Exploiting Spatial Information in Raw Spectroscopic Imagery using Convolutional Neural Networks

J. Zachary Gazak
Air Force Research Laboratory
Justin Fletcher
United States Space Command

Ryan Swindle, Ian McQuaid
Air Force Research Laboratory

ABSTRACT

The detection of closely spaced artificial satellites informs critical tactical decisions for space domain awareness (SDA). One method, spectroastrometry, exploits differences in spectral energy distributions to detect the presence of multiple objects below the resolving power of a telescope. In this work we present a spectroastrometric approach leveraging convolutional neural networks (CNNs) to detect the presence of multiple objects in high resolution spectroscopic observations unresolvable through conventional imaging. This forumlation is tested against a simulated dataset modeled after the Spica II spectrograph. Spica II–an optical instrument mounted on a trunnion port of the 3.6m AEOS telescope on Haleakalā–is being used to build an on sky set of images to validate our approach. Using a dataset of 300,000 simulated observations, we train a CNN to perform classification between CSOs and single objects across distributions of on sky separations, satellite classes, and atmospheric parameters. We quantify the theoretical limits of our approach across differences in CSO spectral energy distribution, relative target magnitudes, and physical target separation at geostationary orbit.

1. INTRODUCTION

By exploiting the wavering flux centroid of one dimensional spatial profiles measured along the dispersion direction of a raw spectrum, astronomers have extracted salient spatial information from spatially unresolved phenomena [1] [14]. This technique−spectroastrometry−allows the detection of spatial structure well below the seeing limit of a ground-based telescope in situations where that structure encodes differing spectral distributions (Fig. 1) [14]. One situation meeting this condition is the reflection of light by spatially non-symmetric material distributions−i.e., resident space objects (RSOs).

Limitations of spectroastrometry have slowed the adoption of this technique. First, instrumental and atmospheric effects, including optical aberrations, dispersion, and alignment errors introduce challenging geometries onto raw imaged spectral frames, such that measuring the centroid of 1D spatial slices is not straightforward. Second, detecting closely spaced objects (CSOs) via spectroastrometry in the field of Space Domain Awareness (SDA) is difficult due to similarity in most RSO spectra. Spectral differences are created by reflection physics of materials and geometries from different satellite designs, with the bulk spectral signal appearing solar. Finally, the separation of two objects in a 2D spectrum is a function of their geometry but also the alignment of the spectrograph slit. An instrument oriented 90° w.r.t. inclination will provide no measurable signal as component spectra fall on top of one another. If orientation is known a priori, detecting the presence of a nearby object is moot.

The advent and success of Convolutional Neural Networks (CNNs) in the field of computer vision offers solutions to most limitations of spectroastrometry. First, instrumental and atmospheric effects are smooth functions of telescope pointing and can be easily fit out by a CNN given enough training data. Second, we need not be concerned with specific physical causes of a spectroastrometric signal. Relaxing that constraint allows the construction of a simple

