Simultaneous Detection, Recognition, and Localization of Geosynchronous Satellites from Ground Based Imagery

J. Zachary Gazak, Ryan Swindle, Matthew Phelps, Justin Fletcher
USSF Space Systems Command (A&AS)

ABSTRACT

The protection of high value government and commercial satellites is predicated on the detection, recognition, and localization of active and defunct satellites, rocket bodies, and debris across all orbital regimes above Earth. In the bulk of those orbital regimes, recognition has lagged the needs of space domain awareness stakeholders, as target distances preclude resolved imagery from ground-based observatories. Recent research has demonstrated that the rich, distance-invariant information encoded into the spectrum of an object's reflected sunlight can fill this need. This concept, named SpectraNet, introduced deep learned spectroscopy as a solution bypassing the dearth of exquisite knowledge of the physics governing reflection against complex materials. By training deep neural networks against a baseline of observations, SpectraNet learns the information content linking spectra to object identity. In this work we introduce and demonstrate a significant advancement in the SpectraNet pedigree which solves a number of limitations with the original proof of concept. By utilizing multi-spectral imaging sensors instead of a traditional longslit spectrograph, this work eliminates the need for a priori knowledge of target location. Further, targets are simultaneously detected and identified, and the stellar background allows for the extraction of astrometric information required for target localization. Finally, this technology enables proliferation by utilizing inexpensive, off the shelf components and small telescopes. In this work we describe the learned backbones for detection and recognition, end to end performance in simulation, and initial on sky observations.

1. INTRODUCTION

Effective space domain awareness (SDA) enables the safe and equitable use of orbital regimes surrounding Earth. Without exquisite catalogs of the growing population of satellites in space, all decisions−from policy to launch to intervention−assume significant risk. While significant advancements in sky coverage and rapid updating of catalogs have been realized in the last few years, the ability to enrich SDA catalogs with identifying information has lagged behind.

To detect objects and maintain catalogs, global networks of ground based optical sensors collect continuously, generating many thousands of images per night. Detection of targets withing those images via deep neural networks has proven to be highly effective in extracting object positions within these images [14, 9, 33]. Tracking objects requires the localization of detections in the greater universe; that is, the extraction of a detection's celestial latitude and longitude (Right Ascension and Declination, respectively). This location is accomplished by detecting stars ubiquitous to all space imagery [13, 15] and matching unique patterns of those stars to previously measured catalog locations [25]. The result provides a mapping from pixel to celestial coordinates.

Finally, as orbital regimes grow increasingly congested and spacecraft increasingly mobile, the ability to identify a spacecraft from information contained in an observation becomes critically important. Relying on correlation of observed object locations to existing catalog entities breaks down in crowded regimes, during and after rendezvous and proximity operations, and following significant maneuvers. The ability to recognize observed satellites has been demonstrated [18] and shown to be effective throughout orbit seasonality [20]. However, the technologies employed in those works required a priori knowledge of a target's location. In this work we introduce the concept of imaging spectroscopy to enable simultaneous detection, localization, and recognition of satellites in spatially unresolved space imagery.

The main contributions of this work include:

- Two large datasets of simulated spectrally enriched space imagery of 110 satellite classes observed over six months, one for multispectral cameras, and the second for multi aperture telescopes.

