

# M2S2: MULTISPECTRAL DATASET FOR MATERIAL SEGMENTATION OF SATELLITES

*Kimmy De Alba, Enrique De Alba, J. Zachary Gazak, Justin Fletcher*

Space Systems Command (A&AS), EO Solutions

## ABSTRACT

Material segmentation of satellites using multispectral imaging can support ground-based sensing and is often less sensitive to range and angular resolution. Many existing methods are validated primarily in simulation, leaving sim-to-real transfer under-evaluated. We introduce M2S2, comprising 12,960 synthetic and 3,024 hardware-in-the-loop (HIL) images across 10 spectral bands (400–900 nm), three satellite geometries, and up to eleven material classes. The dataset systematically varies elevation angles, material complexity, and lighting conditions, with synthetic wavelength-to-RGB approximations and true multispectral HIL captures. As a baseline, adapting Segment Anything (SAM) to multispectral inputs yields statistically significant gains over RGB on synthetic data, with advantages increasing with material count (up to +4.64% macro recall), where 75% of non-two-class settings are significant at  $\alpha = 0.05$ . HIL predictions exhibit uniformly high temporal consistency ( $mTC = 0.9830$ ) with modest lighting effects (weak diffuse exceeding directional by 0.0024). These results suggest M2S2 is useful for characterizing sim-to-real challenges and for studying domain adaptation in satellite material segmentation. The dataset is available at <https://huggingface.co/datasets/e-dealba/M2S2>.

**Index Terms**— Multispectral, Material Segmentation, Segment Anything Model (SAM), Hardware-in-the-loop (HIL).

## 1. INTRODUCTION

Autonomous material segmentation of resident space objects (RSOs) supports ground-based sensing tasks such as functionality assessment, status monitoring, and anomaly screening [1, 2, 3]. While on-orbit material segmentation has been demonstrated [4], ground-based sensors offer advantages including monitoring capabilities, reduced mission complexity, and surveillance of large orbital populations [5, 6, 7]. Traditional methods rely on optical imagery, but their effectiveness is reduced at geostationary orbit (GEO) and beyond due to limited spatial resolution. To address these constraints, spectroscopic techniques have been adopted, offering characterization based on material reflectance spectra that is less sensitive to range and angular resolution [8].

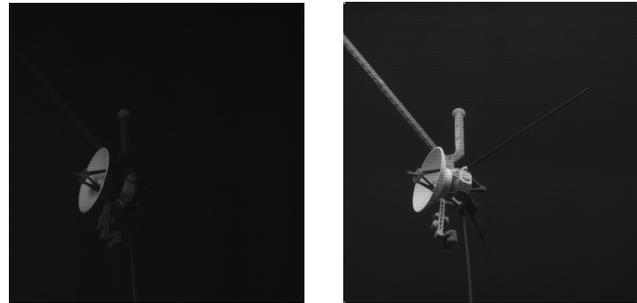
Prior work has presented semantic segmentation methods for space traffic management (STM) using multispectral imaging [9]. However, the sim-to-real transfer problem is under-evaluated on real data, with existing methods only validated on simulated data using advanced simulators due to lack of spectral data [10] and the challenge in obtaining real on-orbit data [11]. A public multispectral dataset can support sim-to-real studies by providing labeled data that are impractical to obtain in operational scenarios. In the past decade, similar datasets like SPEED+ [12] have catalyzed advances in spacecraft pose estimation and navigation.

Following the approach established by SPEED+, we employ a similar methodology to physically recreate and simulate the space environment using a ground-based hardware-in-the-loop (HIL) robotic testbed capable of capturing arbitrary numbers of images of target mockup models. These HIL images serve as a surrogate for spaceborne imagery, enabling ground-based evaluation of segmentation performance across domains that differ from synthetic training data.

