

SpeckleNet: Perceiving Closely-Spaced Objects with Speckle Interferometry

Andrew Vanden Berg
Air Force Research Lab

Ian Cunnyingham
Morph Optic, Inc.

Justin Fletcher
United States Space Force Space Systems Command

ABSTRACT

Speckle interferometry is a well-established astronomical technique that enables detection of closely-spaced objects (CSOs) near the theoretical limit of a telescope's aperture. The application of this technique can mitigate the resolution-limiting effects of the atmosphere. Observations comprise many short exposures which freeze the turbulence of the atmosphere. These frames are then stacked in the spatial-Fourier domain in which the phase errors induced by the atmosphere average out, given many samples. If a CSO is present, a characteristic fingerprint can be observed in the amplitude of the stacked images. The structure of this fingerprint is determined by parameters of the underlying scene, including apparent radial separation and brightness of the objects. In certain limits of these parameters, this fingerprint degrades. There remains an unsolved decision problem, which is to determine the plurality of objects, given a speckle interferometric image. In this work, we formalize this decision problem as a binary classification task: SpeckleNet. We apply modern computer vision approaches, including convolutional neural networks, to classify images by plurality and characterize the performance of each classifier over a range of underlying parameters. The resulting models require only a few milliseconds to infer plurality and can be adapted to any optical telescope, requiring only collection software and post-processing changes.

1. INTRODUCTION

The space domain grows increasingly dynamic. The advent of on-orbit servicing and megaconstellations have simultaneously expanded the population of satellites while also placing them in closer proximity to one another. As such, close encounters between artificial Earth satellites are likely to proliferate. Affordable sensing concepts must be fielded at scale to enable the continuity of space domain awareness (SDA) through these encounters. In this work, we automate the perception of images formed using an established astronomical technique to validate a low-cost optical sensing concept that addresses this need. The observational technique, known as speckle interferometry, requires only traditional ground-based optical instruments and could be readily adopted by many existing sensor systems.

We contribute: 1) a task description that formalizes the perceptual task of speckle interferometry, 2) four synthetic datasets representing a baseline task realization against which future exploitation techniques may be measured, and 3) a novel learned approach to speckle interferometry exploitation complete with comparative results from a variety of model architectures.

Related works are discussed in Section 2, and the task is detailed in Section 3. Section 4 describes the data used in this work and documents the synthetic data generation algorithm, along with its parameterizations and assumptions. The datasets generated to support this work are then described. Section 5 provides a summary of experiments conducted, including a description of the model architectures used, and reports results. Experimental results are quantified using a variety of metrics, and variations in performance are analyzed. We conclude the section with the information necessary to reproduce this work. Finally, we summarize our findings in Section 6 and suggest several topics for future work.

DISTRIBUTION A. Approved for public release: distribution is unlimited. Public Affairs release approval AFRL-2022-3898.

The views expressed are those of the author and do not necessarily reflect the official policy or position of the Department of the Air Force, the Department of Defense, or the U.S. government.

2. RELATED WORKS

Short exposure astronomical imaging was pioneered by Fried in 1966. By constraining the integration time of the imagery to be shorter than the coherence time of the atmosphere, this approach enables the production of images in which atmospheric contributions to wavefront aberrations are static [6]. Taking several such images in quick succession allows an observer to build up a sample of the variation in atmospheric aberration, while the target remains unchanged. These images may then be stacked [1] or down-selected [7] to achieve diffraction-limited image recovery. Speckle interferometry was introduced by Labeyrie, and is built on the intuition that contributions from speckle patterns average out as large numbers of frames are added in the Fourier domain; for objects that have a center of symmetry, as is the case for binary stars and closely-spaced satellite point-sources, what remains is a characteristic fringe pattern [16]. This approach was later extended by Knox and Thompson [13] to enable recovery of phase information from spatially-extended objects. The Knox-Thompson technique has been applied to satellite image recovery [26]. Later work by Lawrence et. al. applied Labeyrie’s technique to the problem of extended object image recovery, and provided the first formulation of the closely-space object perception task addressed in this work [17]. Recent work by Scott et. al. has extended the application of speckle interferometry to on-orbit servicing of satellites in geosynchronous Earth orbit (GEO)[22, 25, 24, 23]. Our work builds on the lineage of these recent works by adding a data-driven baseline task formulation, while also providing a novel, learned solution to that task.

The advent of learned SDA image processing is comparatively recent. Since the introduction of learned deep space object detection [5], several works have expanded upon the initial formulation [29, 4, 3]. Successful learned approaches to a variety of related problems, including satellite identification [8, 27, 9], pose estimation [21, 18], and resolved imagery enhancement and perception [14, 15, 19, 28], have been published. We adopt the closely-spaced object detection measurement approach introduced by Gazak et. al. [10], and expand upon it by introducing an approach that performs the same task without a dependence upon an optical dispersion element. This work may be positioned as an extension of the growing interdisciplinary field of learned SDA data processing.