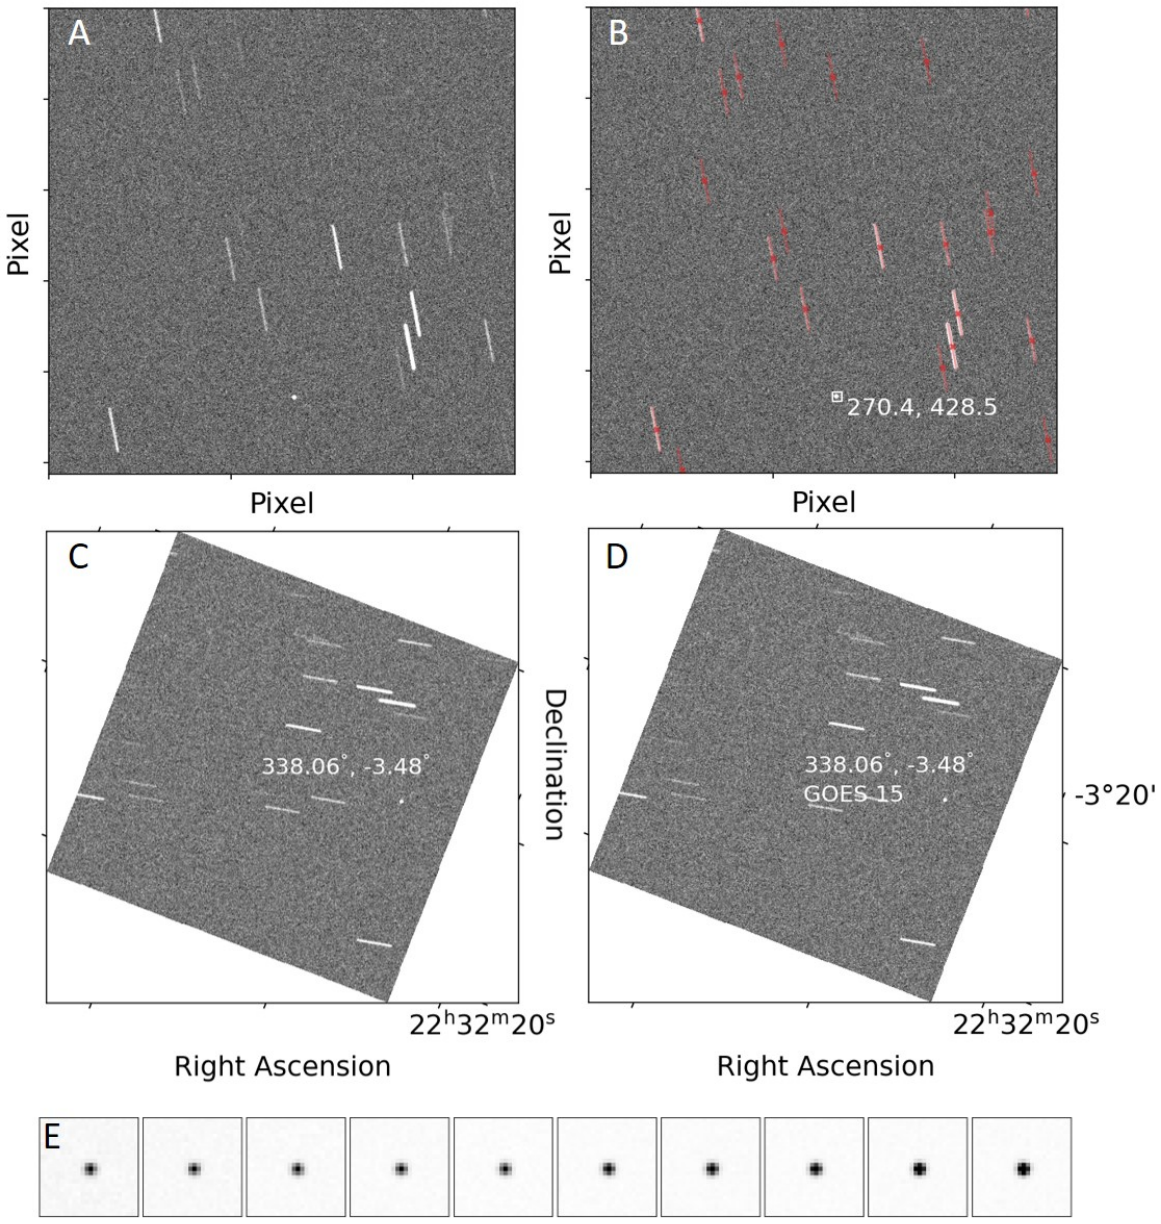

Fig. 1: Simultaneous detection, localization, and recognition. Frame A: A single channel raw frame from a Raven class telescope. Frame B: Star and object detections in pixel space. Frame C: Localization information extracted from star positions. Frame D: Object identifying information extracted from spectral channel of image. Frame E: Cutout spectral channels for detected object used in the recognition step.

- Astronomical observations through multispectral cameras on Raven class telescopes.
- An analysis on the relative effectiveness of multispectral vs multi aperture systems for simultaneous detection, localization, and recognition of satellites from affordable, easy to deploy ground based equipment.

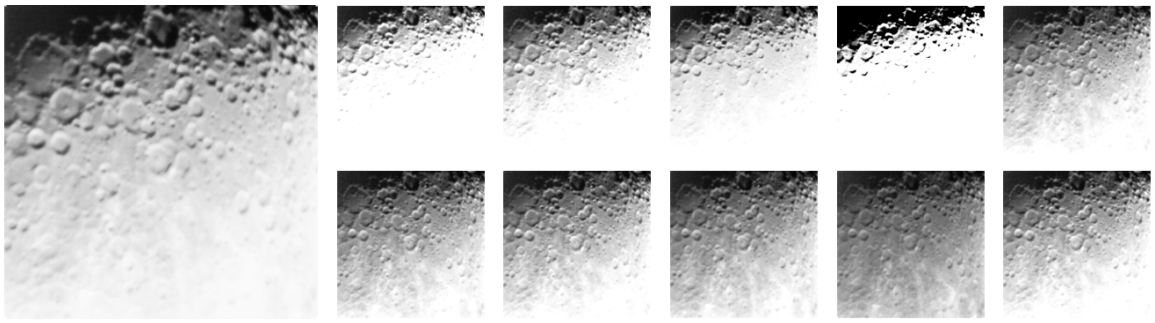

Fig. 2: The large left image is a full frame, 2048 pixel square image of the moon captured through a 0.35 m aperture Raven class telescope with a Silios Toucan multispectral camera. To the right, the extracted ten channels from 400 to 900 nanometers, organized from left to right, top to bottom. Spectral differences in the reflected sunlight are apparent to the naked eye and demonstrate the multi-filter capability of large format Bayer filters.

