branch-point identification approaches across a wide range of optical turbulence conditions, quantified by Rytov numbers ($\mathcal{R}$) ranging from 0 - 2, as well as in the presence of additive noise. To quantify our performance, we use various metrics commonly used for binary classifiers, including the F_1 score and the Receiver Operating Curve (ROC). By and large, we find that the machine learning model outperformed the $D_{4\sigma}$ approach. Implementing machine learning of this caliber into branch-point identification provides the basis for future studies exploring improved adaptive-optics compensation in regions of strong turbulence.

Keywords: branch points, hidden phase, machine learning, atmospheric optics

1. INTRODUCTION

As a laser beam propagates through the atmosphere, turbulence along the path results in the imposition of phase aberrations onto the beam. When the laser propagates over long path lengths or through strong turbulence, these phase aberrations give rise to constructive and destructive interference, which results in variations in irradiance across the beam. We refer to these turbulence-induced irradiance variations as scintillation. Scintillation has been studied by researchers for decades.[1–5] In 1992, Fried and Vaughn published a paper which showed that when scintillation produces regions of total destructive interference, branch points form in the phase function of the complex-optical field.[6] Branch points are singularities in the phase function, about which we observe a 2π circulation in phase. Branch points form in pairs and are connected by a 2π phase discontinuity which we refer to as a branch cut.

Since Fried and Vaughn's seminal paper, researchers have been interested in studying branch points and branch cuts for a number of reasons. First, branch points and branch cuts have punitive implications on adaptive-optics system performance. Researchers have shown that in strong scintillation conditions, adaptive-optics systems must correct for the phase field that results from branch points and branch cuts in order to achieve good system performance.[7–15] Second, branch points give us physical insight into the optical-turbulence environment that created them. For this reason, researchers have been interested in studying branch points as a means to improve our understanding and characterization of laser propagation through atmospheric turbulence.[16–21]

Provided we have a means to measure the complex-optical field of our laser beam, identifying branch points is trivial. As shown by Fried and Vaughn, by taking the derivative of the phase function of the complex-optical field and computing local contour integrals over the pupil, we can readily identify pupil locations where the contour integral yields values of $\pm 2\pi$; suggesting the presence of at least one branch point within the contour.[6]

*Further author information:
E-mail: tdefoor@nd.edu

Unconventional Imaging, Sensing, and Adaptive Optics 2025, edited by Jean Dolne, Santasri Bose-Pillai, Matthew Kalensky, Proc. of SPIE Vol. 13619, 1361918 · 2025 Published by SPIE
0277-786X · doi: 10.1117/12.3063685

By making the contour over which we take our integral small, and rastering these computations over all pupil locations, we are able to obtain a mapping of branch points as a function of pupil locations. However, this approach is predicated on having access to the full phase function of the complex-optical field. While some wavefront sensing approaches such as the self-referencing interferometer, digital holographic wavefront sensor, or phase-retrieval-based wavefront sensors enable such access, these approaches are often either not mature enough to implement at the system level or have practical limitations that have yet to be overcome.

The Shack–Hartmann wavefront sensor (SHWFS), on the other hand, is one type of wavefront sensor that has been used for adaptive optics and optical-turbulence characterization for decades. While generally regarded as a "reliable" wavefront sensor, this adjective is unfitting when the SHWFS is used in strong scintillation conditions. Specifically, scintillation imposes dynamic range challenges[7] and measurement biases due to non-uniform illumination across the pupil.[22–24] Furthermore, under normal operation (i.e., estimating phase gradients which are then used in a least-squares reconstruction algorithm), the SHWFS does not easily enable the identification of branch points.

Given the SHWFS' profuse utilization across many fields (astronomy, free-space optical communications, directed energy) and ease of use, there has been interest in using the SHWFS to detect branch points. Since the SHWFS provides estimates of phase gradients, we can use the circulation of the measured SHWFS phase gradients to identify branch points.[6,25] This approach was explored in detail by Chen.[26] Wild and Le Bigot[27] introduced an alternative SHWFS-based branch-point identification approach referred to as the branch-point potential method. This approach was later demonstrated in experiment by Murphy and Dainty[28] and shown to give improved performance over the circulation of phase gradients approach. In 2000, Tyler demonstrated that slope discrepancy could be used to identify branch points.[29] Lastly, Kalensky[30] demonstrated that when branch points are present within a SHWFS' sub-aperture pupil, their presence prominently affects the resultant SHWFS irradiance pattern. The presence of optical-phase discontinuities and singularities such as branch points induces appreciable far-field irradiance pattern spreading and smearing. Through this, if the measured far-field irradiance distribution is wider than the nominal, un-aberrated case, then it can be assumed that a branch point lies in the phase at that spatial region. As a result, it was shown that the second-moment statistics, quantified by $D_{4\sigma}$, of individual SHWFS irradiance patterns could also be used to identify branch points. This approach was later compared to the circulation of phase gradients approach in Refs. [31–33].

The past few years have seen the integration of machine learning approaches into SHWFS processing. The scope of applications ranges from assisting centroid measurements[34] to augmenting full OPD field reconstruction.[35–41] Additionally, DeFoor et al.[42] demonstrated the utility of an encoder-style convolutional neural network (CNN) in estimating the magnitude of discontinuities in the incident optical phase using SHWFS sub-aperture irradiance pattern images alone in both simulation and experiment. The combination of these applications showcases the potential for extending the applicability of the SHWFS into historically challenging optical sensing environments — as it pertains to this paper, we consider strong scintillation conditions.

Building off of these prior developments, we introduce a new approach for identifying branch points from SHWFS data where we leverage machine learning to identify whether or not a branch point falls within the pupil of each SHWFS sub-aperture. In order to both generate training data to create the model as well as to explore the efficacy of this approach, we utilize wave-optics simulations. Given the similarity of this approach to the beam-spread approach (which utilizes $D_{4\sigma}$ thresholding) introduced in Ref. [30], we compare the performance of our machine-learning-based branch-point identification model with the beam spread method. The following paper outlines our approach. First, section 2 discusses the wave-optics simulation model used to both generate the training and test our branch-point identification approach. In Sec. 3, we introduce our machine learning model. In Sec. 4, we present and discuss our results. Finally, in Sec. 5, we summarize and conclude the paper.

2. WAVE-OPTICS SIMULATIONS

In this section, we discuss the wave-optics simulations that we used to model laser beam propagation through atmospheric turbulence. Here, we modeled a spatially filtered point-source beacon, and allowed this point-source beacon to spherical-wave propagate from the beacon plane to our receiver plane. We accomplished this using angular-spectrum propagation. We modeled atmospheric turbulence using 20 evenly spaced phase screens

generated using the modified von Karman spectrum, where we defined the inner scale of turbulence, ℓ_0, to be 1 mm and the outer scale of turbulence, L_0, to be 100 m. We generated our phase screens using 2048×2048 points subtending 1×1 m^2. For all simulations, we defined the propagation path to be $L = 2.5$ km, the wavelength to be $\lambda = 1$ µm, and the strength of turbulence to be uniform along the propagation path, to emulate a horizontal propagation path. In order to explore different optical-turbulence conditions, we varied the index of refraction structure constant, C_n^2, used to create the screens. In this paper, we quantify the strength of our optical-turbulence conditions using the spherical-wave Rytov number, $\mathcal{R}$, defined as[43]

$$\mathcal{R} = 0.563k^{7/6} \int_0^L C_n^2(z) \left(\frac{z}{L}\right)^{5/6} (L-z)^{5/6} \, \mathrm{d}z, \tag{1}$$

where k is the laser wavenumber. The Rytov number is an indication of scintillation, which is why it is an appropriate metric to characterize our trade space.