3. PROBLEM

We approach the problem of identifying the presence of closely-space objects as a binary classification task. Let X be a dataset of n images with an input image denoted as x_i , where $i \in \{1, 2, \dots, n\}$. Let Y be the truth classifications of the dataset with the classification of an image $y_i \in \{0, 1\}$. Here $y_i = 0$ corresponds to the negative class where only a single object is present in the image, and $y_i = 1$ corresponds to the positive class in which the image contains a CSO. We seek to train a convolutional neural network that represents the functional mapping $f : X \mapsto Y$ in which an input image x_i is mapped to its classification inference y_i .

For the purposes of this work, we only consider the presence of a single CSO in any given image. This classification task could also be extended to a regression problem where the task is to identify the orientation, apparent separation, and delta visual magnitude of the CSO present, and this could be explored in future work.

4. DATASETS

We use simulated data for model training and evaluation and contribute an algorithm for simulating an image composed of speckle images stacked across the spatial-Fourier domain. This algorithm is described below.

Simulated data in this work was modeled after the 3.6m Advanced Electro-Optical System (AEOS) telescope located at the Air Force Maui Optical and Supercomputing (AMOS) Observatory, and we modeled a sensor capable of collecting short-exposure frames in the visible spectrum. This simulation algorithm can be adapted to any optical system, however.

4.1 Simulation

Dataset simulation begins with a blank image B represented by a two-dimensional array of zeros with width and height equal to the desired image size z in pixels. An object is added to the center pixel with a pixel value of 1.0. It is assumed that the objects of observation extends over less than one pixel in the image raster. We call this the *bright object* because any CSO added to the frame will be dimmer in visual magnitude than the bright object. If a CSO is present, a *dim object* (i.e., a CSO) is added to the image with an (x, y) offset from the bright object. The dim object

DISTRIBUTION A. Approved for public release: distribution is unlimited. Public Affairs release approval AFRL-2022-3898.

The views expressed are those of the author and do not necessarily reflect the official policy or position of the Department of the Air Force, the Department of Defense, or the U.S. government.

offset is determined by the desired orientation (radians) and apparent separation at GEO (m) of the dim object, and it is calculated using the pixel-wise instantaneous field of view (IFOV, in arcsecs) of the sensor along with the distance to GEO orbit (m) from the Earth's surface. With a given orientation θ , separation at GEO s , IFOV v , and distance to orbit b , we have the offset:

$$(x, y) = \left(\frac{s}{\frac{vb\pi}{180 \cdot 3600}} \cos(\theta), \frac{s}{\frac{vb\pi}{180 \cdot 3600}} \sin(\theta) \right)$$

The denominator $(vb\pi)/(180 \cdot 3600)$ is a conversion to meters per pixel from the arcseconds per pixel given by the sensor IFOV. Once the pixel offset is determined, the pixel value of the dim object can be calculated:

$$f(\Delta m) = 10^{-0.4|\Delta m|}$$

Here, Δm is the difference in visual magnitude between the bright object at the image center to the dim CSO added. Recall that the bright object begins with a pixel value of 1.0.

We then model the effects of diffraction by convolving an Airy disk with the image based on the telescope aperture diameter, the sensor IFOV, and the filter center wavelength c and bandwidth w . Another blank image of zeros is created with the same desired image size, and the following function is applied to each pixel in the image:

$$g(r) = A \left[\frac{2J_1\left(\frac{\pi r}{R/R_z}\right)}{\frac{\pi r}{R/R_z}} \right]^2$$

In this function, r is the radial distance of the pixel from the maximum of the Airy function, taken to be located at the center of the image. R is the radius of the Airy disk, approximated as $1.22\frac{\lambda}{d}$ for small angles, with wavelength λ and aperture diameter d . $R_z \approx 1.22$, and J_1 is the first order Bessel function of the first kind. This Airy disk is normalized and then convolved with the image B containing the objects.

We then convert the image to its power spectrum PS , defined as follows:

$$PS(\cdot) = |H(\cdot)|^2$$

The power spectrum PS is equal to the squared magnitude of the Fourier transform H of the image. We apply a two-dimensional fast Fourier transform to this image to compute the Fourier transform, and the fringe patterns emerge in the power spectrum when two objects are present in the initial image.

To the power spectrum, we apply atmospheric turbulence and noise. Atmospheric calibration frames were collected on stars of known brightness using the 3.6m AEOS telescope with integration times of 100 msec. These calibration frames were averaged across an assumed atmospheric decorrelation time of 1 sec. The mean and standard deviation of these averaged calibration frames were determined to capture stochasticity due to the atmosphere. We used this same integration time of 100 msec for all simulated data in this work; changing the integration time would require collection of additional calibration data.

Next, we compute the total photon counts present in the image for each object. The spectral flux density ($W/m^2/m$) is calculated from the center of the sensor filter bandwidth and converted to photon flux density (photons/sec/ m^2/m). The total photon count N for one object is then the product of the integration time t (sec), sensor filter bandwidth w (m), aperture area A (m^2), quantum efficiency¹ q , photon flux density p , and the brightness ratio of the object, equivalent to $10^{-0.4m}$ with object visual magnitude m :

$$N = 10^{-0.4m} A w p q t$$

If two objects are present, the photon counts of each object can be added together for the total photons in the image.