2. PREVIOUS WORK

The development of a compact and affordable telescope system allowing for detection, localization, and recognition of spatially unresolved objects in space imagery requires a number of exquisite technologies. First, in the detection step, raw frames from a telescope are processed to extract satellites. This step yields the pixel coordinates of satellites in the frame. Second, in the localization step, field stars are detected and those detections are run through an astrometry solver. This step enables the conversion of pixel-space detections into right ascension and declination coordinates, such that the object is localized in space. Finally, spatial information is extracted from the raw observations and used to classify the observation as a particular satellite, satellite family, or flag as an unknown or immeasurable entity. This step enriches SDA catalogs with identifying information and enables stakeholders to make informed decisions about on orbit scenarios.

The performance of algorithms in each step is critically important; in the detection step, under-performing algorithms can miss objects or create a flood of false positives–queuing costly searches or followup observations. For localization, detecting too few stars can make astrometric fitting impossible, in which case the data is useless to the SDA community. Finally, not conducting object classification relies heavily on orbit correlation–a challenge in increasingly congested orbital regimes–while false classification can raise unnecessary alarms.

An explosion in learned solutions, those which leverage trained deep neural networks instead of hand designed features, has advanced the state of the art in each detection, localization, and recognition steps. We discuss these advances below.

2.1 Detection

The cornerstone work of Fletcher et. al. 2019 [14] demonstrated the effectiveness of computer vision to the detection of spatially unresolved objects in high contrast space imagery. We utilize SatSim [9], a fully differentiable high fidelity space scene generator for studying unusual scenarios such as breakups and collisions, and for enabling advanced training scenarios using pre-training, transfer learning, etc. Previous research explored the application of compact detection algorithms to the detection problem [22], and demonstrated the use of segmentation as a technique to extract detections instead of object detection models [33].

Models are developed and trained with PyTorch [27] and experiments are controlled and monitored with MLFlow [34].

2.2 Localization

StarNet [13] and StarCSP [15] offer learned solutions to detecting streaked stars without using error prone information in the header to dictate streak length and orientation. These pixel star detections are fed into a plate solving code, such as Astrometry.net [25]. In this plate solver, hash indexed quad patterns of stars allow efficient solving of the necessary transforms from pixel coordinates to sky coordinates. This critical step allows the pixel coordinates of detected space objects to be mapped into real world coordinates useful in orbit determination and propagation.

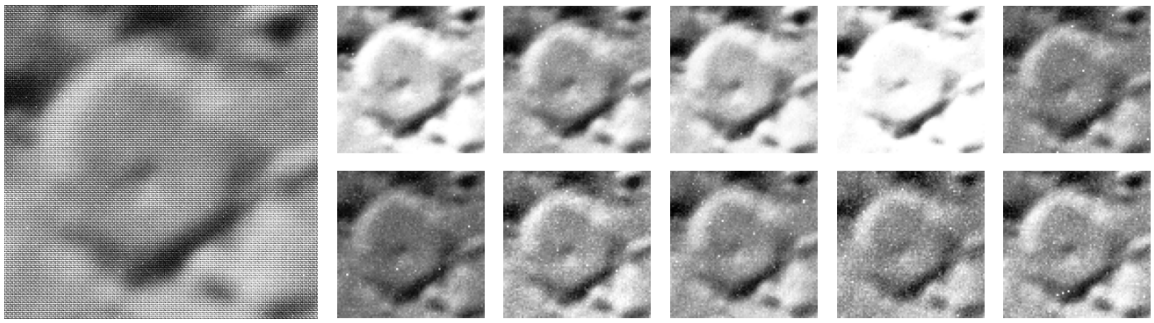

Fig. 3: A zoomed region from the top left of Figure 2. The large left image represents raw camera output, with the visible Bayer pattern. The smaller images to the right are extracted channels. Each extracted channel suffers a loss in spatial resolution over the raw camera footage, but the addition of spectral information provides a salient additional source of information.

2.3 Recognition

With a well understood radiative transfer model, spectroscopic observations can be reproduced with high fidelity [11]. Such models provide detailed physical underpinnings of the observed system [24]. In the study of stellar emission spectra, physics based models can extract the parameters of a star from a single observations [16]. In SDA, asteroid, and planetary astronomy, where incident sunlight is reflected, absorbed, and re-emitted by solid materials [8, 12], the underlying physics is governed by statistical thermodynamics and quantitative model-based spectroscopy is less applicable.

In the field of SDA, anthropogenic materials impart unique spectral signatures onto reflected light and often produce highly specular bidirectional reflection distribution functions (BRDFs) on which material orientation is a significant component. With new materials launched continuously and weathered in earth orbit, models of spectral reflectance are difficult to measure [10]. An important body of work is continuously narrowing the gap between observations and models, but the work is difficult and progress is slow [6].

In the meantime, the effectiveness of computer vision for the recognition satellites in regimes where spatial information is lost has been demonstrated on sky [18], in simulation for long baselines [20], and for large ranges of spectral information density [19]. This, paired with a growing number of spectroscopy ready SDA telescopes [23, 26, 28], makes learned recognition for SDA enrichment a pressing need.