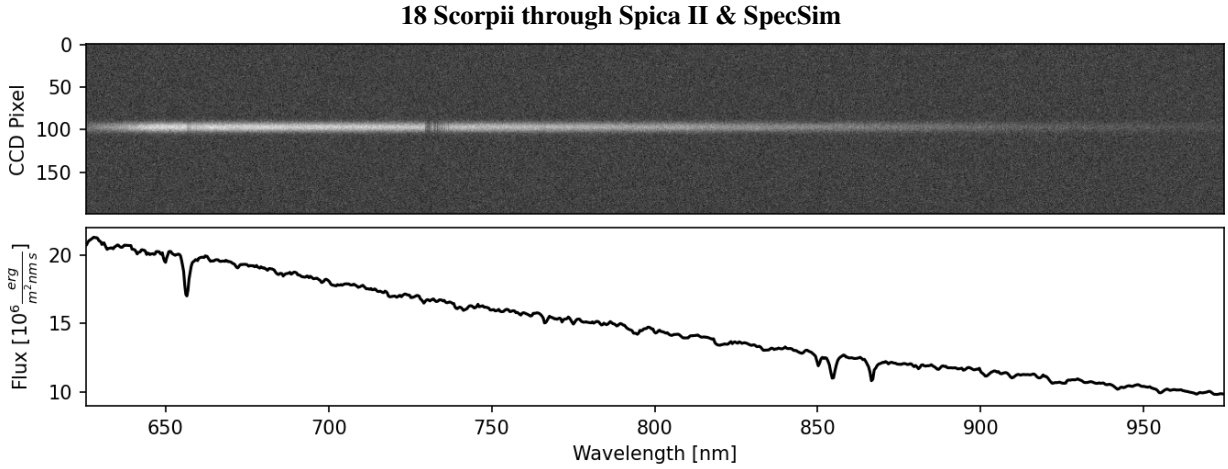

Fig. 1: **Top panel**: A simulated raw Spica II CCD observation of 18 Scorpii in the red bandpass. 18 Scorpii is a close solar analog [3], such that this spectrum and CCD frame are typical of resident space objects reflecting solar radiation. This frame is simulated using as input the measured spectrum of 18 Scorpii (bottom panel) [2] passed through our FPA simulation code, SpecSim (§4.2). Given an input spectra, SpecSim models atmospheric, telescope, and instrument effects to produce a raw output frame. We utilize raw frames such as these to train the CNN used in this paper.

classifier−single or double object−and frees a CNN's task capacity to focus on detecting the presence of signals mapping to these classifications. Finally, multiple exposures of a single target under different slit rotations can overcome the problem of unknown inclination, and may provide a measurement of that inclination.

We have developed an end-to-end prototype for detecting CSOs in geostationary orbit including a convolutional neural network to extract information from raw spectroscopic frames (§3), a modeling package for simulating raw focal plane output (§4), and a training and validation framework (§5).

In this study we focus on an existing optical spectrograph, Spica II, mounted on the trunnion port of the AEOS 3.6m telescope atop Haleakalā. While this study does not contain on sky validation of our computer vision solution, this instrument is currently being used to build a dataset necessary to do so. We utilize it's design configuration to simulate the dataset on which this work is based (Fig. 1, §4.2).

2. RELATED WORKS

Astronomical applications of CNNs cover an important problem set in deep learning−the extraction of information in the low signal to noise regime. Problems of low signal are arguably the most difficult in the field of computer vision, but nonetheless critical as complexity of the SDA environment increases and the accumulation of data accelerates. These two factors produce tactical scenarios untenable by human analysts and opportunities for deep learning solutions. In the astronomy community, similar problems of data storage and analysis are emerging as large survey projects come online [9][7].

This work is motivated by an application of neural networks to ground-based electro-optical imagery in the SDA domain [5]. Here we expand on the success of that work into spectroscopic imagery.

3. FORMALIZATION

3.1 Deep Learning

Deep neural networks are large layered collections of parameterized nonlinear functions tuned to encode an optimized mapping between input vectors and target output vectors through the back-propagation of errors between those target outputs and network predictions[10]. The specific geometry of those layers is called a network's architecture, and is problem specific.