¹This work assumes a quantum efficiency of 0.8.

DISTRIBUTION A. Approved for public release: distribution is unlimited. Public Affairs release approval AFRL-2022-3898.

The views expressed are those of the author and do not necessarily reflect the official policy or position of the Department of the Air Force, the Department of Defense, or the U.S. government.

To add the atmospheric turbulence to the power spectrum, we generate an atmospheric turbulence frame of our desired image size, where the value of each pixel is sampled from a normal distribution with mean and standard deviation equal to those calculated from the atmospheric calibration data. We take the element-wise product of the power spectrum PS and the atmospheric turbulence AT and normalize with the total photon count:

$$PS \odot AT \frac{N^2}{\max(PS \odot AT)}$$

Next, we add read noise to the power spectrum. With a sensor read noise variance e , the standard deviation of the read noise for the power spectrum PS_{std} is expressed as follows:

$$PS_{std} = \sqrt{\frac{e N z^2}{\text{(number of frames)}}}$$

Finally, we conduct a fast Fourier transform shift to move the zero-frequency component to the center of the power spectrum. This simulation process is repeated for each sample in the dataset, and the entire algorithm is outlined in Alg. 1.

Algorithm 1 Image simulation

```

1:  $B \leftarrow \text{Zeros}(z, z)$ 
2:  $B(\frac{z}{2}, \frac{z}{2}) \leftarrow 1.0$ 
3: if  $CSO$  present then
4:    $(x, y) \leftarrow \text{Offset}(s, v, b, \theta)$ 
5:    $B(\frac{z}{2} + x, \frac{z}{2} + y) \leftarrow 10^{-0.4|\Delta m|}$ 
6: end if
7:  $K \leftarrow \text{AiryDisk}(z, d, v, c, w)$ 
8:  $B^* \leftarrow \text{Convolve}(B, K)$ 
9:  $PS \leftarrow |\text{2DFastFourierTransform}(B^*)|^2$ 
10:  $AT \leftarrow AT_{mean} + AT_{std} \text{StandardNormal}(z, z)$ 
11:  $N_1 \leftarrow \text{PhotonCount}(m_1, A, w, p, q, t)$ 
12: if  $CSO$  present then
13:    $N_2 \leftarrow \text{PhotonCount}(m_2, A, w, p, q, t)$ 
14:    $N \leftarrow N_1 + N_2$ 
15: else
16:    $N \leftarrow N_1$ 
17: end if
18:  $PS^* \leftarrow PS \odot AT \frac{N^2}{\max(PS \odot AT)}$ 
19:  $PS_{std} \leftarrow \sqrt{\frac{e N z^2}{\text{(num. frames)}}}$ 
20:  $PS^* \leftarrow PS + PS_{std} \text{StandardNormal}(z, z)$ 
21:  $PS^* \leftarrow \text{FastFourierTransformShift}(PS^*)$ 

```

Example images from the simulation algorithm are shown in Fig. 1. The image on the left shows two objects with visual magnitudes of 12.0 separated by 50m; the image on the right shows one object with a visual magnitude of 12.0. The fringe patterns, while faint, are present in the image with two objects. We elected to forgo any image preprocessing at the end of the simulation algorithm that would highlight or emphasize any fringe patterns to the human eye and instead allow the models to learn patterns from the raw, simulated images. This decision was supported by empirical results from initial experiments.

4.2 Parameters

We simulated four datasets for this work using identical input parameters, but varying the sample size. These dataset parameter distributions are outlined in Table 1. An image size of 400 pixels was used for all datasets. All possible orientations of the CSO relative to the center object were considered. The minimum radial offset of 0m was chosen to

DISTRIBUTION A. Approved for public release: distribution is unlimited. Public Affairs release approval AFRL-2022-3898.

The views expressed are those of the author and do not necessarily reflect the official policy or position of the Department of the Air Force, the Department of Defense, or the U.S. government.