After our point-source beacon propagates from the beacon plane, through our atmospheric turbulence phase screens, and reaches our receiver plane, we apply an aperture transmittance function, and collimate the complex-optical field. For all simulations, we defined our aperture diameter to be $D = 0.5$ m. In general, these simulations followed the procedures outlined in Refs. [44–47]. Using the phase function of the complex-optical field obtained after propagation, we apply the circulation of phase gradients approach to identify any branch points contained within our pupil. We treat the branch points that we identify during this step as our known branch points.

In order to model a SHWFS, we generate a mask of thin-lens transmittance functions which represent an array of lenslet sub-apertures. We apply this SHWFS mask to our received complex-optical field. Subsequently, we perform angular-spectrum propagation to obtain a simulated SHWFS image. For the simulations that we conducted here, we defined 64×64 circular sub-apertures subtending across our 1×1 m^2 simulation domain. This resulted in SHWFS sub-apertures with a diameter of 15.6 mm. The focal length of the thin-lens transmittance functions used to create our SHWFS mask was $f = 25$ m. We present example simulated SHWFS imagery in the first row of Fig. 1 and the associated "hidden" phase component in the second row of Fig. 1. Here, each column represents example data when $\mathcal{R} = 0,\ 0.25,\ 0.75$ from left to right, respectively. From this figure, it is clear that increasing $\mathcal{R}$ increases the number of branch points present across the pupil.

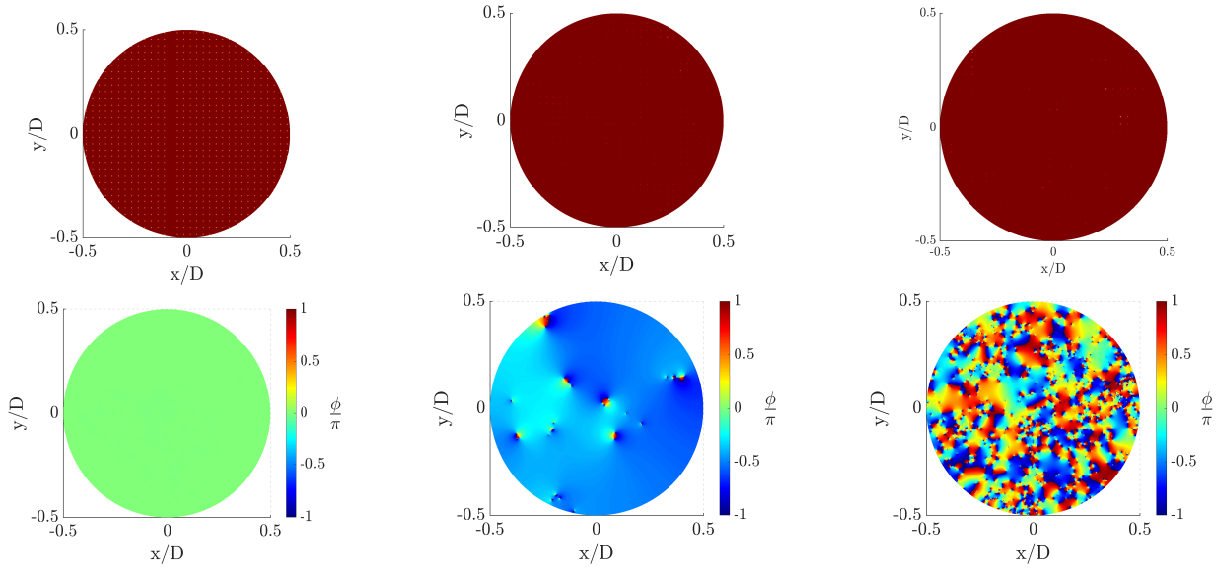

(a) $\mathcal{R} = 0$ **(b)** $\mathcal{R} = 0.25$ **(c)** $\mathcal{R} = 1.5$

Figure 1: Simulated SHWFS images (top row) and accompanying hidden phase (bottom row) for various $\mathcal{R}$ values.

To train the machine learning model discussed in Sec. 3, we generated SHWFS imagery for different optical-turbulence conditions where we varied $\mathcal{R}$ from 0.25 - 0.65. For each $\mathcal{R}$, we generated 35 independent realizations.

We trained our model on individual SHWFS sub-aperture irradiance pattern images, and as a result, there was no need to extend the training dataset into higher or lower $\mathcal{R}$ values. This is because our optical-turbulence conditions, spanning from $\mathcal{R}$ = 0.25 - 0.65, already offered plenty of diversity. However, the dataset used for testing the machine learning model and $D_{4\sigma}$ approach consisted of SHWFS images spanning $\mathcal{R}$ values from 0 - 2, with a total of 200 independent realizations for each $\mathcal{R}$ to ensure statistically meaningful results. More information on the training and testing data will be discussed in the next section.

3. MACHINE LEARNING MODEL

In this section, we describe the architecture of the proposed machine learning model used for branch-point identification. The model used for this application follows an encoder-style convolutional neural network (CNN) structure. This specific model type was used as it has been proven resilient to noise and useful in numerous image processing and classification applications.[42,48–50] The model was developed to receive a single, SHWFS sub-aperture irradiance pattern image and classify it as having zero or at least one branch point present in the spatial field associated with the image. The specific structure used in this analysis is presented in Fig. 2. Through this, we define our positive class to be the identification of a branch point and our negative class to be the designation of no branch point located in the spatial location of the incident phase corresponding to the input image. The model used here consists of 4 convolutional layers, each with a 3×3 kernel that extracts the features. Additionally, each layer is separated by a combination of batch normalization layers and Relu activations to stabilize training and introduce controlled non-linearity. Additional layers include max and global average pooling to downsize the data before being sent to a fully connected dense layer with a softmax activated output consisting of a two-class confidence score identifying whether the input image covers no (0) or at least one (1) branch point, spatially. A detailed description including layer sizes can be found in Table 1.

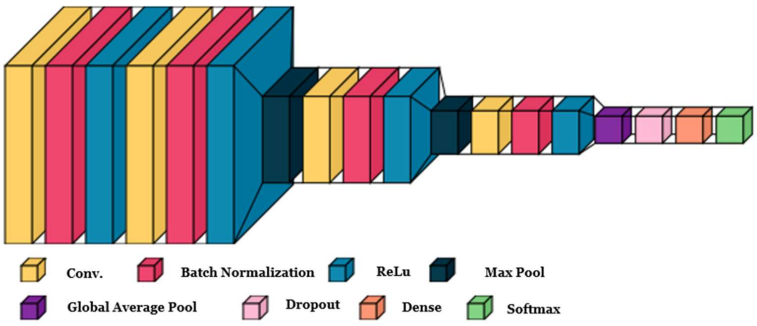

Figure 2: CNN model architecture.

The model was trained on a set of single sub-aperture irradiance patterns pulled from a full SHWFS image simulated in the fashion outlined in Sec. 2. To ensure a balanced training dataset, the number of images housing 0 branch points was limited to only 500 more than the number of images housing branch points, resulting in 42,922 images after an 80% training, 10% validation, and 10% testing, split of the data. Additionally, to introduce robustness, every training image was augmented by randomly adding x and y offsets between ± 2 pixels, spatial scaling between 0.8 and 1.2$\times$ the nominal image size, as well as additive, zero-mean Gaussian noise at variance levels of 0, 0.00015, and 0.0003. The training progress was monitored using categorical cross-entropy as its loss function to optimize the model classification.

Additional testing data was generated also following the method described in Sec. 2 to provide unseen data for more rigorous testing of the ML model. As the goal was to evaluate the performance of the machine learning model against the proven reliable beam spread thresholding methods, the test data conditions were expanded to $\mathcal{R}$ values from 0 - 2 and simulated additive Gaussian noise at variance levels of 0, 0.00015, and 0 (referred to as no, low, and high, noise levels). Additionally, as stated previously, 200 simulations for each $\mathcal{R}$ condition were generated to ensure statistical convergence. The results of this comparison are presented in Sec. 4.

