Improving Spectral-Based Estimation of Space Object Orientation

Matthew Phelps, Thomas Swindle, J. Zachary Gazak, Justin Fletcher
Space Systems Command
United States Space Force
Andrew Vandenberg
Air Force Research Laboratory

ABSTRACT

Accurately measuring the orientation of a given satellite can provide vital information in identifying the spacecraft's operational status and predicting its propagation into a future state. Measuring such a quantity using ground-based sensors has, however, proven to be an immense challenge for the majority of space objects of interest, especially those in geosynchronous orbit (GEO) with large orbital radii. In the GEO regime, conventional resolved imaging techniques are ill-equipped to resolve the spatial distribution of material about a satellite due to limited spatial resolution. In this work, we continue exploring the application of ground-based spectrometry to estimate the rotational state of a space object. Building on previous works of spectroscopy applied to tasks like positive-id and pose estimation, we train deep convolution-based neural networks on simulated spectroscopic renders to estimate the rotational state. By increasing the spectral range and extending the capacity of the pose-prediction head, we obtain an appreciable reduction in rotational error. We present current leading performance results on spectroscopic pose estimation and discuss inherent limitations.

1. INTRODUCTION

The primary objective of Space Domain Awareness (SDA) is to obtain a complete specification of the current physical state of all space objects orbiting Earth [1] and predict the evolution of probable future states. The space volumes of typical interest in SDA span from low Earth orbit (LEO), to medium Earth orbit (MEO), to geosynchronous Earth orbit (GEO), to cislunar, and X-GEO beyond. To begin prescribing the state of an object, we first specify the motion of the center of mass. This comprises seven degrees of freedom, which may be represented as six orbital elements (e.g. Kepler elements or two-line-element set) or the state vector $(\mathbf{r}, \dot{\mathbf{r}})$, in addition to the time epoch t. To further specify the physical state, one must also describe the motion about the object's center of mass. This may be achieved by measuring the instantaneous rotational state and angular velocity, which is the objective of this work. Finally, to complete the kinematic specification (sans mass), one must also account for the complex articulations of man-made objects, which necessarily adds additional degrees of freedom describing the kinematic chain of joints.

With the focus of this work being the estimation of the rotational state of a space object, we note that the quantity of interest is the rotational matrix $\mathbf{R}$. As a member of the special orthogonal group in three dimensions, SO(3), the matrix $\mathbf{R}$ must obey

$$\mathbf{R}\mathbf{R}^{\mathbf{T}} = \mathbb{I}, \qquad \det \mathbf{R} = 1. \tag{1}$$

While the 3x3 matrix $\mathbf{R}$ is the canonical description of 3D rotation, there are a number of alternative representations that may be used, some more practical than others. These include quaternions, axis-angle, Euler angles, and the Stieffel 6D vector. [12], [18], and [8] provide an overview these representations, with emphasis on deep learning applications and continuity properties.

In order to measure the rotational state using ground-based sensors, the use of optical reflecting telescopes to collect observations of sufficient spatial resolution is feasible if the sensor-object distance is adequately small. For the

[1] As humans continue their expansion into the cosmos, the set of relevant orbits is naturally extended to include regions about the moon, Mars, sun, etc.

large aperture telescopes used in SDA tasks, in practice this means the object must primary reside in LEO. Here, with the distribution of an object's surface area about its center being acceptably resolved, one can leverage all the advancements in modern computer vision – methods based on natural imagery – to deduce the object's orientation.

What about objects that lie beyond LEO? To begin estimating rotation for objects at larger orbital radii, it appears necessary to begin consideration of non-imaging sensor concepts. One sensor modality that is particularly suited for large distance measurements is that of spectroscopy. The defining virtue of spectroscopy is that, while the flux decays according to $\frac{1}{r^2}$, the relative spectral features themselves are distance invariant. As an illustrative example, Fig. 1 exhibits the 1D reflection spectra for a single space object residing at LEO (top) and GEO (bottom) orbits. In Fig. 1 (right) we can also observe the counterpart imaging renders, with the object at GEO comprising only a handful of pixels.