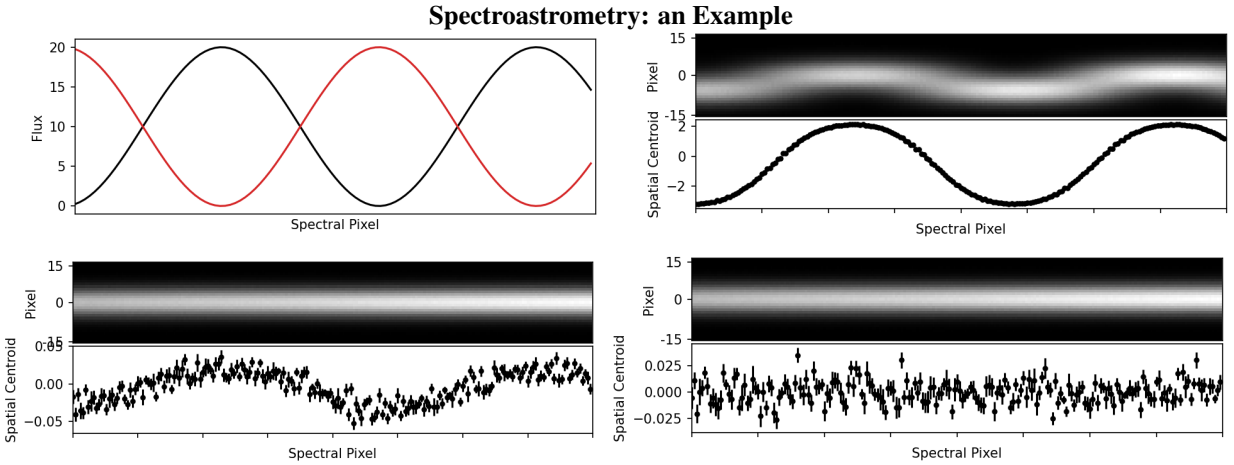

Fig. 2: A contrived spectroastrometry example. **Top left**: two synthetic sinusoidal spectra offset by π. **Top right**: (top) a characteristic slice of a raw 2D spectroscopic image simulated from the two sinusoidal spectra for the Spica II spectrograph using SpecSim (See §4.2) in $1''.5$ seeing separated by $1''.0$ (250m at ∼GEO orbit), (bottom) centroid measurements and errors for each vertical slice of the above 2D image. Note the obvious, visual separation of these objects and the signal dominated sinusoidal motion of the pixel centroid along the spectral axis. This motion of the centroid is the spectroastrometric signal. **Bottom left**: same plots repeated for a separation of $0''.01$ (2m at ∼GEO orbit). Despite being far below the $1''.5$ atmospheric seeing, the spectroastrometric signal remains significant against the uncertainty of the measured centroids. **Bottom right**: same plots repeated for objects with no ($0''.0$) separation. At this point, the spectroastrometric signal disappears.

For this work, we adopt a deep Residual Attention Network (AttnResNet) architecture [13]. This network layers multiple attention blocks within a Residual Network [6], allowing the Residual Network to develop location-aware features. The stacked structure of the Residual Attention Network develops different types and scales of attention during training [13]. The network reduces task complexity by learning a combination of different locations and scales of attention best suiting a classification task [11]. We hypothesize that this formulation suits spectroscopy, a field in which a minority of observed pixels contain a majority of the pertinent information (see Fig. 1), but where outlying regions may still contribute.

3.2 Performance Quantification

For classification problems, precision and recall are measured on a per-class basis, and are defined as

$$\text{precision}_\text{C} = \frac{N_{TP_C}}{N_{TP_C} + N_{FP_C}} \tag{1}$$

$$\text{recall}_\text{C} = \frac{N_{TP_C}}{N_{TP_C} + N_{FN_C}} \tag{2}$$

Where N_{TP_C} gives the number class C examples correctly classified, N_{FP_C} is the number of examples incorrectly classified as C, and N_{FN_C} represents class C examples which are improperly classified. The F_{1_C} scores, representing a harmonic average between precision and recall, are then calculated as

$$F_{1_C} = 2 * \frac{\text{precision}_\text{C} * \text{recall}_\text{C}}{\text{precision}_\text{C} + \text{recall}_\text{C}} \tag{3}$$

Accuracy is calculated across the entire classifier, and is defined as

$$\text{accuracy} = \frac{N_{TP}}{N_{TP} + N_{FP}} \tag{4}$$

For completeness we report the set of precision, recall, and F_1 scores as well as accuracy for the full experiment. In the analysis of model performance in §5, we report accuracy only within the set of examples of paired CSO objects which allows for an exploration of the physical conditions under which model accuracy falters.

4. DATASETS

This work utilizes high fidelity simulated data. For each of 300,000 frames, we sample one or two satellite models from a set of 152, place those models in random orientations and orbits, calculate their reflection radiometry (§4.1), and propagate that signal through simulated atmosphere, telescope, and instrument physics (§4.2).

We sample from a set of 152 satellite models. Each frame draws one or two models from this set, places those models at a random orientation and orbit, and computes the reflected spectrum and raw spectrograph focal plane array (FPA) output. The M&S environment allows significant control over simulated data. An input solar spectrum is supplied as input, allowing datasets to span multiple spectral resolutions.