Table 1: Summary of CNN architecture.

Layer #	Layer Type	Description
0	Image Input	$32 \times 32 \times 1$ image (input)
1	2-D Convolution	32 filters, 3×3, stride $[1, 1]$
2	Batch Normalization	Batch normalization
3	ReLU	ReLU activation
4	2-D Convolution	32 filters, 3×3, stride $[1, 1]$
5	Batch Normalization	Batch normalization
6	ReLU	ReLU activation
7	Max Pooling	2×2 pool, stride $[2, 2]$
8	2-D Convolution	64 filters, 3×3, stride $[1, 1]$
9	Batch Normalization	Batch normalization
10	ReLU	ReLU activation
11	Max Pooling	2×2 pool, stride $[2, 2]$
12	2-D Convolution	128 filters, 3×3, stride $[1, 1]$
13	Batch Normalization	Batch normalization
14	ReLU	ReLU activation
15	Global Avg. Pooling	2-D global average pooling
16	Dropout	Dropout with 30% rate
17	Fully Connected	2 neurons
18	Softmax	Softmax activation (output)

4. RESULTS AND DISCUSSION

In this section, we present the results from a comparison between the CNN machine learning model and the beam spread approaches for branch-point identification across various optical-turbulence conditions. To measure far field irradiance pattern spread, $\mathrm{D}_{4\sigma}$ was calculated for each SHWFS sub-aperture irradiance pattern area of interest (AOI). This beam spread metric determines a beam diameter equal to four standard deviations of the mean for the far-field irradiance pattern. $\mathrm{D}_{4\sigma\mathrm{x}}$ is calculated by,

$$D_{4\sigma_x} = 4\sqrt{\frac{\sum_{i=1}^{n}(x_i - \bar{x})^2 I_i}{\sum_{i=1}^{n} I_i}}, \tag{2}$$

where n is the total number of pixels in the sub-aperture irradiance pattern image, I_i is the irradiance value at each pixel location, x_i is the spatial x-coordinate value of the pixel, and $\bar{x}$ is the centroid pixel x-coordinate. Here the image pixels are raster-ordered to simplify calculation. A similar equation exists for $\mathrm{D}_{4\sigma_\mathrm{y}}$. If the overall spread $\left(D_{4\sigma} = \sqrt{D_{4\sigma_x}^2 + D_{4\sigma_y}^2}\right)$ is over some set threshold with respect to the un-aberrated case, the method flags the AOI as containing at least one branch point. It is often convenient to normalize $\mathrm{D}_{4\sigma}$ by the theoretical diffraction-limited spot size of the sub-apertures, $\mathrm{D}_{\mathrm{DL_{SA}}}$, which is calculated by,

$$D_{DL_{SA}} = 2\frac{\lambda f}{D_{SA}}, \tag{3}$$

where λ is the wavelength of light, f is the focal length of the sub-apertures, and D_{SA} is the diameter of the sub-aperture lenslet. In this study, we test three levels of thresholding fidelity: liberal ($\frac{D_{4\sigma}}{D_{DL_{SA}}} = 2.5$), moderate ($\frac{D_{4\sigma}}{D_{DL_{SA}}} = 2.75$), and conservative ($\frac{D_{4\sigma}}{D_{DL_{SA}}} = 3.0$).

Detection results for the three noise levels are presented in Fig. 3, where subplots (a)–(c) represent increasing additive noise. Here, total branch points (dark blue bar) corresponds to the total number of branch points detected using the circulation of phase gradient in the computed hidden phase; this is taken as the ground truth. The number of "obtainable" branch points (orange bar) corresponds to the total number of branch points that can be detected using the "at least one per AOI" approach. For example, if a SHWFS sub-aperture AOI spatially

covers 4 branch points, this counts as one. By this logic, the maximum possible number of branch points that can be detected is the sum total number of SHWFS sub-apertures across our pupil that satisfy the binary condition of whether or not at least one branch point exists within the sub-aperture pupil. As a result, perhaps the most immediately noticeable feature in Fig. 3 is the large disparity between the total number of branch points and every other bar on the figure. Our ability to identify branch points is indeed limited to the spatial resolution of our SHWFS. "Net" branch points were calculated by taking the number of true positives (TP) and subtracting the number of false positives (FP). A true positive means that the method correctly detected the presence of a branch-point-affected AOI and a false positive means that a branch point was detected in an AOI when there was not one present. Presenting net detected branch points in this fashion penalizes over prediction and is more indicative of method performance rather than just presenting TPs.