3. THIS WORK

In this work we concentrate on the recognition step in SDA pipelines. The novel aspect of this work is the utilization of sensing concepts which enable detection and localization. Unlike previous SpectraNet work [18, 31, 19, 20] which focuses on longslit spectrographs as the enabling technology, here we explore the use of imaging focal plane arrays designed to provide simultaneous spectral information. This distinction moves SpectraNet powered telescopes from a dedicated followup instrument, capable of recognition once an object is detected and an orbit is determined by other means, to a self contained unit of SDA in which the discovery and recognition steps are tightly coupled.

We test five model backbones, ResNet-18, -34, and -50 [21], HybridSN [30], and PhotConv [19]. Residual networks are celebrated for high performance and rapid training due to the application of residual skip connections and batch normalization. For each of the ResNet backbones we utilize pre-trained weights and modify the initial ResNet layer to accept our data cube shape, that is 10 x 10 spatial pixels and 10 spectral channels. HybridSN is a model built to enhance feature extraction from hyperspectral data cubes by combining 3D convolutions for spectral information and 2D convolutions for spatial features [30]. The model is especially effective on hyperspectral imagery of earth's surface; we apply the same technology to space imagery. Finally, we adopt the PhotConv model developed for few channel photometric measurements to train on spectral data cubes; this is done by compressing the spatial information in our detections to 1D arrays of 10 spectral channels.

One Month Challenging Dataset			
Network	Dragonfly	Toucan	Toucan Defocused
ResNet50	0.673 ± 0.006	0.218 ± 0.012	0.457 ± 0.020
ResNet34	$\mathbf{0.696 \pm 0.002}$	$\mathbf{0.247 \pm 0.003}$	0.468 ± 0.010
ResNet18	0.687 ± 0.004	0.238 ± 0.005	0.522 ± 0.009
HybridSN	0.664 ± 0.010	0.206 ± 0.010	$\mathbf{0.571 \pm 0.005}$
PhotConv1d	0.620 ± 0.001	0.100 ± 0.002	0.368 ± 0.015
One Month Simple Dataset			
Network	Dragonfly	Toucan	Toucan Defocused
ResNet50	0.776 ± 0.008	0.197 ± 0.032	0.615 ± 0.014
ResNet34	0.793 ± 0.007	0.307 ± 0.022	0.639 ± 0.026
ResNet18	0.796 ± 0.005	0.238 ± 0.036	0.635 ± 0.018
HybridSN	$\mathbf{0.827 \pm 0.002}$	$\mathbf{0.334 \pm 0.002}$	$\mathbf{0.743 \pm 0.004}$
PhotConv1d	0.747 ± 0.003	0.164 ± 0.005	0.506 ± 0.001

Table 1: Classifier accuracies across five networks, three instrument designs, and two difficulty scenarios. Maximum per-instrument performances appear in bold.

4. DATASETS

Models are trained on a dataset of 100,000 simulated examples, 10,000 validation and 10,000 test examples. Two class sets are used, one in which 110 classes include multiple populated satellite families. This challenging dataset contains six DirecTV designs, six AMC designs, five THEMIS, four NOAA, four DMSP, and so on. Within these families of satellites, differences are minimal as overall designs stay consistent, but models are trained to recognize each individual design. The second set of 70 classes contains a single entry from each satellite family. In both sets, roughly half of the satellites modeled articulate solar panels to maximize solar flux, and the dataset contains 11 rocket bodies (Figures 5 and 6).

In this work we simulate two hardware setups for collecting simultaneous many filter imagery; a multispectral camera and a multi aperture telescope. The multispectral camera allows for a simpler optical setup, in which a single telescope is equipped with a single instrument. However, the camera hardware introduces complexities at the image plane and camera efficiency in currently available commercial offerings is limited. In contrast, the drawback of multi aperture telescope lies in the hardware and software complexity, with accurate temporal triggering required across multiple co-mounted separate apertures on a single mount. At the image plane, each aperture sits behind a different filter and receives higher efficiency and non-adulterated point spread functions.

4.1 Toucan Multispectral Camera

In commercial color CCDs, simultaneous color information is collected using a Bayer filter [7] in which the focal plane is divided into two by two blocks containing one red, one blue, and two green filtered pixels. The specific design and matrix size of Bayer filters can be extended for industrial and scientific purposes, and current commercial offerings from Ximea [2, 3, 4] and Silios [1] include nine, sixteen, and twenty-five filter setups ranging from roughly 400 to 900 nanometers in bandwidth.

Previous work has demonstrated that five filters are adequate to learn spectral positive identification, and that increasing spectral coverage is more important than increasing number of filters [19]. For this reason we choose to simulate the Silios Toucan, which holds 10 filters between 400 and 900 nanometers in a 16 pixel Bayer pattern. In Figures 2 and 3 the properties of this multispectral camera are visualized. The example frames were captured using a Toucan camera on a 0.35 meter aperture Raven class telescope [32] deployed in Antofagasta, Chile.