Set	Parameter	Value
Scene Geometry	Angular separation [″]	U(0.0, 1.5)
	Slit-Inclination angle [°]	0.0
	Range [km]	35800
	Objects per frame	U([1, 2])
Atmosphere [12] [8]	Seeing FWHM	U(0.4, 2.0)
	Airmass	U([1.0, 1.5, 2.0, 2.5, 3.0])
	PWV	0.5
	Observatory altitude [m]	3060
Instrument	Grating [grooves/mm, blaze]	150, 800
	Pixel Pitch [μm]	20
	Spectral resolution [nm]	6.5
	Dark Current [e^{-1}/ph/s]	0.005
	Read Noise [e^{-1} RMS]	3
	Bias [ADU/pix]	600
	PSF mode	Gaussian
	Gain	1.0
Exposure	Target N_{ADU} [ph/pixel column]	1000

4.1 Radiometry

An input solar spectrum is propagated through existing, state of the art radiometric software. Inputs into this simulation code include target model and orientation, orbit, time of observation, material BRDF information, and sensor location. The output is a one dimensional, integrated spectrum of the light reflected towards the sensor location.

4.2 Focal Plane Array (SpecSim)

The reflected spectrum is propagated through our in house spectroscopic focal plane array (FPA) modeling code, SpecSim. SpecSim is a python codebase with easily configurable atmospheric, telescope, and instrument definitions. Parameters and values adopted/varied in this study are shown in Table 4.

We adopt the Cerro Paranal Advanced Sky Model [12] [8] to provide atmospheric transmission and emission based on parameters including precipitable water vapor (PWV), airmass, and observatory altitude.

5. EXPERIMENTS

In §5.1 we discuss the overall performance of this model. In §5.2−§5.4, the discussion focuses on model performance against apparent separation, magnitude, and spectral differences. For these sections, we consider only examples containing CSO signatures.

Class	Precision	Recall	F1	Accuracy
Single	0.921	1.000	0.959	-
CSO	0.999	0.909	0.952	-
Overall	-	-	-	0.956

Table 1: Tabulated performance parameters. In classification problems, precision, recall, and F_1 scores are calculated on a per-class basis, while accuracy represents the overall performance of the model.

5.1 Overall Classifier Performance

The classifier trained for this work delivers an impressive 95.5% accuracy (Table 1), with strong performance on both single satellite and CSO examples. We visualize the performance of this model in a confusion matrix (Fig. 3), where values on the diagonal represent correct classifications.

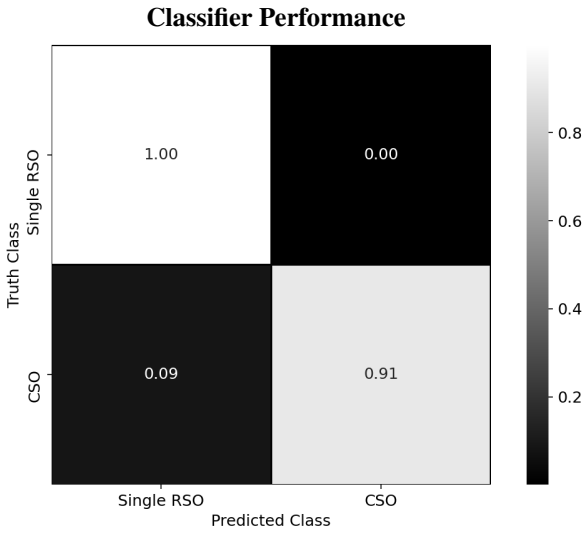

Fig. 3: Confusion matrix for the classifier in this work. Top left and bottom right squares represent correct classifications.