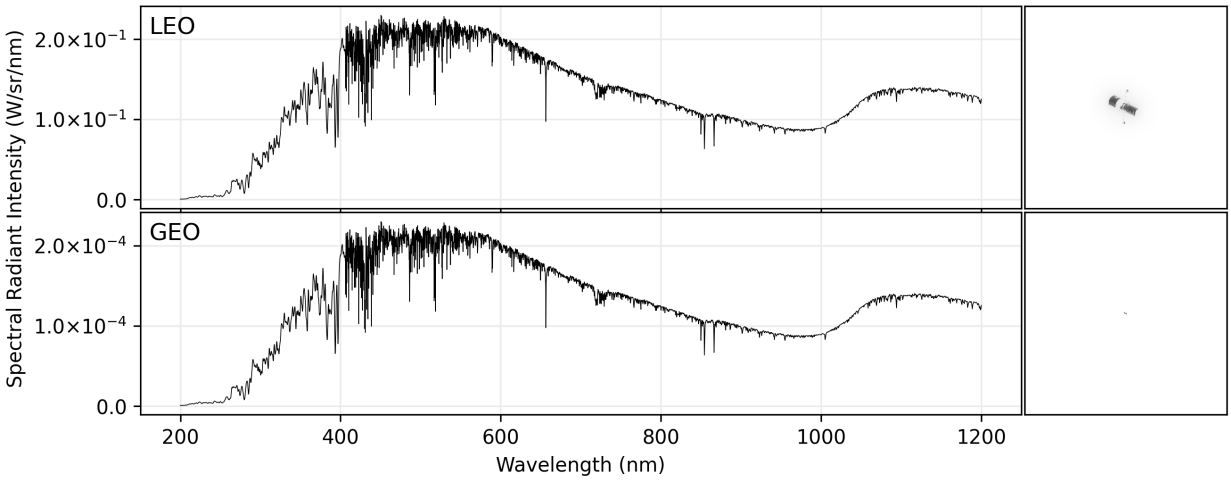

Fig. 1: Simulated reflection spectrum associated with a spacecraft at LEO (top) and GEO (bottom) orbit. (Right) associated optical image renders, inverted in color. As the spacecraft is repositioned from LEO to GEO, we observe a critical loss of information in the associated optical images (right) whereas the spectral information is perfectly preserved (left).

The connection between spectra and rotation lies within the bidirectional reflectance distribution function (BRDF), which captures the information related to an object's atomic material properties, surface normals, and reflection geometry associated with the sun-object-sensor system. As an object rotates, the sensor will sample different values of the BRDF's associated with each material surface. These values are both angle and wavelength dependent. In aggregate, spectra themselves encode an incredibly dense amount of information regarding the interaction of the atoms comprising a material with electromagnetic radiation.

We note that the map from spectra to orientation is highly nonlinear and non-isomorphic. In order to make statistical estimations of **R** using spectra as input, we propose the use of deep neural networks to iteratively learn the underlying mapping through signals received from supervised training. In the following we describe our methods in constructing a network specific to this task and detail an experiment measuring its efficacy.

2. RELATED WORKS

The basis of work exploiting spectral features to learn SDA relevant tasks started with the introduction of SpectraNet [3]. Using real on-sky data captured with a long-slit spectrograph on the 3.6m AEOS telescope, SpectraNet performs positive-id of satellites using a small number of samples per class and demonstrates compelling efficacy with accuracies exceeding 70%. This network is convolution-based, employing a wide-kernel variant of ResNet [7]. This body of work continues with [5] which uses spectra to discrimate beetween closely separate objects (CSO's) and [4] which explores the application of incremental learning of positive-id in an environment with constrained collection cadences. Increasing the density of spectral information, [14] investigates the use of simulated spectral polarimetry

to perform positive-id, yielding improved accuracy relative to [3]. [15] investigates the use of slitless spectrometry, which simultaneously admits multi-object identification and positive-id of satellites within the field of view.

Following the architecture of SpectraNet, [12] continues the application of spectra towards orientation estimation, just as we do here. Here, wide-kernel ResNets are modified to perform direct regression of the rotational state. Accounting for the continuity properties of rotation representations and their role in optimization methods, [12] employs the use of the 6D Stieffel representation. This work serves as the baseline of comparison and is contrasted to our results in Table 1.