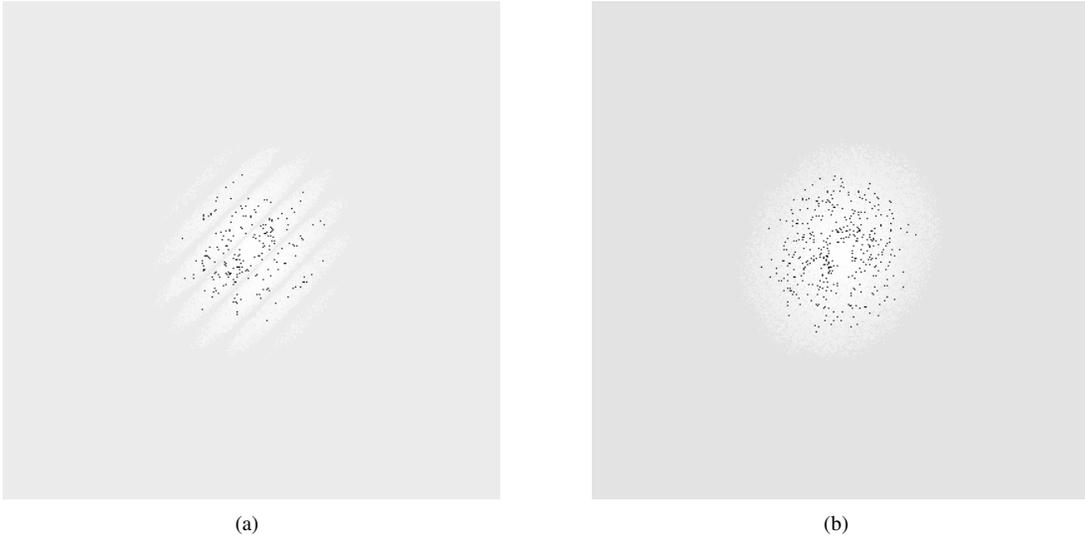


Fig. 1: Simulated images with (a) CSO present and (b) only a single object

test the limits of model performance - distinguishing between two objects with no separation would not be possible, but we wanted to evaluate where performance degradation occurred. The radial offset maximum of 200m was selected to understand performance as the objects become resolved visually without speckle interferometry techniques. We considered dim objects with visual magnitudes between 10 and 18 and varied the delta magnitude (and therefore the bright object magnitude) from 0 to 7. Satellites in GEO orbits often have visual magnitudes between 14-18, but can be brighter than 10 visual magnitude depending on the satellite orientation or when glinting occurs. The fringe patterns that exist in the power spectrum when two objects are present are most apparent when the objects have the same brightness. We varied the delta visual magnitude up to 7 to again characterize where model performance degrades.

Table 1: Dataset parameters

Parameter	Distribution
CSO Present	Uniform{True, False}
Radial Offset (m)	Uniform(0, 200)
Dim Object Magnitude	Uniform(10, 18)
Delta Visual Magnitude	Uniform(0, 7)
Orientation (rad)	Uniform($-\pi$, π)

For each sample in a dataset, the distribution parameters were sampled, and then the simulation algorithm was applied. We generated four datasets with these parameters, with sample sizes of 20K, 100K, 400K, and 1M (1000K). The 20K-dataset was split into train, test, and validation with splits of 0.5, 0.25, and 0.25, respectively. The other three datasets used splits of 0.8, 0.1, and 0.1, respectively. All samples were stored as PNGs with a bit depth of 16.

5. EXPERIMENTS

This section outlines the experiments conducted in this work. Accuracy is the primary metric we considered given the task was binary classification. We also report precision, recall, and F_1 results. We compute the true positives N_{TP} , false positives N_{FP} , and false negatives N_{FN} . A positive prediction corresponds to predicting that a CSO is present, and a negative prediction corresponds to predicting that only a single object is present. Precision, or specificity, is the fraction that the model is correct when a CSO is predicted, and it is defined as follows:

$$\text{precision} = \frac{N_{TP}}{N_{TP} + N_{FP}}$$

DISTRIBUTION A. Approved for public release: distribution is unlimited. Public Affairs release approval AFRL-2022-3898.

The views expressed are those of the author and do not necessarily reflect the official policy or position of the Department of the Air Force, the Department of Defense, or the U.S. government.

Recall, or sensitivity, is the fraction of CSOs that are captured by positive model predictions, and it is defined as follows:

$$\text{recall} = \frac{N_{TP}}{N_{TP} + N_{FN}}$$

We also report F_1 score, which is a metric that combines precision and recall, and it is defined as follows:

$$F_1 = 2 \cdot \frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}}$$

We also evaluate the model performance across the dataset parameters, namely those of radial offset and delta visual magnitude.

5.1 Models

In this work, we employed convolutional neural networks (CNNs) to the binary classification task at hand, and we utilized ResNets [11] as our model backbone, a standard for classification tasks across the field of computer vision. We compared the performance of ResNet-34, -50, -101, and -152 to evaluate the impact of model capacity on performance.

5.2 Results

Accuracy measures on the respective test sets are reported in Table 2. The highest accuracy of 80.3% was found with the ResNet-34 model trained on 1M samples. Performance across different ResNet models is relatively consistent. The increased model capacity of ResNet-101 and -152 does not translate to higher performance. This suggests that either the underlying task is not too complex and therefore does not require increased model capacity, or additional hyperparameter tuning is necessary for ResNet-101 and ResNet-152 in order to boost performance. Models with higher capacity tend to require additional tuning, especially related to regularization. However, the fact that the lower capacity ResNet-34 and ResNet-50 perform well means that these models could be more easily deployed to edge computing nodes where computational power may be limited.

As the dataset size increases, performance increases but not continually. Lower performance on the 20K-sample dataset suggests that not enough samples were presented to the model to effectively learn the task; however, increasing the dataset size by an order of magnitude from 100K to 1M samples does not drastically increase performance.