5.2 Apparent Separation

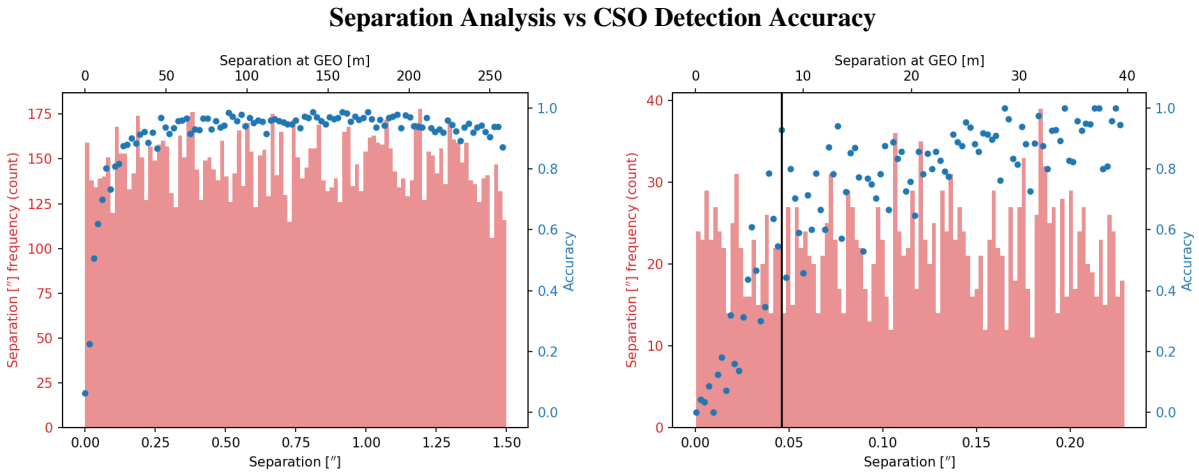

Fig. 4: CNN accuracy as a function of angular separation of CSO targets.

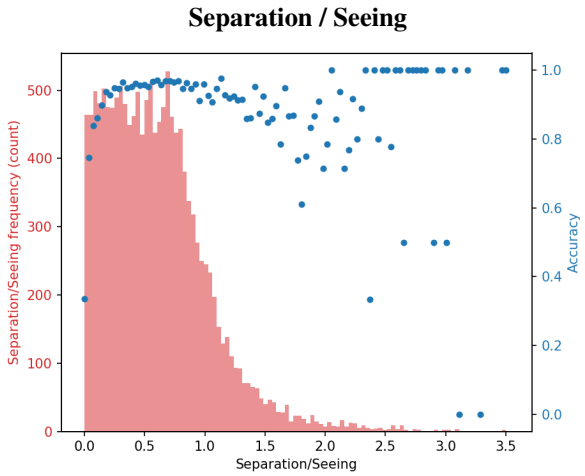

Fig. 5: CNN accuracy as a function of the ratio between target angular separation and seeing (blue dots), and population of targets within given ratios. The CNN performs with high accuracy at target separations well below that of the seeing FWHM.

The theoretical limit for the separation of two distant objects in the presence of real optics is directly related to the overall size of the optics. This diffraction limit goes as $\sim \lambda/D$, where λ is the wavelength of observation and D is the diameter of the primary optic. This limit is generally dominated by atmospheric seeing. Seeing, introduced by atmospheric turbulence, forces point sources into smooth, wide point spread functions(PSFs) in long exposures. In general, resolving below an atmosphere's seeing occurs at great expense, by launching space-based telescopes or using complex and costly adaptive optics systems. Spectroastrometry avoids these limitations by focusing on repeated, single measurements−flux centroids− instead of requiring the reconstruction of a resolved image.

AttnResNet accuracy across all separations ($0 - 1''.5$) and a narrow range ($0 - 0''.25$) is plotted in Figure 5.2. At separations greater than $0''.25$ (physical separations at GEO of $>$ 40m), AttnResNet performs with high accuracy. As expected, this accuracy falls off as separation decreases. Still, the model remains performant below $0''.045$, the diffraction limit of a 3.6m telescope.

In Figure 5.2 we provide a visualization of performance against the ratio between angular separation and atmospheric seeing, and note that AttnResNet performs with high accuracy well below the seeing limit of an observation.

5.3 Delta Magnitude

In CSO discrimination techniques−and in the field of SDA−the difference in magnitude (Δ_m) is of critical importance. Tactically relevant technologies must be able to discriminate between objects with significant differences in brightness, for example, a small interloper near a larger, high priority satellite. In Figure 5.3 we plot performance against overall distribution of simulated apparent magnitudes, and more interestingly, across Δ_m. Because this study increases simulated exposure time to increase spectral S/N, magnitude is not expected to affect performance. In the case of Δ_m, a clear trend is present. AttnResNet performs at highest accuracy on targets separated by up to $\Delta_m = 2$, and drops slowly off after that, remaining performant out to the limit of this study, $\Delta_m = 8$. Future work will simulate more observations at higher magnitude differences to more adequately measure performance in this regime.