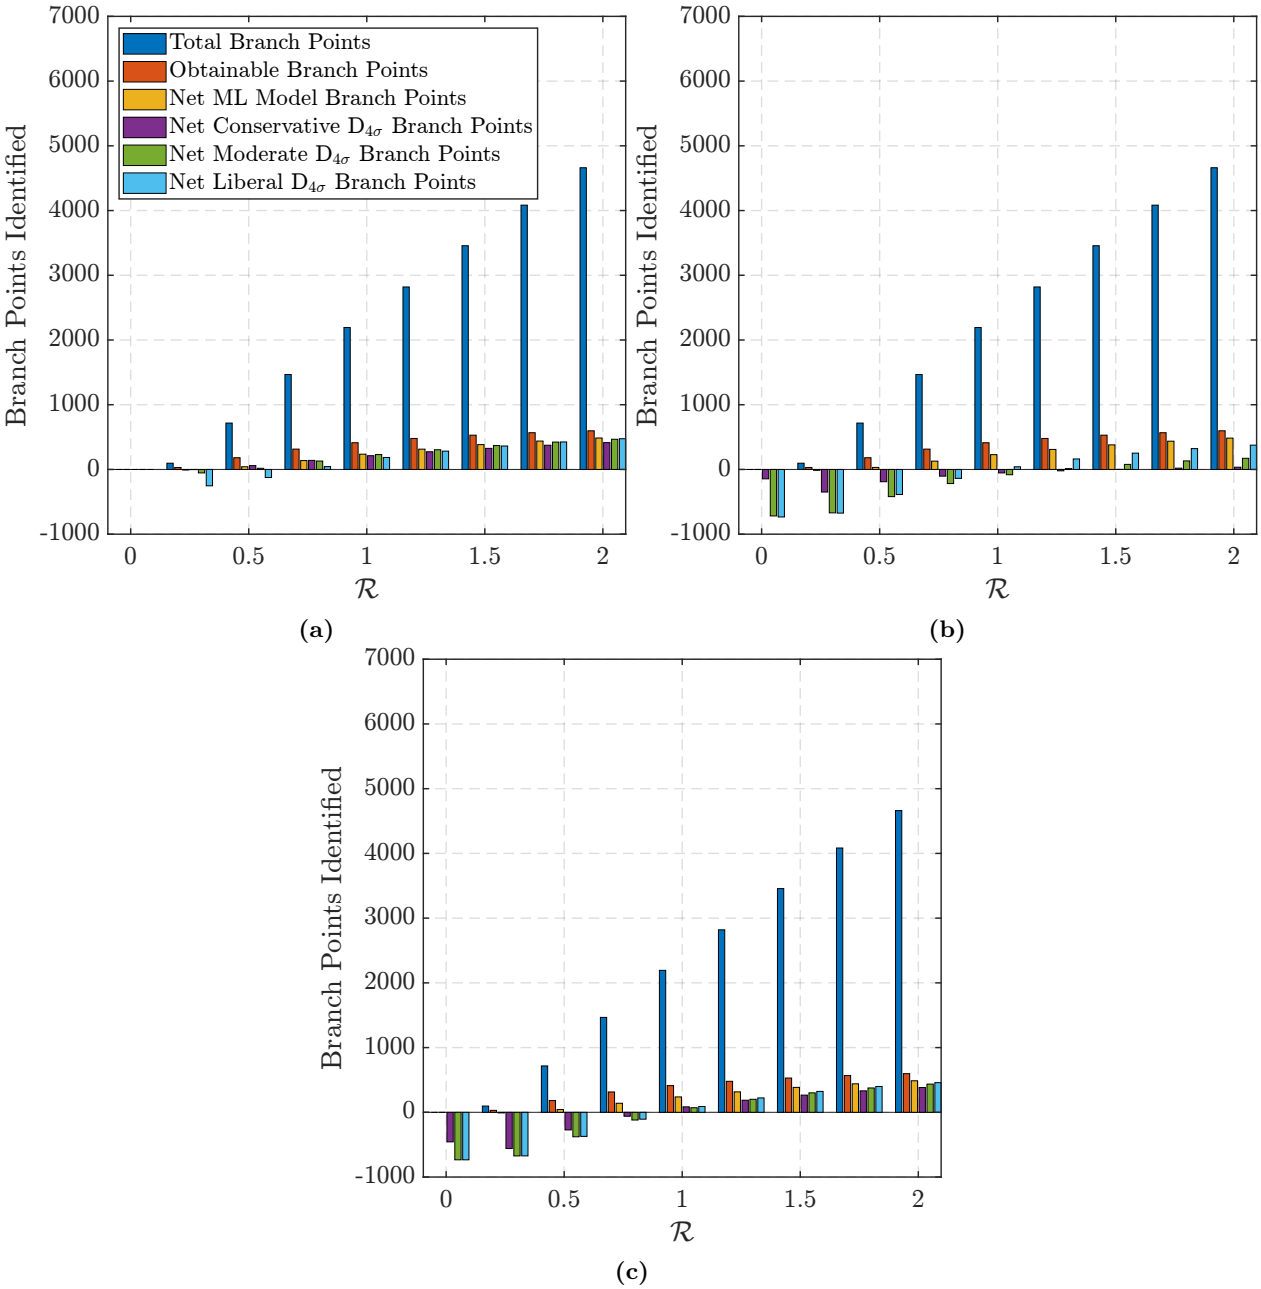

Figure 3: Detection results across all $\mathcal{R}$ values for (a) no, (b) low, and (c) high noise levels

It is clear in Fig. 3a that all methods behave similarly across all turbulence conditions. Taking this plot at face value would suggest no advantage for the use of the machine learning model (yellow bar) over the $D_{4\sigma}$ thresholding methods (purple, green, and cyan bars for conservative, moderate, and liberal thresholding, respectively). The performance discrepancy only slightly improves in the low (Fig. 3b) and high (Fig. 3c) additive noise cases. In these cases, there are more over-predictions in the lower $\mathcal{R}$ conditions where the addition of noise is causing excessive irradiance pattern spreading.

While useful, these plots do not tell the entire story. Figures 4, 5, & 6 show the results of true positive rate (TPR), false positive rate (FPR), and false negative rate (FNR) as a function of $\mathcal{R}$ for the cases of no, low, and high additive noise, respectively. We present TPR as the number of TP normalized by the average number of "obtainable" branch points across all 200 independent realizations for each $\mathcal{R}$ case. The FPR was calculated by normalizing the number of FP by the number of AOIs which did not have a branch point. The FNR was calculated by normalizing the number of false negatives by the number of "obtainable" branch points.

Presented in Fig. 4 are the TPR, FPR, and FNR results as a function of $\mathcal{R}$ for the no additive noise case. Focusing on the $D_{4\sigma}$ thresholding in Fig. 4a, the ordering of methods is as expected. The liberal thresholding (blue diamonds) has the highest TPR, followed by the moderate (green crosses), and trailed by the conservative (red squares). Both the machine learning (ML) model (black asterisks) and liberal thresholding beam-spread approach have high reported values of TPR, signifying that these methods are best at finding branch-point-affected AOI's as they exist. However, the TPR metric does not penalize the inclusion of false positives (FP) and missed detections (FN), which are plotted in Figs. 4b & 4c, respectively. Again, focusing on the $D_{4\sigma}$ thresholding methods in Fig. 4b, the liberal thresholding has the highest rate of false positives due to its low threshold value, conservative thresholding has the lowest rate as its threshold was chosen to reduce false positives, and the moderate thresholding results lie in the middle between the two. The ML model exhibits similar behavior in FPR as the liberal thresholding. This trend is also evident in the FNR curve plotted in Fig. 4b, where the liberal thresholding and ML model have the lowest FNR across $\mathcal{R}$ values, followed by moderate and conservative results. Although exhibiting slight improvement in these plots, the ML model performs relatively closely to the liberal thresholding method.

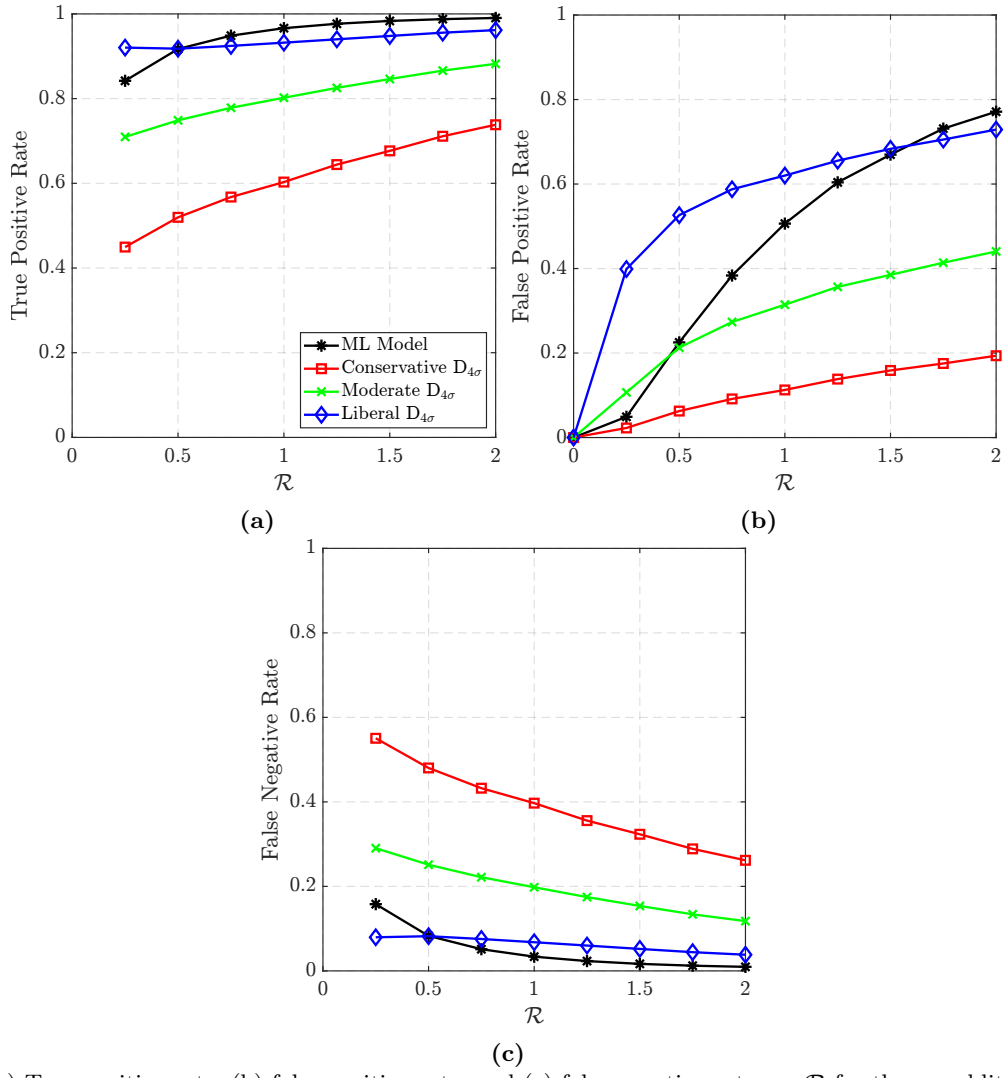

Figure 4: (a) True positive rate, (b) false positive rate, and (c) false negative rate vs. $\mathcal{R}$ for the no additive noise case.