Table 2: Performance metrics

Architecture	Dataset Size	Accuracy (%)	Precision	Recall	F_1
ResNet-34	20K	73.6	0.855	0.572	0.685
	100K	77.4	0.906	0.616	0.734
	400K	80.0	0.944	0.631	0.757
	1M	80.3	0.952	0.639	0.765
ResNet-50	20K	74.1	0.842	0.595	0.697
	100K	78.2	0.910	0.633	0.747
	400K	80.0	0.944	0.638	0.761
	1M	80.2	0.952	0.637	0.763
ResNet-101	20K	74.1	0.828	0.610	0.702
	100K	78.1	0.905	0.634	0.746
	400K	79.3	0.935	0.631	0.753
	1M	79.6	0.961	0.617	0.751
ResNet-152	20K	74.2	0.894	0.552	0.683
	100K	77.9	0.940	0.602	0.734
	400K	79.6	0.958	0.620	0.753
	1M	80.0	0.960	0.626	0.757

ResNet-34 tended to perform the best across the various ResNet architectures tested, achieving the highest accuracy, recall, and F_1 score. The highest precision of 0.961 was achieved by ResNet-101 trained on the 1M-sample dataset. A

DISTRIBUTION A. Approved for public release: distribution is unlimited. Public Affairs release approval AFRL-2022-3898.

The views expressed are those of the author and do not necessarily reflect the official policy or position of the Department of the Air Force, the Department of Defense, or the U.S. government.

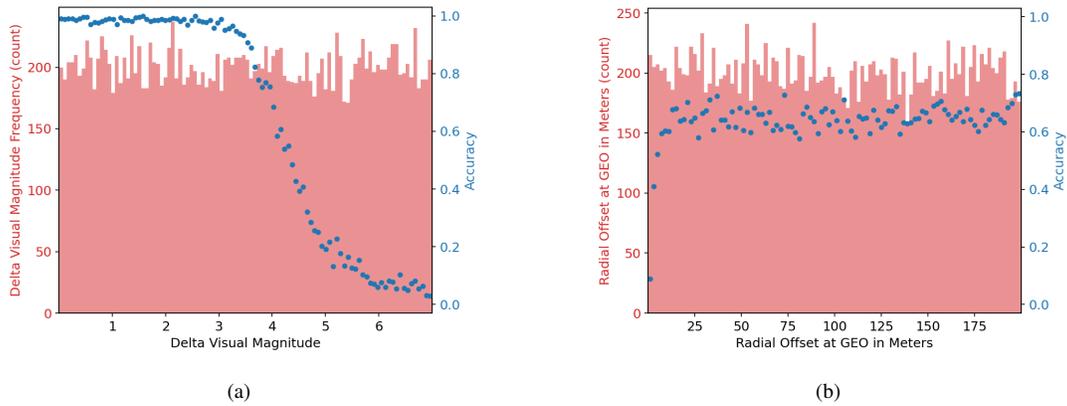


Fig. 2: Performance across (a) delta visual magnitude and (b) radial offset

high precision is desired in this problem because it means that there are very few false positives. Because a positive prediction indicates the presence of a previously unknown CSO, positive predictions (true or false) are likely to trigger follow up by the user. Fewer false positives would therefore reduce the burden of unwarranted follow up.

The highest recall of 0.639 was achieved by ResNet-34 trained on the 1M-sample dataset. This result means the model is making a substantial number of false negative predictions. While this is undesirable, it is not surprising given the nature of the problem and the parameters of the datasets. As the delta visual magnitude increases, particularly above 4.0, the fringe patterns will be increasingly faint, so we would not expect the model to detect nonexistent fringes. However, this work did not conduct an exhaustive hyperparameter search, so finetuning the model hyperparameters could improve recall.

The performance of the ResNet-50 model trained with the 400K-sample with various delta visual magnitudes and radial offsets is reported in Fig. 2. The plot on the left shows the model performs extremely well with low delta visual magnitude, and performance begins to deteriorate when the delta visual magnitude rises above 3. As the delta visual magnitude increases, the model accuracy drops to zero. These images will have faint or nonexistent fringe patterns, so this is evidence the model is learning to classify images without fringe patterns as single objects. The plot on the right in Fig. 2 also reveals that the model achieves a consistent performance when the apparent separation at GEO between the two objects is high, and performance degrades below 10m of separation.

Examining the confusion matrix of the ResNet-50 model trained on the 400K-sample dataset, shown in Fig. 3, we see the model predicts a single object roughly twice as often as a CSO. With a dataset balanced across the single object and CSO classes, this leads to lower recall, as previously discussed, but it means the model is accurately identifying the absence of fringe patterns present with a single object or when the two objects have a high delta visual magnitude.