5.4 Spectral Distribution

We adopt the Kullback-Leibler (KL) distance,

$$D_{KL}(f_1||f_2) = \sum_{\lambda} f_1(\lambda) \ln \frac{f_1(\lambda)}{f_2(\lambda)} \tag{5}$$

as a quantitative measurement of difference between two spectral flux vectors $f_{1,2}$ both normalized to unity. Because $D_{KL}(f_1||f_2) \neq D_{KL}(f_2||f_1)$, we always assign f_1 to be the brighter object as measured before normalization.

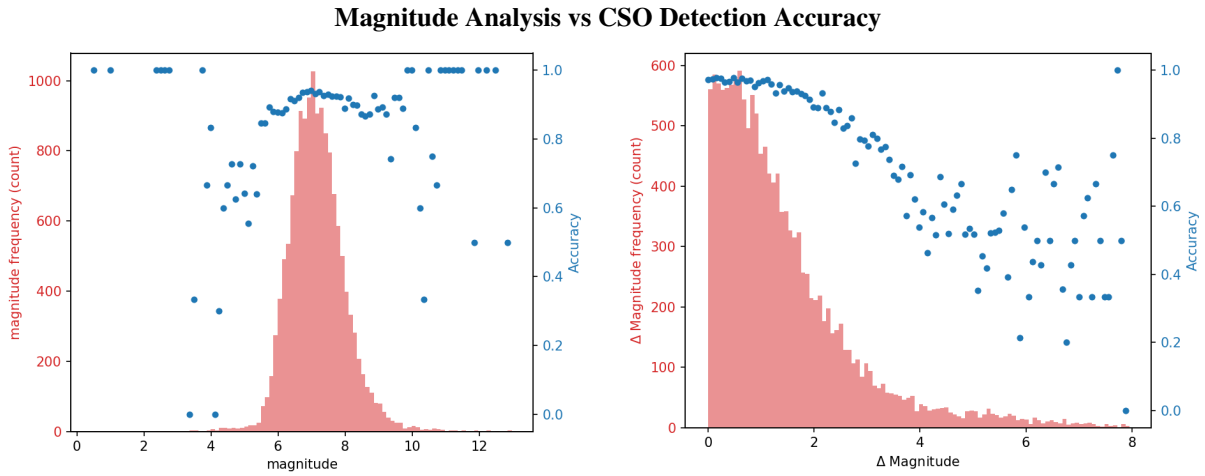

Fig. 6: CNN accuracy as a function of apparent and Δ magnitude of CSO targets. The effects of magnitude are minimal in this work as simulated exposure time was adjusted to return a consistent number of counts across the dataset.

In the dataset generated for this study, the effects of D_{KL} are overwhelmed by examples which have high angular separation and low D_{KL}, masking out an expected drop in overall classifier accuracy as $D_{KL} \to 0$. At separations near or above the atmospheric seeing, even CSOs with identical spectra are detectable (Fig. 1). The effect as $D_{KL} \to 0$ is more apparent when plotted against angular separation as in Fig. 7.