Figures 5 & 6 present TPR, FPR, and FNR results as a function of $\mathcal{R}$ for the low and high additive noise cases. It is worth noting that the $D_{4\sigma}$ thresholding values were adapted to address the addition of noise and to provide a more rigorous benchmark for ML model performance characterization. These new threshold values were chosen in the same fashion as for the no noise case: a liberal threshold chosen to maximize detections, a conservative threshold chosen to minimize false positives, and a moderate threshold chosen to be the middle ground between the two. Conversely, nothing was changed about the ML model used. For the low noise case presented in Fig. 5, the tested thresholds were $\frac{D_{4\sigma}}{D_{DL_{SA}}} = 4.25$ for liberal, $\frac{D_{4\sigma}}{D_{DL_{SA}}} = 4.5$ for moderate, and $\frac{D_{4\sigma}}{D_{DL_{SA}}} = 4.75$ for conservative. The values used for the high noise case presented in Fig. 6 were 4.75, 5, and 5.25, respectively. However, despite the authors' best attempts to extend the usefulness of the $D_{4\sigma}$ approach into simulated noise conditions, the ML model outperforms all the beam-spread approaches and exhibits the same behavior across all $\mathcal{R}$ conditions for TPR, FPR, and FNR as shown in the no noise case in Fig. 4. Additionally, from Figs. 5 & 6, we see that the moderate and liberal thresholding results are nearly flagging every AOI as containing at least one branch point; as shown by the high TPR (Figs. 5a & 6a) and FPR (Figs. 5b & 6b) and low FNR (Figs. 5c & 6c).

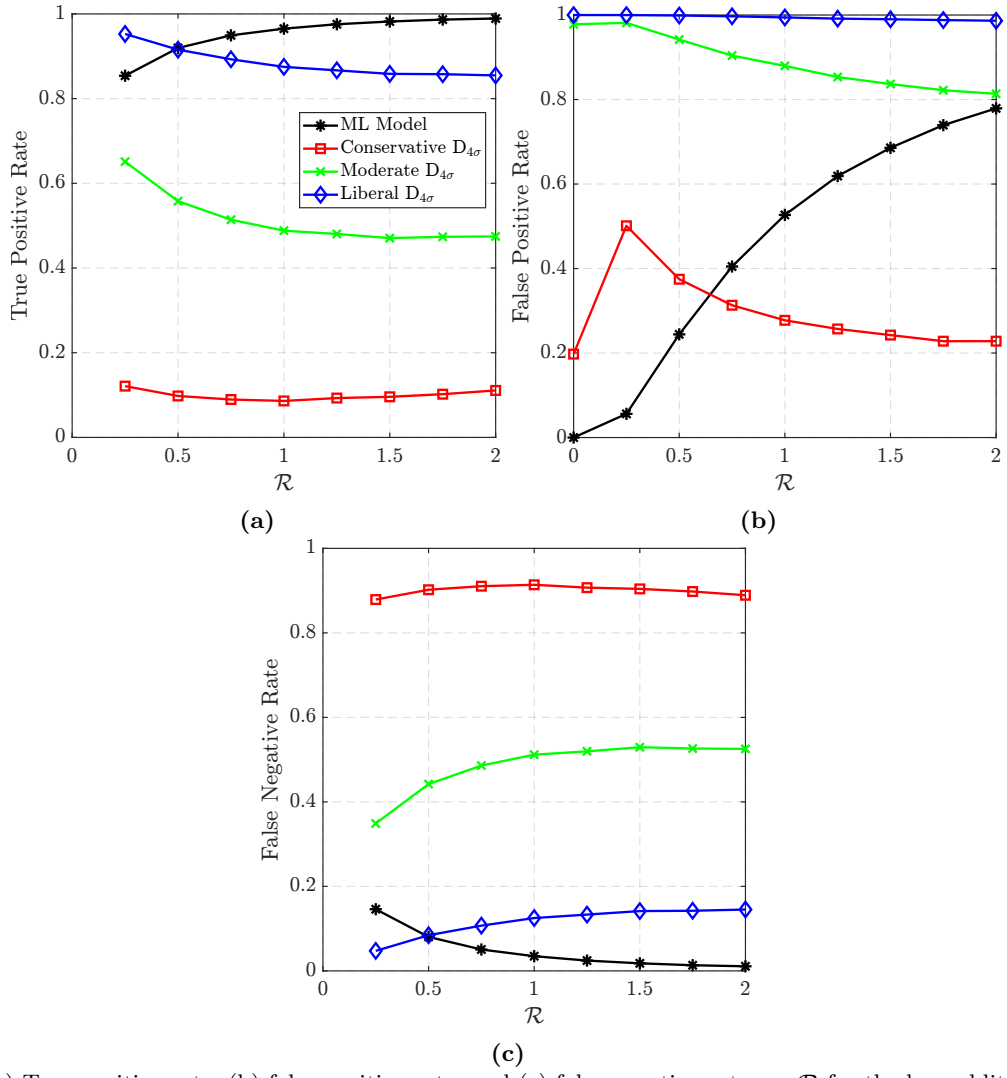

Figure 5: (a) True positive rate, (b) false positive rate, and (c) false negative rate vs. $\mathcal{R}$ for the low additive noise case.