Fig. 4 shows the performance of the ResNet-50 model trained the four datasets with varying delta visual magnitude and apparent separation. We utilized a custom x -axis to detail model performance for low separations. As the dataset size increases, the model performance becomes more consistent. The red region in the lower right-hand corner of each plot reveals that the model achieves extremely high accuracy for low delta magnitudes when the apparent separation greater than 4m, including in the 4-8m range below the diffraction-limiting capability of the sensor. Below 4m of separation, the model occasionally still detects a CSO, but performance is closer to random. Model performance degrades above 4.2 delta visual magnitude, but accuracy scores remain high from 2.8-4.2 delta visual magnitude.

Thus far, we have reported performance with a default threshold of 0.50 for positive predictions. Fig. 5 shows the precision-recall curves of each model architecture trained on the 400K-sample dataset. All model architectures perform similarly, and each max F_1 , denoted F_1^* , is approximately equal to 0.78, with precision values ≈ 0.8 and recall values ≈ 0.76 . Relative to the performances reported in Table 2, there is an improvement to recall which comes with a drop to precision. However, performance does not change drastically.

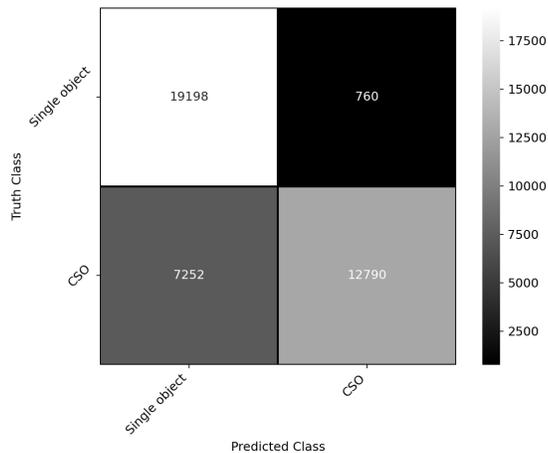


Fig. 3: Confusion matrix

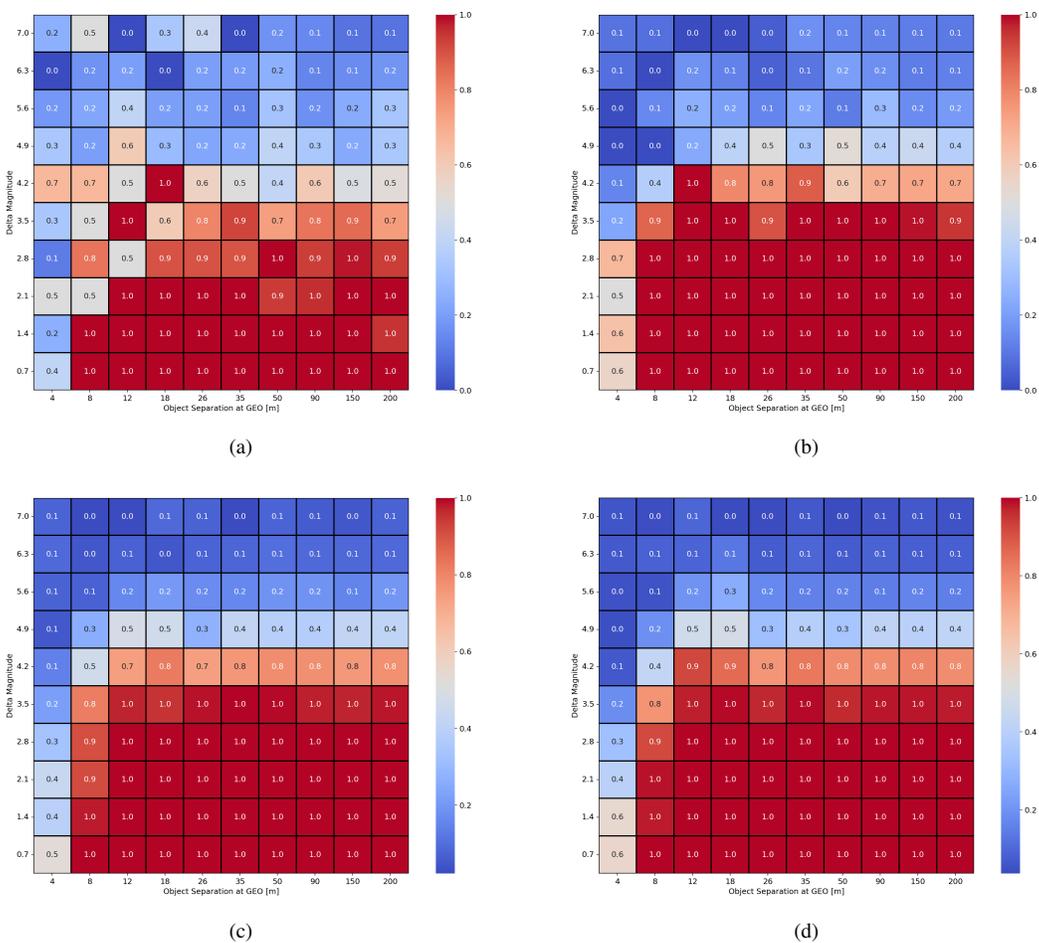


Fig. 4: Delta visual magnitude vs separation with dataset sizes of (a) 20K, (b) 100K, (c) 400K, and (d) 1M

DISTRIBUTION A. Approved for public release: distribution is unlimited. Public Affairs release approval AFRL-2022-3898.