Computer vision techniques suited for pose estimation fall into two categories: those that are supervised by an intermediate representation such as keypoints, and those that perform direct pose estimation. These intermediate representations exploit the use of ancillary ground truth labels associated with the geometric structure of the resolved image. As such, this body of work cannot be applied to non-resolved sensor concepts. With respect to direct pose estimation, there are a number of convolution-based works, notably [10] and [13].

Finally, we note the existing work of using light curves to perform attitude estimation [1], [2]. It worth noting that the technique of photometry can viewed as summing the spectral irradiance over all wavelengths into a single scalar and thus integrates out all the wavelength dependent information encoded in spectra.

3. METHODS

3.1 Dataset

Similar to [12], we construct a dataset comprised of 1D spectral renders of a space object at various orientations which serves as the ground truth labels in the supervised training setup. The SO(3) space defining the rotational states is sampled in an equi-distant manner and comprises a total of 5e4 observations. Using an advanced radiometric simulator, we model the optical system after the 10m Keck optical telescope [16] located at the summit of Mauna Kea, Hawaii. Relative to [12], we expand the spectral range further into the infrared (IR) to cover 320 – 1010 nm with an average spectral resolution $\Delta\lambda = 0.28$ nm. The space object of interest is based on a 3D material surface model (with associated BRDF's) of the Galaxy-15 telecommunications satellite [6] positioned in geostationary orbit at 35000 km. In order to measure the upper bound of the performance of a spectral-based pose estimator, the geometry of the system was designed to represent an absolute minimal complexity baseline configuration; namely, the solar phase angle is fixed at 45°, amounting to a sun-object-sensor system behaving as a fixed rigid body.

3.2 Architecture

The architecture of the neural network (Fig 2) comprises two modules: a convolutional neural network (CNN) to efficiently process the raw 1D spectra and form meaningful representations, and a pose prediction head to estimate the direct rotational state. Building on the work of [12], we employ the use of wider initial kernels and strides in order to target spectral features that span a larger band of wavelengths. Unlike conventional resolved imagery, we note that the spatial axis actually represents the variation across wavelength, and thus the convolutional inductive bias of spatial-equivariance may be an inherent limitation in spectral applications. Nonetheless, due to the parameter efficiency afforded by convolutional kernels and the wide success evident in time-series applications, we employ CNN's to serve as the backbone of the architecture.

Taking in the 1D reflectance spectrum input (a tensor of shape (1, 3584)), the CNN learns a vector representing the spectral embedding for a single sample (e.g. encoded vector, latent vector). The spetral embedding (shape (1, 512)) is unified with a tensor comprised of a grid of SO(3) vector embeddings through an efficient broadcast and sum operation [10]. The spectral-SO(3) tensor then passes through a number of fully connected layers (MLP) to finally yield a vector of unnormalized logit scores. By injecting the discritized grid of SO(3) embeddings, we are able to model the rotation estimator as a conventional classification task. Hence, we utilize the cross entropy loss $\mathscr{L}$ which, in our implementation, effectively normalizes the logits with a softmax layer and computes the negative log likelihood.