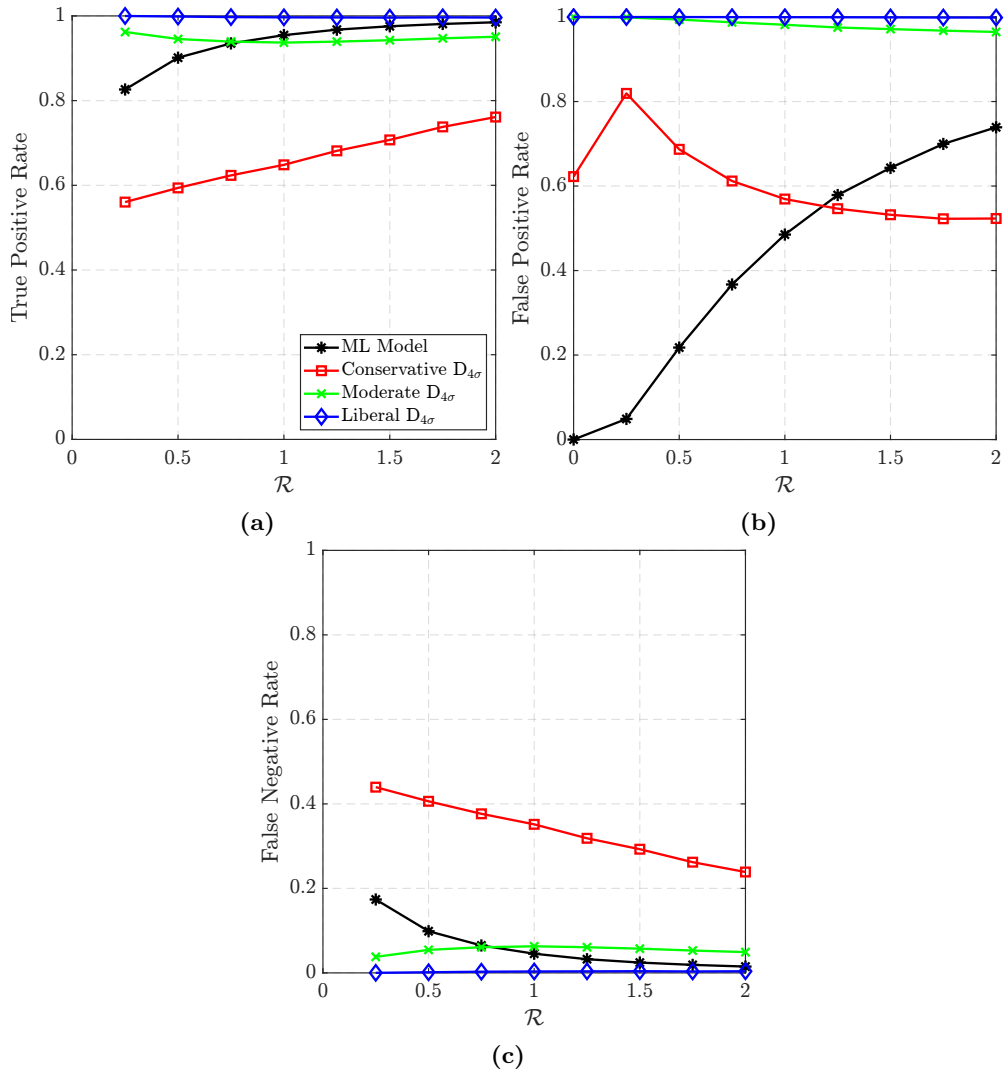

Figure 6: (a) True positive rate, (b) false positive rate, and (c) false negative rate vs. $\mathcal{R}$ for the high additive noise case.

However, the TPR, FPR, and FNR presented in Figs. 4-6, although useful, do not fully describe the reliability and generalizability of the applied identification methods as they are rate based statistics and calculated at single operating points ($\mathcal{R}$). Accordingly, we now turn to two metrics: the F_1 score and the Receiver Operating Curve (ROC). First, the F_1 score quantifies how accurate the model is by representing the reliability of the methods' prediction. It is calculated as the harmonic mean between precision and TPR (also often referred to as recall).[48] Precision quantifies the accuracy of branch-point detections as

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}. \tag{4}$$

The F_1 score, given by

$$F_1 = 2 \cdot \frac{\text{Precision} \times \text{TPR}}{\text{Precision} + \text{TPR}}, \tag{5}$$

is maximized when both precision and TPR are high, signifying a good balance between the two and a reliable model. Figure 7 presents F_1 score as a function of $\mathcal{R}$ for all three levels of noise. It is clear in these plots that the ML model outperforms the $D_{4\sigma}$ methods, proving to have a better balance between correct detections while not having too many missed detections. Additional focus should be paid to Fig. 7a. Here we see that the ML

model has the highest F_1 score across all $\mathcal{R}$ values, while the liberal has the lowest score. This is contrary to the results presented in Fig. 4, where the ML model and liberal thresholding method revealed similar performance. The plots of TPR, FPR, and FNR are single point, rate-based metrics and represent raw count statistics and therefore, are not fully indicative of the overall fidelity of the model. For these reasons, the F_1 score serves as a more reliable performance metric.

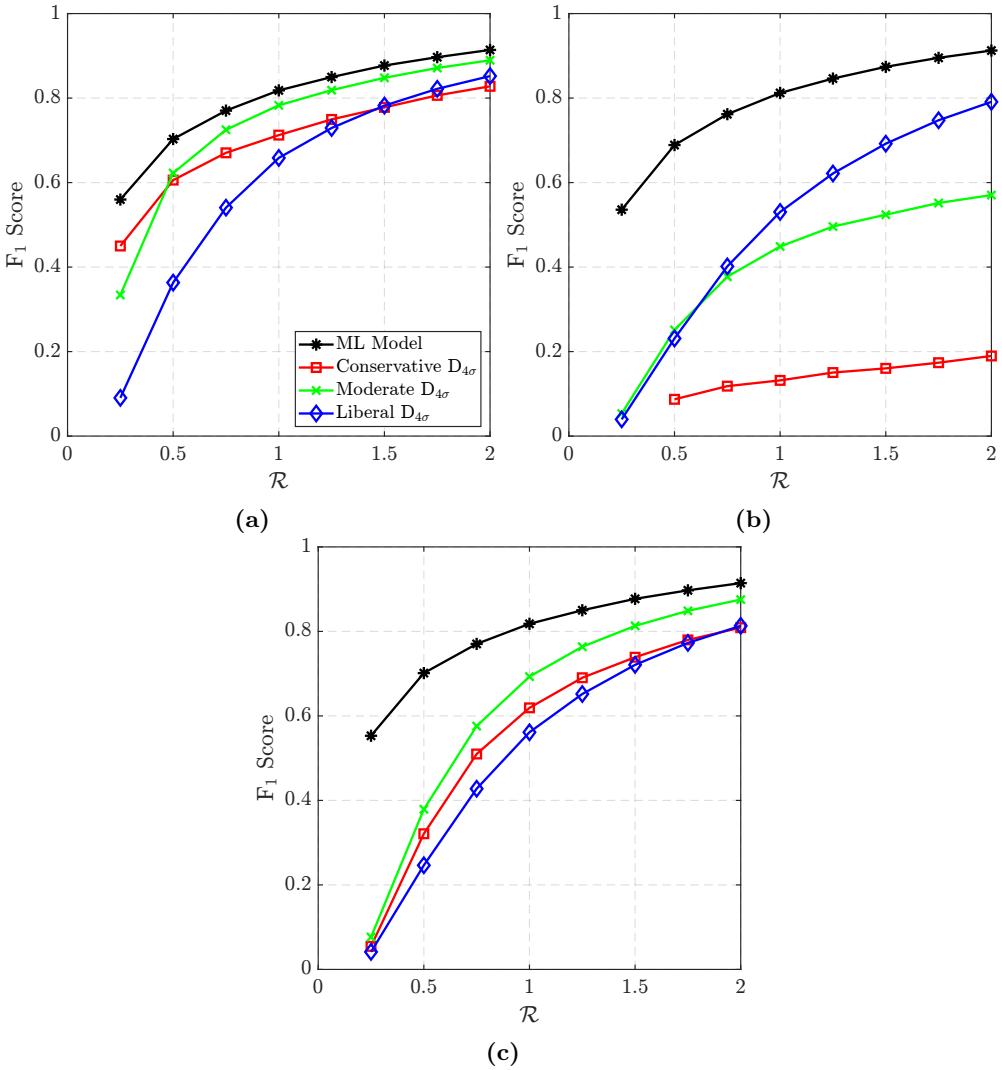

Figure 7: F_1 score vs. $\mathcal{R}$ for (a) no, (b) low, and (c) high noise levels.

Additionally, plots of ROC for all tested methods for each noise level are presented in Fig. 8. The ROC is succinctly a plot of true positive rate (TPR) vs. false positive rate (FPR). Presenting results in this manner highlights how well each method can distinguish between a branch-point affected AOI and not, despite $\mathcal{R}$ variation. Model performance is based on how closely the area under the curve (AUC) approaches 1, or how close the curve hugs the top left corner of the plot domain. The higher the AUC, the better the model is at classifying. For all three noise conditions, the ML model outperforms all the $D_{4\sigma}$ methods with AUC scores of 0.7003, 0.7004, and 0.6613 for the no, low, and high noise cases, respectively. For the no noise case, the $D_{4\sigma}$ AUC scores are 0.4903, 0.2298, 0.1068 for liberal, moderate, and conservative thresholds, respectively. Since the ROC for the $D_{4\sigma}$ methods in low and high noise are visibly erroneous, there is no need to report their AUC scores. Combining the results, it is easy to see the ML model's resilience and fidelity in identifying branch points from SHWFS data, proving to be both reliable in its predictions (high F_1), consistent across $\mathcal{R}$ (high ROC AUC),

and robust to additive noise when compared to the tested irradiance pattern spreading approach though $D_{4\sigma}$ thresholding.