The views expressed are those of the author and do not necessarily reflect the official policy or position of the Department of the Air Force, the Department of Defense, or the U.S. government.

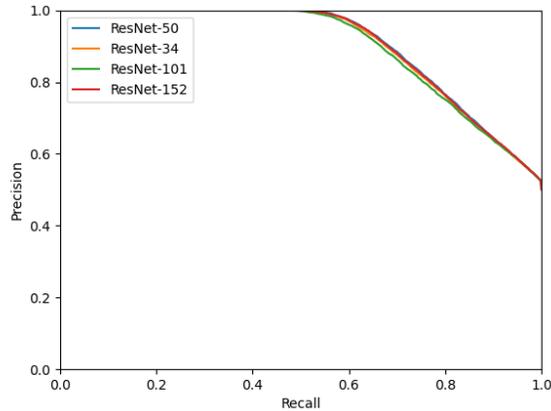


Fig. 5: Precision-recall curves

5.3 Reproducibility

We implemented models and conducted experiments with the MMClassification Python package from OpenMMLab [2], which is built on top of PyTorch [20]. In particular, Python version 3.7, PyTorch 1.6, and MMClassification 0.23 were utilized. Models were trained on NVIDIA Tesla V100 GPUs, each with 32GB of memory. We used a batch size of 256 for ResNet-34 and -50 models and a batch size of 128 for ResNet-101 and -152. We trained models with the Adam optimizer [12] for 200 epochs on the 20K-sample and 100K-sample datasets and 100 epochs for the 400K and 1M datasets. The goal of this work was to explore the feasibility of leveraging CNNs for this classification task, so we did not conduct an exhaustive hyperparameter search.

6. CONCLUSION

In this work, we presented a simulation algorithm for generating synthetic speckle interferometric images. We then demonstrated the capability of convolutional neural networks to excel in the binary classification task of identifying whether a CSO is present in a given image. The best model achieved 80.3% accuracy in a challenging dataset with a precision of 0.952. The models consistently detect CSOs down to 4m of apparent separation at GEO and delta visual magnitudes of up to 4.2. Below separations of 4m and with delta magnitudes above 4.2, model performance degrades.

Higher capacity models did not achieve higher performance, but additional hyperparameter tuning could improve performance of these models. Increasing the dataset size improved performance, but performance leveled off with a dataset size above 400K samples.

This work leveraged simulated data modeled after a single sensor in the 3.6m AEOS optical path. Other sensor parameters should be studied to identify when CSO detection is possible; a detailed sensor parameter study could inform the design and deployment of new sensors in ground-based sensor networks for the purpose of CSO detection. Future work should also investigate model performance on real data. Lastly, this classification problem could be extended to a regression task wherein the model would predict CSO separation, visual magnitude, and orientation. Incorporating more than one CSO is another possible extension, and the existing model architectures could be modified for more general classification tasks.

Acknowledgements

The authors wish to thank Brandoch Calef for his assistance in the development of the simulation algorithm code.