While multispectral cameras are ideal for resolved imagery in well lit environments, the pixel layout can provide uneven spectral sampling and introduce artifacts in measured spectra for the narrow point spread functions (PSFs) of astronomical imagery. The natural distribution of a PSF dictates that central pixels receive higher counts than those

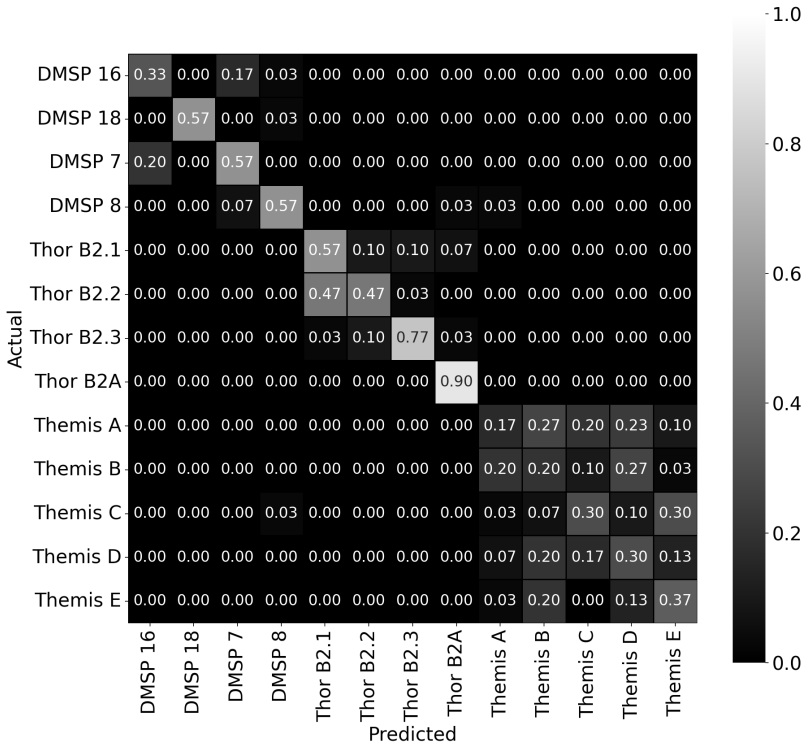

Fig. 4: Multiaperture Dragonfly performance across three families of satellites. In this subset confusion matrix, we isolate rocket bodies from the full dataset. For this reason, rows need not sum to unity. Here we demonstrate the performance of the Dragonfly performance on challenging satellite families, including Themis, Thor, and DMSP. In comparison to single and simpler objects (see for example Figure 5), satellite families suffer significant confusion as the spectral differences between classes are limited.

radially more distant. When each pixel is collecting unique spectral information, the received filter flux is dictated not only by the radiative transfer of the sun-satellite system, but also by the relative position of the pixel with respect to the imaged PSF.

For clarity, we refer to the multispectral experiments as Toucan−the sensor simulated in this work−for the remainder of this paper.

4.2 Dragonfly Multi Aperture Telescope

Multi aperture telescopes such as the Dragonfly system [5] allow for numerous co-aligned apertures on a single robotic telescope mount. Originally designed to observe extended, low surface brightness sources like galaxies, multi aperture systems also pave the way for multiple simultaneous photometric measurements. To that end, a number of systems have been developed in the SDA community to monitor multi color light curves of satellites [29, 28, 23]. These systems have up to four apertures, fewer than necessary for spectral recognition [19].

To enable a direct comparison between multi aperture and Toucan systems, we simulate a robotic telescope with the same filters as a Toucan camera of §4.1; 10 apertures with unique filters on each aperture. Due to the dedicated aperture per filter, multi aperture systems require smaller apertures per filter. We hypothesize that multi aperture systems will not suffer the same PSF-related issues as Toucan sensors. The drawback of multi aperture systems is the complexity of instrument design, alignment, and operation.

For clarity, we refer to multi aperture experiments as Dragonfly, after their namesake instrument, for the remainder of this paper.