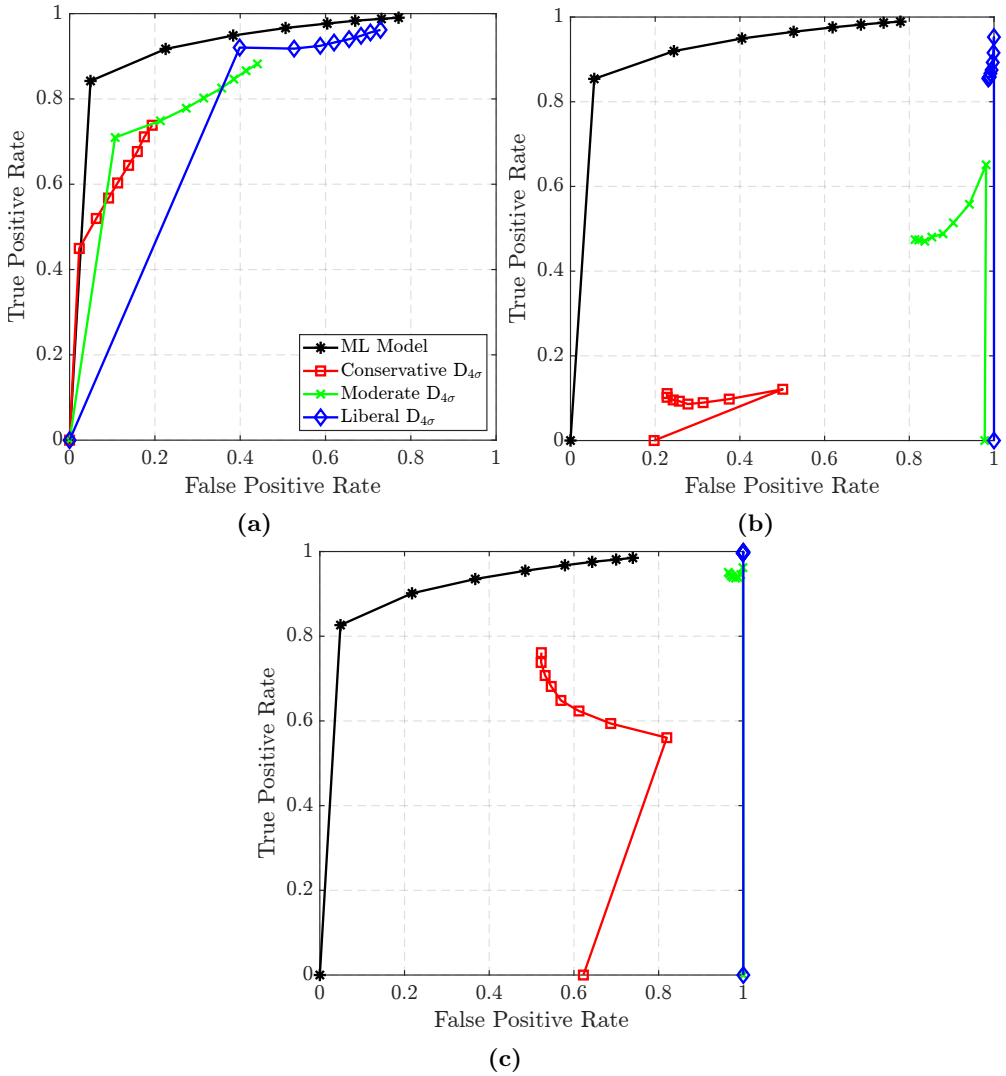

Figure 8: Receiver Operating Curve for (a) no, (b) low, and (c) high noise levels.

5. SUMMARY AND CONCLUSIONS

In this paper, we proposed the use of a CNN for branch-point detection using SHWFS data, as it is proven computationally inexpensive and robust to noise environments. The CNN model utilized in this study was trained to receive a single SHWFS sub-aperture irradiance pattern image and estimate whether or not a branch point was located at that pupil location. This model was tested alongside the $D_{4\sigma}$ method for identifying branch points, first presented in Ref. 30. To conduct this study, we leveraged wave-optics simulations to generate SHWFS imagery in varying optical turbulence conditions, which we characterized by the spherical-wave $\mathcal{R}$.

We computed performance metrics such as the true positive, false positive, and false negative rates for both methods. Taking these rate-based results at face value shows that the ML model performs marginally better than the $D_{4\sigma}$ method when used with a liberal thresholding value. However, these metrics do not fully capture the reliability of the method. Therefore, two additional metrics were explored; the F_1 score and the ROC AUC. Across all values of $\mathcal{R}$, the machine learning model outperforms the $D_{4\sigma}$ methods in F_1, proving more reliable. Additionally, the ROC AUC for the machine learning model was the highest at 0.7003 proving it to be a better

classifier when compared to the $D_{4\sigma}$ methods. We also found that adding Gaussian noise to the data does not affect the performance of the machine learning model, but greatly affects the performance of the $D_{4\sigma}$ methods. In general, this paper has shown the viability of machine learning for identifying branch points from SHWFS data. Currently, this analysis is performed solely on simulated data; however, future testing will seek to apply our methods to experimentally collected data. At large, this work helps to inform efforts looking to develop branch-point-tolerant adaptive-optics schemes which leverage the SHWFS.

Funding Sources

Grateful acknowledgment is given to the Joint Directed Energy Transition Office (JDETO) and Office of Naval Research (ONR). Their sponsorship of Notre Dame's Aero Optics Group and NSWCDD's Independent Applied Research (IAR) program (ONR award document number N0001425GI00781), respectively, has made this research possible.

Disclosures

The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright notation thereon.

Distribution Statement A. Approved for public release: distribution is unlimited.

REFERENCES

[1] Little, C., "A diffraction theory of the scintillation of stars on optical and radio wave-lengths," *Royal Astronomical Society* **111** (1950).

[2] Clifford, S. F., Ochs, G. R., and Lawrence, R. S., "Saturation of optical scintillation by strong turbulence," *JOSAA* **64** (1974).

[3] Yura, H. T., "Physical model for strong optical-amplitude fluctuations in a turbulent medium," *JOSAA* **64** (1974).

[4] Hill, R. J., "Theory of saturation of optical scintillation by strong turbulence: plane-wave variance and covariance and spherical-wave covariance," *JOSAA* **72** (1983).

[5] Andrews, L. C., Phillips, R. L., Hopen, C. Y., and Al-Habash, M. A., "Theory of optical scintillation," *JOSAA* **16** (1999).

[6] Fried, D. L. and Vaughn, J. L., "Branch cuts in the phase function," *Applied Optics* **31**, 2865 (may 1992).

[7] Barchers, J. D., Fried, D. L., and Link, D. J., "Evaluation of the performance of hartmann sensors in strong scintillation," *Applied Optics* **41**, 1012 (feb 2002).

[8] Roggemann, M. C. and Koivunen, A. C., "Wave-front sensing and deformable-mirror control in strong scintillation," *Journal of the Optical Society of America A* **17**, 911 (may 2000).

[9] Steinbock, M. J., Hyde, M. W., and Schmidt, J. D., "LSPV+7, a branch-point-tolerant reconstructor for strong turbulence adaptive optics," *Applied Optics* **53**, 3821 (jun 2014).

[10] Spencer, M. F. and Brennan, T. J., "Deep-turbulence phase compensation using tiled arrays," *Optics Express* **30**, 33739–33755 (aug 2022).

[11] Hyde, M. W., McCrae, J. E., Kalensky, M., and Spencer, M. F., "Hidden phase in two-wavelength adaptive optics," *Applied Optics* (2024).

[12] Kalensky, M., Getts, D., Banet, M. T., Burrell, D. J., Hyde, M. W., and Spencer, M. F., "Limitations of beam-control compensation," *Opt. Express* **32**, 42301–42317 (Nov 2024).