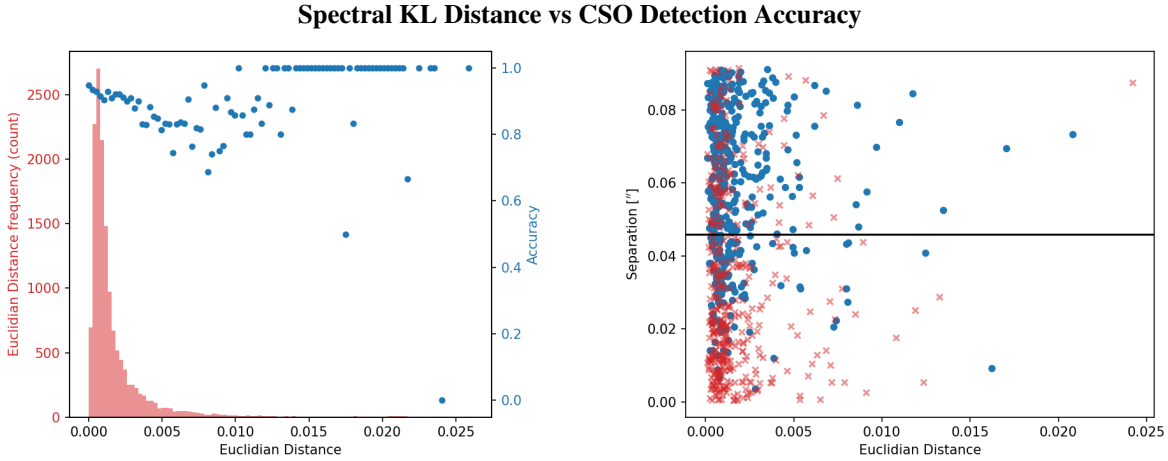

Fig. 7: Left: Population histogram of CSO simulations vs KL distance between their component spectra (red, left axis, §5.4), and model accuracy across those bins (blue, right axis). Paradoxically, the results show increasing accuracy correlated with decreasing spectral difference. In practice, however, this is a product of sampling bias towards similar reflection spectra, and a product of the difficulty of spectroastrometry for reflection spectra discussed in §1. Right: True positive CSO detections (blue circles) overplotted with false positive classifications as single objects (red X's), with the diffraction limit for a 3.6m telescope denoted as a solid horizontal line. When angular separation is large enough, a CNN can detect CSOs even with nearly identical spectral signatures. As objects move closer together, dropout based on KL distance increases dramatically as expected.

6. CONCLUSIONS AND FUTURE WORK

In this report we have demonstrated an Attention Residual Network CNN framework for detecting CSOs in geostationary orbit by exploiting spectroastrometric signals. The analysis of this trained AttnResNet shows excellent per-

formance to close separations, detecting ∼50% of CSOs simulated at separations of ∼5 meters at geostationary orbit and observed with $N_{counts} \simeq 1000$ per spectral pixel column by a low resolution spectrograph on a 3.6m telescope. For targets down to $m_V \sim 15$, this level of signal is achievable in minutes using the existing Spica II instrument simulated in this work.

Opportunities for future spectroastrometry work abound. A subset of particular interest:

On-sky Validation: Using Spica II, an observed spectral catalog is being populated to evaluate this approach with and without transfer learning on real data.

Inclination extraction: We will design a multi predictor head CNN capable of consuming multiple spectra at varying slit-inclination angles and predicting object class and, in the case of a CSO detection, projected position angle of the separation axis.

Single object articulations: In practice, RSOs on the size scale of the limiting separation distance of a spectroastrometric device (∼5m @ GEO for Spica II) will contain their own spectroastrometric signals. This offers both challenges (disentangle a single large object from two closely spaced ones) and opportunities (extract salient target information from a single object) which both warrant exploration.

Rotational spectroastrometric imaging: Explore the additional information gain by continuous rotation of a slit against targets of interest. By collating spectra at many rotation angles, a dataset will contain spectroscopic information vs radial extent and position−in effect, a hyperspectral image signature extracted from spatially unresolved data. The extent to which this information can be used to inform the structure and material distribution of satellites should be explored.

Hardware opportunities: We plan to explore the effects of hardware (spectral resolution, bandpass, telescope size, location, etc) to develop a theoretical framework to utilize for future instrument designs.

Autonomous spectroscopy: Spectroscopic instruments tend to be high effort and low yield; requiring dedicated, trained observers and returning data covering a low percentage of observing time. Future instrument designs taking full advantage of the power of computer vision will require the opposite: low effort (or autonomous), high yield instruments capable of delivering the thousands of observations needed.

More generally, we hypothesize that spectroscopy is poised for considerable gains with the adoption of CNN driven technologies. In resolved and unresolved imaging, neural networks provide highly performant, real time solutions for existing technologies requiring hours or days-long data reduction and consumption by human analysts. Spectroscopic technologies, however, are largely absent. This is due to a non-intuitive information encoding: an expert in spectroscopy still operates on spatial intuition, while spectroscopic images resolve in the energy domain. Even reduced one dimensional spectra, while easier to understand, still contain information beyond what physics based algorithms can model. While the complexity of spectroscopic images largely precludes analyst-in-the-loop exploitation technologies, these images contain dense information regarding the material composition, construct, and orientation of RSOs. CNNs are particularly well suited to extract this information from properly assembled training datasets as deep neural networks do not require physics-based or human-intuition driven rule sets to map input data to tactically relevant information.