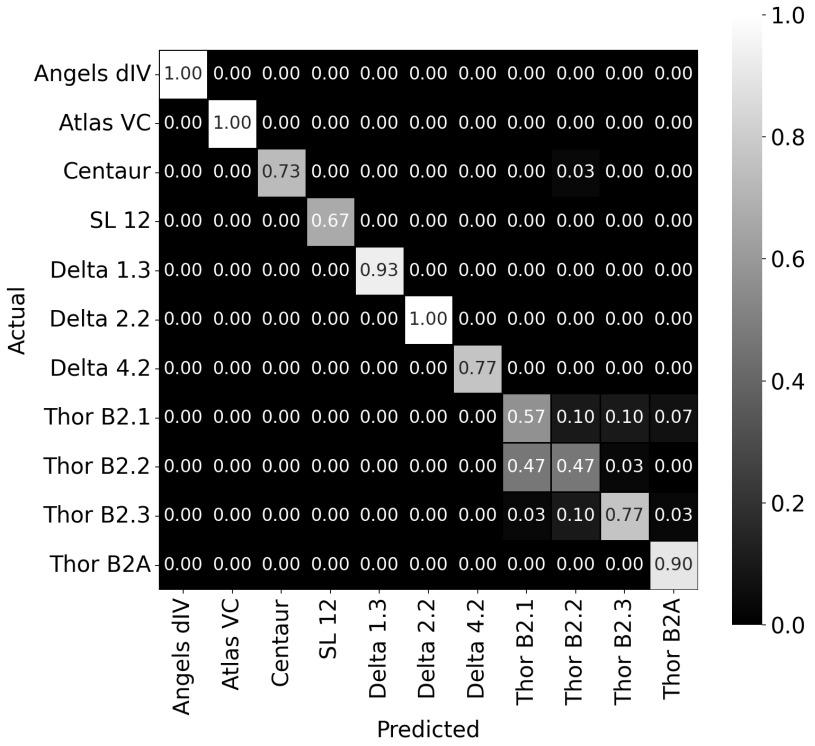

Fig. 5: Multiaperture Dragonfly performance across rocket bodies. In this subset confusion matrix, we isolate rocket bodies from the full dataset. For this reason, rows need not sum to unity. Here we demonstrate the performance of the Dragonfly performance on rocket bodies, including the Thor family. This figure is presented for comparison to Toucan multispectral performance in Figure 6.

5. EXPERIMENTS

In each experiment, detections are sliced out of simulated imagery to construct data cubes with 10x10 spatial size and 10 spectral channels. This is the exact mode to be used in an active SDA scenario, where the preceding detection step allows for cube extraction. For the Toucan systems, the frame is first reconstructed from 2048x2048 pixels to 512x512x10. The Dragonfly data is naturally in this latter cube form.

The compact nature of each example allows for rapid training even on large datasets. For each model backbone discussed in §3 we run a 200 trial hyperparameter experiment to maximize performance per model and data type. The results of that experiment are shown in Table 1.

Performance measurements in this work relate to spectral recognition alone. In practice, additional accumulated information can provide enhanced performance, such as position angle and seasonal phase, telescope pointing location, and so forth.

Dragonfly class multi aperture telescopes performed admirably, with maximum overall classification accuracy of 82.7%. This is well within the performance expected based on earlier, less realistic simulations [19], and shows that filter based photometry using on sky imagery (instead of the simulated one dimensional spectra) should unlock the final step in a single instrument with combined detection, localization, and recognition capabilities. For the challenging dataset in which multiple generations of satellite designs are represented, the Dragonfly performance drops roughly 13% to 69.6%. Given that the performance drop reflects confusion within satellite families (Figure 4), the lower performance is consistent with expectation.

Initial tests with the Toucan design and challenging dataset show poor performance of only 24.7%, unacceptable for a

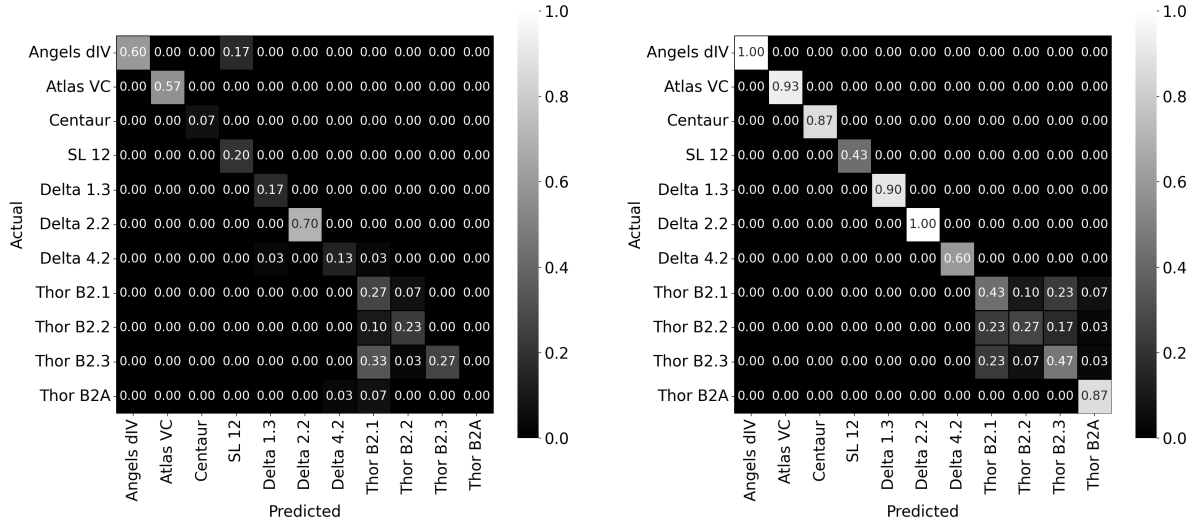

Fig. 6: Effects of defocusing on multispectral instruments. Subset confusion matrices isolate a narrow selection of classes for visibility, and thus rows and columns may not sum to unity. Rocket bodies in the focused Toucan experiment (left) and the defocused experiment (right). As is apparent in Table 1, the performance of the Toucan class instrument is significantly improved when defocusing, which spreads target PSFs over a larger number of pixels and promotes stability in Bayer spectral sampling.