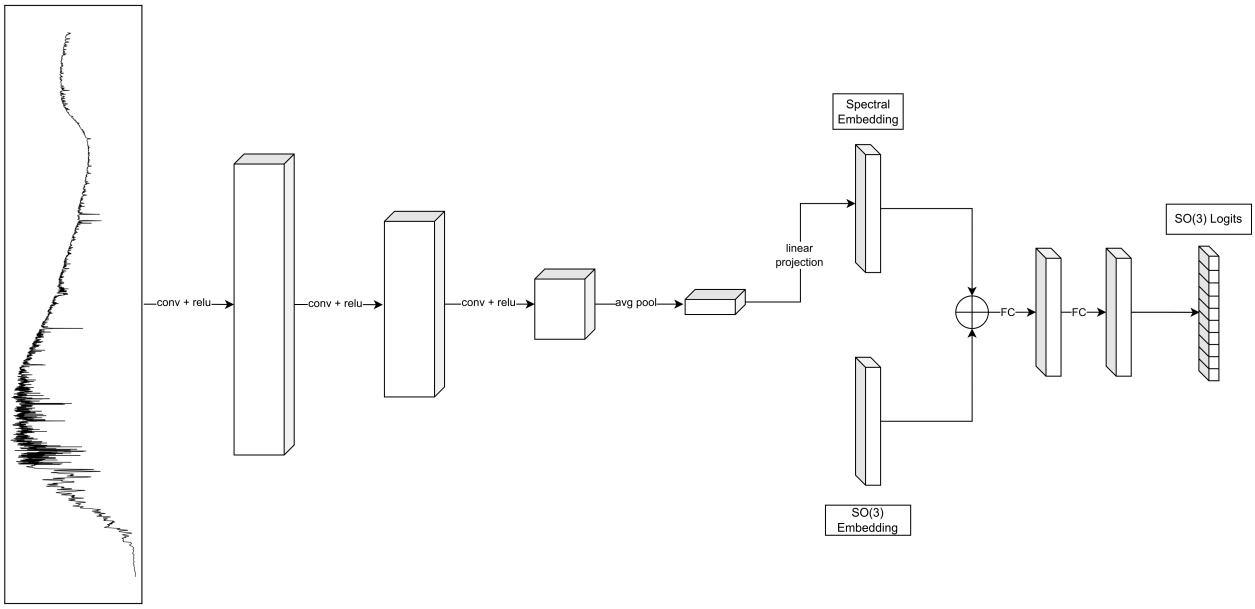

Fig. 2: Convolution + SO3MLP architecture.

3.3 Training

To train the rotation estimator, each run performs total of 2.2e6 single-sample gradient evaluations, utilizing a series of Nvidia A100 (80GB) GPU's. In our compute environment, this amounted to 5e4 train steps with a batch size of 44. A linear warmup (10%) cosine decay learning rate schedule is used, with a peak learning rate of 1e-4. Optimization is performed using AdamW [9], which also performs the duties of weight decay. Models are trained using the PyTorch framework [11] and analyzed with MLflow [17].

4. EXPERIMENTS

Given the 5e4 observations of the satellite at various rotational orientations, we perform an 85/15 train/val split and evaluate the performance on both the previous baseline of [12] and the Conv-SO(3)MLP architecture of Fig. 2. The performance is measured using SO(3) geodesic error

$$\Phi_1(\mathbf{R}_1, \mathbf{R}_2) = \cos^{-1}\left[\tfrac{1}{2}(Tr(\mathbf{R}_1\mathbf{R}_2^T) - 1)\right]. \tag{2}$$

which serves as a proper distance metric in SO(3) space, measuring the geodesic distance between the predicted and ground truth rotation. The range of the error metric spans from $[0°, 180°]$ and the performance of a rotation estimator choosing an orientation at random is approximately 126°.

Model	Spectral Range	$\Delta\lambda$	Conv. Kernels	Conv. Strides	$\downarrow$ SO(3) Geodesic Error
ResNet-18W + SO3MLP	200 - 1200 nm	0.28 nm	[28, [3]×18]	[12, [2]×5]	102.38
ResNet-18W [12]	320 - 1010 nm	0.30 nm	[28, [3]×18]	[12, [2]×5]	95.56
Conv + SO3MLP	200 - 1200 nm	0.28 nm	[28, [3]×5]	[12, [2]×5]	78.07
Conv + SO3MLP	200 - 1200 nm	0.28 nm	[28, [3]×7]	[12, [2]×7]	**65.64**

Table 1: Performance results across varying convolution-based architectures, measured by SO(3) geodesic error.

We first note that while the baseline of [12] has an appreciable amount of error remaining, its distance from the random baseline indicates a valuable learning of the underlying mapping between spectra and rotation. As we directly follow-on the work of [12] and incorporate the SO(3)MLP pose module, we observe an sizable decrease in performance

relative to the baseline. Continued experiment revealed that even with heavy regularization or smaller ResNet modules (e.g. ResNet-10), the use of vanilla CNN's consistently outperformed the ResNet counterparts when used in conjuction with the SO(3)MLP estimation head.