[13] Kalensky, M., Burrell, D. J., Spencer, M. F., Banet, M. T., Oesch, D. W., and Getts, D., "Branch-point compensation in extended-beacon adaptive optics," in [*Unconventional Imaging, Sensing, and Adaptive Optics 2024*], Bose-Pillai, S. R., Dolne, J. J., and Kalensky, M., eds., SPIE (Oct 2024).

[14] Kalensky, M., Burrell, D. K., Banet, M. T., and Spencer, M. F., "Hidden-phase compensation in extended-beacon adaptive optics," *Applied Optics* **64**(18) (2025).

[15] Kalensky, M., "Adaptive-optics limitations for laser beam propagation along slant paths," *Applied Optics* **64** (March 2025).

[16] Oesch, D. W., Sanchez, D. J., and Matson, C. L., "The aggregate behavior of branch points - measuring the number and velocity of atmospheric turbulence layers," *Optics Express* **18**, 22377 (oct 2010).

[17] Oesch, D. W., Sanchez, D. J., and Tewksbury-Christle, C. M., "Aggregate behavior of branch points - persistent pairs," *Optics Express* **20**, 1046 (jan 2012).

[18] Oesch, D. W., Tewksbury-Christle, C. M., Sanchez, D. J., and Kelly, P. R., "Aggregate behavior of branch points: characterization in wave optical simulation," *Optical Engineering* **51**, 106001 (oct 2012).

[19] Oesch, D. W. and Sanchez, D. J., "Photonic orbital angular momentum in starlight," *Astronomy Astrophysics* **567**, A114 (jul 2014).

[20] Beck, J. R., Bos, J. P., Brennan, T. J., and Spencer, M. F., "Wave-optics investigation of branch-point density," *Optical Engineering* **61** (apr 2022).

[21] Voitsekhovich, V. V., Kouznetsov, D., and Morozov, D. K., "Density of turbulence-induced phase dislocations," *Applied Optics* **37**, 4525 (jul 1998).

[22] Holmes, R. B., "Scintillation-induced jitter of projected light with centroid trackers," *Journal of the Optical Society of America A* **26**, 313 (jan 2009).

[23] Mitchell, E. W., Burrell, D. J., Hyde, M. W., Driggers, R. G., and Spencer, M. F., "Scaling laws for the noise-equivalent angle and c-tilt, g-tilt anisoplanatism due to scintillation," *Applied Optics* **64**(18) (2025).

[24] Hyde, M. W., Mitchell, E. W., and Spencer, M. F., "Closed-form expressions for the centroid-tilt error due to scintillation," *Optics and Lasers in Engineering* **194**(109166) (2025).

[25] Fried, D. L., "Branch point problem in adaptive optics," *Journal of the Optical Society of America* **15**, 2759 (oct 1998).

[26] Chen, M., Roux, F. S., and Olivier, J. C., "Detection of phase singularities with a shack-hartmann wavefront sensor," *Journal of the Optical Society of America A* **24**, 1994 (jun 2007).

[27] Wild, W. J. and Bigot, E. O. L., "Rapid and robust detection of branch points from wave-front gradients," *Optics Letters* **24**, 190 (feb 1999).

[28] Murphy, K. and Dainty, C., "Comparison of optical vortex detection methods for use with a shack-hartmann wavefront sensor," *Optics Express* **20**, 4988 (feb 2012).

[29] Tyler, G. A., "Reconstruction and assessment of the least-squares and slope discrepancy components of the phase," *Journal of the Optical Society of America* **17**, 1828 (oct 2000).

[30] Kalensky, M., "Branch-point identification using second-moment shack hartmann wavefront sensor statistics," *Applied Optics* **62**, G101 (may 2023).

[31] Kalensky, M., Oesch, D., Bukowski, T. J., Miller, K., and Getts, D., "Branch-point detection using a shack-hartmann wavefront sensor," in [*Unconventional Imaging, Sensing, and Adaptive Optics 2023*], Bose-Pillai, S. R., Dolne, J. J., and Spencer, M. F., eds., SPIE (oct 2023).

[32] Kalensky, M., Oesch, D. W., Bukowski, T. J., Miller, K., and Getts, D., "Comparison of branch-point detection approaches using a shack hartmann wavefront sensor," *Optical Engineering* **62** (Dec. 2023).

[33] Oesch, D. W., Kalensky, M., and Sexauer, M. S., "Branch point detection in smash measurements using second-moment wavefront sensor statistics," in [*Unconventional Imaging, Sensing, and Adaptive Optics 2024*], Bose-Pillai, S. R., Dolne, J. J., and Kalensky, M., eds., SPIE (Oct 2024).

[34] 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).

[35] 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).

[36] 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).

[37] Noel, Z. A., Bukowski, T. J., Gordeyev, S., and Rennie, R. M., "Shack-Hartmann wavefront reconstruction by deep learning neural network for adaptive optics," in [*Unconventional Imaging, Sensing, and Adaptive Optics 2023*], Dolne, J. J., Spencer, M. F., and Bose-Pillai, S. R., eds., **12693**, 126930G, International Society for Optics and Photonics, SPIE (2023).

[38] Sahba, S., Wilcox, C. C., McDaniel, A., Shaffer, B. D., Brunton, S. L., and Kutz, J. N., "Wavefront sensor fusion via shallow decoder neural networks for aero-optical predictive control," in [*Interferometry XXI*], North-Morris, M. B., Creath, K., and Porras-Aguilar, R., eds., **12223**, 1222303, International Society for Optics and Photonics, SPIE (2022).

[39] Sahba, S., *Machine Learning for Aero-Optical Wavefront Characterization and Forecasting*, PhD thesis, University of Washington (2023).

[40] He, Y., Liu, Z., Ning, Y., Li, J., Xu, X., and Jiang, Z., "Deep learning wavefront sensing method for Shack-Hartmann sensors with sparse sub-apertures," *Optics Express* **29**, 17669–17682 (May 2021). Publisher: Optica Publishing Group.

[41] Hu, L., Hu, S., Gong, W., and Si, K., "Deep learning assisted Shack–Hartmann wavefront sensor for direct wavefront detection," *Optics Letters* **45**, 3741–3744 (July 2020). Publisher: Optica Publishing Group.

[42] DeFoor, T. E., Bukowski, T. J., Kalensky, M., and Gordeyev, S., "Phase discontinuity classification from single-aperture irradiance patterns using machine learning," in [*Unconventional Imaging, Sensing, and Adaptive Optics 2024*], **13149**, 156–167, SPIE (Oct. 2024).

[43] Tatarskii, V. I., [*Wave propagation in a turbulent medium*], Dover Publications, Inc, Mineola, New York (2016).

[44] Martin, J. M. and Flatté, S. M., "Intensity images and statistics from numerical simulation of wave propagation in 3-d random media," *Applied Optics* **27**, 2111 (jun 1988).

[45] Martin, J. M. and Flatté, S. M., "Simulation of point-source scintillation through three-dimensional random media," *Journal of the Optical Society of America* **7**, 838 (may 1990).

[46] Schmidt, J., [*Numerical simulation of optical wave propagation with examples in MATLAB*], SPIE, Bellingham, Wash (2010).

[47] Voelz, D. G., [*Computational Fourier Optics: A MATLAB Tutorial*], SPIE Press, Bellingham, Washington (2011).

[48] Géron, A., [*Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems*], O'Reilly Media, Dar es Salaam (2024).

[49] Krizhevsky, A., Sutskever, I., and Hinton, G. E., "ImageNet Classification with Deep Convolutional Neural Networks," in [*Advances in Neural Information Processing Systems*], **25**, Curran Associates, Inc. (2012).

[50] Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A., "Going deeper with convolutions," in [*2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*], 1–9 (June 2015). ISSN: 1063-6919.