SpectraNet technology. Despite a random performance on this dataset of 0.9%, such low classification accuracies do not provide information to stakeholders which lowers the risk of action. On the single satellite class dataset, performance jumps to only 33.4%. As discussed in §4.1, we attribute these low performance numbers to the confounding variable of PSF position relative to Bayer filter alignment (see §4.2).

To test this, we generate a second dataset in which we heavily defocus the Toucan system with the expectation that a more diffuse PSF will decrease the PSF placement artifacts and increase the number of Bayer pattern pixel matrices covered per observation. This technique is fairly common in astronomy, but is generally applied to maximize flux and thus decrease read noise in exoplanet transit work [17]. We find that this increases the performance of the Toucan system by an impressive 30-40% and suggest that this presents opportunities for the Toucan system to reach performance similar to that of the Dragonfly telescope. In this work we still see a discrepancy favoring the Dragonfly design, with the Toucan multispectral performance on the challenging dataset of 57.1% (-12.5%) and for the simpler dataset of 74.3% (-8.4%). Despite this, the performance numbers are promising and suggest that both multispectral and multiaperture systems should be prototyped and tested on sky.

In these experiments, the best performing model backbone was the HybridSN model, which combines 3D convolutions for spectral features and 2D convolutions for spatial features. The 3D convolutions provide inductive biases for spectral data and offer a promising new line of research to maximize performance on spectrally aware space imagery.

6. CONCLUSIONS

In this work we demonstrate that "single shot" space domain awareness, in which a single telescope can detect, localize, and recognize spatially unresolved satellites is possible with technologies available today. Unlocking this ability allows mainstream SDA applications of spectral recognition which is otherwise impossible due to impractical single mission sensors. The ideal instrument for this type of work follows a Dragonfly class design; a single robotic telescope mount with multiple apertures. A simpler design in which a single aperture feeds a multispectral camera also offers promise and the roughly ten percent drop in recognition accuracy may be overcome with proper instrument configuration and novel network backbones.

[1] CMS series | Multispectral Cameras.
[2] XIMEA - Hyperspectral Snapshot USB3 camera 15 bands 600-860nm.
[3] XIMEA - Hyperspectral Snapshot USB3 camera 16 bands 460-600nm.
[4] XIMEA - Hyperspectral Snapshot USB3 camera 24 bands 665-960nm.
[5] Roberto G. Abraham and Pieter G. van Dokkum. Ultra-Low Surface Brightness Imaging with the Dragonfly Telephoto Array. *Publications of the Astronomical Society of the Pacific*, 126:55, January 2014. ADS Bibcode: 2014PASP..126...55A.
[6] Gregory Badura, Elena Plis, and Christopher Valenta. *Extending Laboratory BRDF Measurements towards Radiometric Modeling of Resident Space Object Spectral Signature Mixing*. September 2021.
[7] Bryce E. Bayer. Color imaging array, July 1976.
[8] Schelte J. Bus and Richard P. Binzel. Phase II of the Small Main-Belt Asteroid Spectroscopic Survey. A Feature-Based Taxonomy. *Icarus*, 158:146–177, July 2002. ADS Bibcode: 2002Icar..158..146B.
[9] Alexander Cabello and Justin Fletcher. SatSim: a synthetic data generation engine for electro-optical imagery of resident space objects. In Khanh D. Pham and Genshe Chen, editors, *Sensors and Systems for Space Applications XV*, page 6, Orlando, United States, June 2022. SPIE.
[10] Heather Cowardin, Jacqu7eline Reyes, Elena Pils, Ryan Hoffman, Gregory Badura, Jainisha Shah, Sydney Collman, Miles Bengtson, Daniel Engelhart, and Timothy Scott. Spectral Characterization of Modern Spacecraft Materials. In *AMOS*, 2022.
[11] Ben Davies, Rolf-Peter Kudritzki, and Donald F. Figer. The potential of red supergiants as extragalactic abundance probes at low spectral resolution. *Monthly Notices of the Royal Astronomical Society*, 407:1203–1211, September 2010. ADS Bibcode: 2010MNRAS.407.1203D.
[12] Francesca E. DeMeo, Richard P. Binzel, Stephen M. Slivan, and Schelte J. Bus. An extension of the Bus asteroid taxonomy into the near-infrared. *Icarus*, 202:160–180, July 2009. ADS Bibcode: 2009Icar..202..160D.
[13] V. Felt and J. Fletcher. Seeing Stars: Learned Star Localization for Narrow-Field Astrometry. 2023.
[14] Justin Fletcher, Ian McQuaid, and Peter Thomas. Feature-Based Satellite Detection using Convolutional Neural Networks. In *AMOS*, page 11, 2019.
[15] J. Z. Gazak and J. Fletcher. StarCSP: Crowd Counting Stars in Rate Track Imagery. 2023.
[16] J. Zachary Gazak, Ben Davies, Rolf Kudritzki, Maria Bergemann, and Bertrand Plez. QUANTITATIVE SPECTROSCOPIC *J* -BAND STUDY OF RED SUPERGIANTS IN PERSEUS OB-1. *The Astrophysical Journal*, 788(1):58, May 2014.
[17] J. Zachary Gazak, John A. Johnson, John Tonry, Diana Dragomir, Jason Eastman, Andrew W. Mann, and Eric Agol. Transit Analysis Package: An IDL Graphical User Interface for Exoplanet Transit Photometry. *Advances in Astronomy*, 2012:e697967, June 2012. Publisher: Hindawi.
[18] J Zachary Gazak, Ian McQuaid, Ryan Swindle, Matthew Phelps, and Justin Fletcher. SpectraNet: Learned Recognition of Artificial Satellites from High Contrast Spectroscopic Imagery. In *WACV*, page 9, 2022.
[19] Jonathan Z. Gazak, Matthew Phelps, Ryan Swindle, Cody Shaw, Zach Funke, and Justin R. Fletcher. The potential for satellite recognition from ground-based filter photometry. In David W. Messinger and Miguel Velez-Reyes, editors, *Algorithms, Technologies, and Applications for Multispectral and Hyperspectral Imaging XXIX*, page 12, Orlando, United States, June 2023. SPIE.
[20] Jonathan Z. Gazak, Matthew Phelps, Ryan Swindle, Andrew Vanden Berg, and Justin R. Fletcher. Model validity dynamics in learned spectroscopic recognition of satellites. In David W. Messinger and Miguel Velez-Reyes, editors, *Algorithms, Technologies, and Applications for Multispectral and Hyperspectral Imaging XXIX*, page 11, Orlando, United States, June 2023. SPIE.
[21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. December 2015.
[22] Jarred Jordan, Daniel Posada, David Zuehlke, Angelica Radulovic, Aryslan Malik, and Troy Henderson. Satellite Detection in Unresolved Space Imagery for Space Domain Awareness Using Neural Networks, July 2022. arXiv:2207.11412 [cs].
[23] Harrison Krantz, Eric C Pearce, Louis Avner, and Kris Rockowitz. Chimera: a high-speed three-color photometer for satellite characterization. 2018.
[24] R. P. Kudritzki, N. Castro, M. A. Urbaneja, I. T. Ho, F. Bresolin, W. Gieren, G. Pietrzyński, and N. Przybilla. A Spectroscopic Study of Blue Supergiant Stars in the Sculptor Galaxy NGC 55: Chemical Evolution and Distance. *The Astrophysical Journal*, 829:70, October 2016. ADS Bibcode: 2016ApJ...829...70K.
[25] Dustin Lang, David W. Hogg, Keir Mierle, Michael Blanton, and Sam Roweis. Astrometry.net: Blind Astromet-