7. REFERENCES

- [1] F. M. Cady and R. H. T. Bates. Speckle processing gives diffraction-limited true images from severely aberrated instruments. *Optics Letters*, 5(10):438–440, October 1980.
- [2] MMClassification Contributors. Openmmlab’s image classification toolbox and benchmark. <https://github.com/open-mmlab/mmlclassification>, 2020.
- [3] Garrett Fitzgerald, Vijayan Asari, and Ruixu Liu. Geosynchronous satellite detection and tracking with WFOV camera arrays using spatiotemporal neural networks (GEO-SPANN). In *Pattern Recognition and Tracking XXXIII*, volume 12101, pages 25–40. SPIE, 2022.
- [4] Garrett Fitzgerald, Zachary Funke, Alexander Cabello, Vijayan Asari, and Justin Fletcher. Toward deep-space object detection in persistent wide field of view camera arrays. In *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference*, 2021.
- [5] Justin Fletcher, Ian McQuaid, Peter Thomas, Jeremiah Sanders, and Greg Martin. Feature-Based Satellite Detection using Convolutional Neural Networks. In *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference*, Maui, HI, September 2019.
- [6] D. L. Fried. Optical Resolution Through a Randomly Inhomogeneous Medium for Very Long and Very Short Exposures. *JOSA*, 56(10):1372–1379, October 1966.
- [7] David L. Fried. Probability of getting a lucky short-exposure image through turbulence. *JOSA*, 68(12):1651–1658, 1978.
- [8] J. Zachary Gazak, Ian McQuaid, Ryan Swindle, Matthew Phelps, and Justin Fletcher. SpectraNet: Learned Recognition of Artificial Satellites from High Contrast Spectroscopic Imagery. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 4012–4020, 2022.
- [9] J. Zachary Gazak, Ian McQuaid, Brandon Wolfson, and Justin Fletcher. Incremental learning of novel resident space object spectral fingerprints. In *Proceedings of AMOS Tech Conference*, 2021.
- [10] J. Zachary Gazak, Ryan Swindle, Zachary Funke, Matthew Phelps, and Justin Fletcher. Closely spaced object detection utilizing spatial information in spectroastrometric observations. In *Sensors and Systems for Space Applications XV*, volume 12121, pages 180–193. SPIE, 2022.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- [12] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR (Poster)*, 2015.
- [13] Keith T. Knox and Brian J. Thompson. Recovery of images from atmospherically degraded short-exposure photographs. *The astrophysical journal*, 193:L45–L48, 1974.
- [14] Trent Kyono, Jacob Lucas, Michael Werth, Brandoch Calef, Ian McQuaid, and Justin Fletcher. Machine learning for quality assessment of ground-based optical images of satellites. *Optical Engineering*, 59(5):051403, January 2020.
- [15] Trent Kyono, Jacob Lucas, Michael Werth, Ian McQuaid, and Justin Fletcher. Towards Real-Time Image Reconstruction Using Deep Neural Networks. In *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference*, Maui, HI, September 2020.
- [16] A. Labeyrie. Attainment of Diffraction Limited Resolution in Large Telescopes by Fourier Analysing Speckle Patterns in Star Images. *Astronomy and Astrophysics*, 6:85, May 1970.
- [17] T. W. Lawrence, D. M. Goodman, E. M. Johansson, and J. P. Fitch. Speckle imaging of satellites at the U.S. Air Force Maui Optical Station. *Applied Optics*, 31(29):6307–6321, October 1992.
- [18] Jacob Lucas, Trent Kyono, Michael Werth, Nicole Gagnier, Zackary Endsley, Ian McQuaid, and Justin Fletcher. Estimating Satellite Orientation through Turbulence with Deep Learning. In *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference*, Maui, HI, September 2020.
- [19] Jacob Lucas, Michael Werth, Trent Kyono, Ian McQuaid, and Justin Fletcher. Automated Interpretability Scoring of Ground-Based Observations of LEO Objects with Deep Learning. In *2020 IEEE Aerospace Conference*, pages 1–7, March 2020.
- [20] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing*

DISTRIBUTION A. Approved for public release: distribution is unlimited. Public Affairs release approval AFRL-2022-3898.

The views expressed are those of the author and do not necessarily reflect the official policy or position of the Department of the Air Force, the Department of Defense, or the U.S. government.

- Systems*, volume 32. Curran Associates, Inc., 2019.
- [21] Matthew Phelps, J. Zachary Gazak, Thomas Swindle, Justin Fletcher, and Ian McQuaid. Inferring Space Object Orientation with Spectroscopy and Convolutional Networks. *AMOS (Sept. 2021)*, 2021.
 - [22] R. Scott and P. Eng. Medium and small aperture speckle interferometry for geostationary on-orbit-servicing space surveillance. In *Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*, page 43, 2017.
 - [23] Robert Scott. *A Method for Optical Tracking of On-Orbit Servicing Operations in Geostationary Orbit Using Speckle Interferometry*. PhD thesis, Carleton University, 2016.
 - [24] Robert L. Scott and Alex A. Ellery. Speckle interferometry tracking of on-orbit servicing in geostationary orbit. *Journal of Spacecraft and Rockets*, 53(3):433–447, 2016.
 - [25] Robert Lauchie Scott and Alex Ellery. An approach to ground based space surveillance of geostationary on-orbit servicing operations. *Acta Astronautica*, 112:56–68, 2015.
 - [26] Paul D. Shubert. Satellite imaging with speckle interferometry. In *Digital Image Synthesis and Inverse Optics*, volume 1351, pages 575–587. SPIE, November 1990.
 - [27] Ryan Swindle, Zach Gazak, Matthew Phelps, Ian McQuaid, and Justin Fletcher. Learned satellite identification using low-resolution spectropolarimetry. In *Polarization: Measurement, Analysis, and Remote Sensing XV*, volume 12112, pages 110–119. SPIE, 2022.
 - [28] Michael Werth, Trent Kyono, Jacob Lucas, Ian McQuaid, Michael Brannon, and Justin Fletcher. Quality-Weighted Iterative Deconvolution (QWID). In *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference*, Maui, HI, September 2020.
 - [29] Douglas Woodward, Celeste Manughian-Peter, Timothy Smith, and Elizabeth Davison. Pixelwise Image Segmentation with Convolutional Neural Networks for Detection of Resident Space Objects. 2021.

DISTRIBUTION A. Approved for public release: distribution is unlimited. Public Affairs release approval AFRL-2022-3898.

The views expressed are those of the author and do not necessarily reflect the official policy or position of the Department of the Air Force, the Department of Defense, or the U.S. government.