7. REFERENCES

[1] Jeremy A. Bailey. Spectroastrometry: a new approach to astronomy on small spatial scales. In Sandro D'Odorico, editor, *Optical Astronomical Instrumentation*, volume 3355 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 932–939, July 1998.

[2] Ralph C. Bohlin, Karl D. Gordon, and P. E. Tremblay. Techniques and Review of Absolute Flux Calibration from the Ultraviolet to the Mid-Infrared. *Publications of the Astronomical Society of the Pacific*, 126(942):711, August 2014.

[3] G. Cayrel de Strobel. Stars resembling the Sun. *AAPR*, 7(3):243–288, January 1996.

[4] Hilda Deborah, Noel Richard, and Jon Yngve Hardeberg. A Comprehensive Evaluation of Spectral Distance Functions and Metrics for Hyperspectral Image Processing. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 3224(6), June 2015.

[5] Justin Fletcher, Ian McQuaid, Peter Thomas, Jeremiah Sanders, and Greg Martin. Feature-Based Satellite De-

tection using Convolutional Neural Networks. *AMOS Conference Proceedings*, 2019.

[6] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. *arXiv e-prints*, page arXiv:1512.03385, December 2015.

[7] Z. Ivezic, T. Axelrod, W. N. Brandt, D. L. Burke, C. F. Claver, A. Connolly, K. H. Cook, P. Gee, D. K. Gilmore, S. H. Jacoby, R. L. Jones, S. M. Kahn, J. P. Kantor, V. V. Krabbendam, R. H. Lupton, D. G. Monet, P. A. Pinto, A. Saha, T. L. Schalk, D. P. Schneider, M. A. Strauss, C. W. Stubbs, D. Sweeney, A. Szalay, J. J. Thaler, J. A. Tyson, and LSST Collaboration. Large Synoptic Survey Telescope: From Science Drivers To Reference Design. *Serbian Astronomical Journal*, 176:1–13, June 2008.

[8] A. Jones, S. Noll, W. Kausch, C. Szyszka, and S. Kimeswenger. An advanced scattered moonlight model for Cerro Paranal. *Astronomy Astrophysics*, 560:A91, December 2013.

[9] Nicholas Kaiser, Herve Aussel, Barry E. Burke, Hans Boesgaard, Ken Chambers, Mark R. Chun, James N. Heasley, Klaus-Werner Hodapp, Bobby Hunt, Robert Jedicke, D. Jewitt, Rolf Kudritzki, Gerard A. Luppino, Michael Maberry, Eugene Magnier, David G. Monet, Peter M. Onaka, Andrew J. Pickles, Pui Hin H. Rhoads, Theodore Simon, Alexander Szalay, Istvan Szapudi, David J. Tholen, John L. Tonry, Mark Waterson, and John Wick. Pan-STARRS: A Large Synoptic Survey Telescope Array. In *Survey and Other Telescope Technologies and Discoveries*, volume 4836 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 154–164, December 2002.

[10] Yann Lecun, Leon Bottou, Genevieve Orr, and Klaus-Robert Müller. Efficient backprop. 08 2000.

[11] Volodymyr Mnih, Nicolas Heess, Alex Graves, and Koray Kavukcuoglu. Recurrent Models of Visual Attention. *arXiv e-prints*, page arXiv:1406.6247, June 2014.

[12] S. Noll, W. Kausch, M. Barden, A. M. Jones, C. Szyszka, S. Kimeswenger, and J. Vinther. An atmospheric radiation model for Cerro Paranal. I. The optical spectral range. *Astronomy Astrophysics*, 543:A92, July 2012.

[13] Fei Wang, Mengqing Jiang, Chen Qian, Shuo Yang, Cheng Li, Honggang Zhang, Xiaogang Wang, and Xiaoou Tang. Residual attention network for image classification. *CoRR*, abs/1704.06904, 2017.

[14] E. Whelan and P. Garcia. *Spectro-astrometry: The Method, its Limitations, and Applications*, volume 742, page 123. 2008.