ric Calibration of Arbitrary Astronomical Images. *The Astronomical Journal*, 139:1782–1800, May 2010. ADS Bibcode: 2010AJ....139.1782L.

[26] Louis Lischwe, TPG Wijnen, and ir. M. Rodenhuis. Spectropolarimeter for Satellite Identification. page 12, 2022.

[27] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. page 4, 2019.

[28] Tamara E Payne, Philip Castro, Luke Weisenbach, Ethan VanTilburg, Matthew Wilson, Trenton Godar, Veronica Wiley, James Frith, and Scott P Milster. Peacock: A Persistent Wide-Field-Of-View Simultaneous Multispectral System Based on COTS Hardware. page 17, 2022.

[29] Tamara E Payne, Philip J Castro, Joseph W Moody, Elizabeth A Beecher, Matthew D Fisher, and Roberto I Acosta. A Discrimination Analysis of Sloan and Johnson Photometric Systems for Non-Resolved Object Characterization. 2016.

[30] Swalpa Kumar Roy, Gopal Krishna, Shiv Ram Dubey, and Bidyut B. Chaudhuri. HybridSN: Exploring 3-D–2-D CNN Feature Hierarchy for Hyperspectral Image Classification. *IEEE Geoscience and Remote Sensing Letters*, 17(2):277–281, February 2020. Conference Name: IEEE Geoscience and Remote Sensing Letters.

[31] Ryan Swindle, Zach Gazak, Matthew Phelps, Ian McQuaid, and Justin Fletcher. Learned satellite identification using low-resolution spectropolarimetry. In David B. Chenault and Meredith K. Kupinski, editors, *Polarization: Measurement, Analysis, and Remote Sensing XV*, page 17, Orlando, United States, June 2022. SPIE.

[32] Paul Sydney, John Africano, Amy Fredericks, Kris Hamada, and Vicki Soo Hoo. Raven automated small telescope systems. In *SPIE Imaging Technology and Telescopes*, October 2000.

[33] Douglas Woodward, Celeste Manughian-Peter, Timothy Smith, and Elizabeth Davison. Pixelwise Image Segmentation with Convolutional Neural Networks for Detection of Resident Space Objects. 2021.

[34] Matei Zaharia, Andrew Chen, Aaron Davidson, Ali Ghodsi, Sue Ann Hong, Andy Konwinski, Siddharth Murching, Tomas Nykodym, Paul Ogilvie, Mani Parkhe, Fen Xie, and Corey Zumar. Accelerating the Machine Learning Lifecycle with MLflow. page 7, 2021.