Fusing vanilla convolutional layers with the SO3MLP head, we find a significant reduction in geodesic error, improving upon the previous baseline by 32%. Additionally, we find that variation of the kernels and strides themselves have a significant impact on the performance. We hypothesize that spatially smaller representations formed through the CNN emphasize more relevance of features along the channel dimension. As the spatial dimension is always averaged pooled to size unity, representations with larger spatial extent may suffer a degradation in information in the average reduction process. In effect, this would suggest that maximal down-sampling via strides followed by appropriate kernel sizing in the CNN would yield higher performances, as is reflected in the study of Table 1.

5. CONCLUSION

In order to improve the estimation of the physical state of a space object orbiting Earth beyond the motion of an object's center of mass, one must establish methods to effectively measure the rotational motion *about* the center of the mass. By leveraging spectroscopy as a sensing method for objects at large orbital radii, this work continues from [12] to investigate the viability of estimating orientation from optical spectra. To improve upon prior baselines, we a) increase the spectral range and resolution and b) introduce a new spectra-SO(3) architecture. With these modifications, we establish a significant improvement in performance, reducing the SO(3) geodesic error by a factor of 0.32. As this line of work is carried into the future, we believe further steps must be taken to reduce the delta between simulation and application in order to reach viability in a physical setting. Leading sources of simulation gap are variation of sun-sensor-geometry and generalization to realistic focal-plane based 2D spectrographs.

6. REFERENCES

[1] G. Badura, C. Valenta, and B. Gunter. Convolutional Neural Networks for Inference of Space Object Attitude Status. *AMOS Conference Proceedings*, 2020.

[2] Gregory Badura, Christopher Valenta, Brian Gunter, and Brian Shoffeitt. Multi-scale convolutional neural networks for inference of space object attitude status from detrended geostationary light curves. 02 2021.

[3] 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.

[4] 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.

[5] Z. Gazak, J. Fletcher, T. Swindle, and I. McQuaid. Exploiting Spatial Information in Raw Spectroscopic Imagery using Convolutional Neural Networks. *AMOS Conference Proceedings*, 2020.

[6] Doyle Hall. Amos galaxy 15 satellite observations and analysis. In *Advanced Maui Optical and Space Surveillance Technologies Conference*, page E12, 2011.

[7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

[8] Du Q Huynh. Metrics for 3d rotations: Comparison and analysis. *Journal of Mathematical Imaging and Vision*, 35(2):155–164, 2009.

[9] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.

[10] Kieran Murphy, Carlos Esteves, Varun Jampani, Srikumar Ramalingam, and Ameesh Makadia. Implicit-pdf: Non-parametric representation of probability distributions on the rotation manifold. *arXiv preprint arXiv:2106.05965*, 2021.

[11] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.

[12] 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.

[13] Martin Sundermeyer, Zoltan-Csaba Marton, Maximilian Durner, Manuel Brucker, and Rudolph Triebel. Implicit 3d orientation learning for 6d object detection from rgb images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 699–715, 2018.

[14] 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.

[15] Peter Thomas, J Zachary Gazak, Ryan Swindle, Zachary Funke, and Justin Fletcher. Slitless field spectroscopy for learned simultaneous object detection and recognition. In *Algorithms, Technologies, and Applications for Multispectral and Hyperspectral Imaging XXVIII*, volume 12094, pages 109–119. SPIE, 2022.

[16] Marcos A van Dam, David Le Mignant, and Bruce A Macintosh. Performance of the keck observatory adaptive-optics system. *Applied Optics*, 43(29):5458–5467, 2004.

[17] Matei Zaharia, Andrew Chen, Aaron Davidson, Ali Ghodsi, Sue Ann Hong, Andy Konwinski, Siddharth Murching, Tomas Nykodym, Paul Ogilvie, Mani Parkhe, et al. Accelerating the machine learning lifecycle with mlflow. *IEEE Data Eng. Bull.*, 41(4):39–45, 2018.

[18] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5745–5753. IEEE, 2019.