In this paper, we propose M2S2, a publicly available benchmark Multispectral dataset for Material Segmentation of Satellites. M2S2 comprises 12,960 synthetic and 3,024 hardware-in-the-loop (HIL) images, captured with a TOUCAN snapshot camera spanning 10 spectral bands (400–900 nm). It features three satellite geometries, up to eleven material classes. Sample images are shown in Figure 1 of synthetic RGB and multispectral captures alongside HIL images.

For the synthetic domain, we utilized the open-source 3D graphics software Blender to render analogous scenarios: elevation angles, material assignments, and satellite geometries. HIL images were captured using the TOUCAN Snapshot Multispectral Camera under three different illumination conditions: weak diffuse, strong diffuse, and directional lighting. M2S2’s images feature three distinct satellite geometries—Voyager, ISS, and Hubble—with each satellite incorporating up to 11 materials across three different material combinations.

We provide a baseline performance study focused on material segmentation, while noting the dataset’s broader potential for applications including 3D reconstruction [13, 14] and pose estimation [15].



**Fig. 1: Example Raw Images of 10-Material Voyager Satellite from M2S2.** From left to right: synthetic, RGB; synthetic, multispectral; HIL, directional; HIL, strong diffuse.

## 2. M2S2 DATASET

### 2.1. Setup

The dataset was collected using a TOUCAN Snapshot Multispectral Camera (SILIOS Technologies), which captures images across 10 spectral bands spanning 400–900 nm with an average bandwidth of 30–50 nm. The camera integrates a custom Bayer-like mosaic filter on a 4.2 MP CMOS sensor ( $2048 \times 2048$  raw resolution,  $512 \times 512$  spectral resolution,  $5.5 \mu\text{m}$  pixel pitch) with a maximum acquisition rate of 65 Hz. Exposure times are adjustable from 100  $\mu\text{s}$  to 5 s with 10-bit digitization. For operation, we employed the COLOR SHADES LAB SDK, which enables display of raw and processed images, spectral normalization, crosstalk correction, and hypercube extraction. During data collection, the TOUCAN was mounted on a SmallRig mini tripod with a fixed clamp for stability and fitted with a vintage 50 mm f/2 M42-mount lens. Diffuse illumination was provided by full-spectrum LED bars (24 W, 5000 K,  $\sim 2200$  lm each). For weak diffuse lighting, three bars were activated ( $\sim 6600$  lm total); for strong diffuse lighting, five bars were used ( $\sim 11000$  lm). Directional illumination was achieved with a LED spotlight (50 W, 5000 K,  $\sim 5000$  lm), oriented at a controlled incident angle to create single-source, high-contrast conditions.

### 2.2. Satellite Mockups

The dataset includes mockups of three satellites: Voyager, International Space Station (ISS), and Hubble. We fabricated mockups with a Bambu A1 Mini 3D printer to create controlled targets with material and color diversity. A range of eleven PLA-based filaments were employed, including PolyLite PLA (Dark Grey, Olive Brown, Lime Green, Red, Azure Blue, Cold White, Wood) and PolyLite PLA Pro (Metallic Blue, Silver, Gold, Metallic Magenta). These variations in surface color and reflectance were selected to emulate distinct material signatures for multispectral segmentation. For each satellite geometry, we offer three material

**Table 1: Material Usage Across M2S2 Satellite Builds (B1/B2/B3).** All materials are PolyLite PLA filaments.

Material	Voyager			ISS			Hubble		
	B1	B2	B3	B1	B2	B3	B1	B2	B3
Pro Silver	✓	✓	✓	✓	✓	✓	✓	✓	✓
Dark Grey		✓	✓		✓	✓		✓	✓
Pro Metallic Blue			✓		✓	✓		✓	✓
Pro Gold		✓	✓		✓	✓			
Cold White		✓	✓		✓	✓			
Pro Metallic Magenta			✓					✓	✓
Red			✓			✓			
Azure Blue			✓			✓			
Olive Brown			✓			✓			
Lime Green			✓			✓			
Wood						✓			
<b>Total Materials</b>	<b>1</b>	<b>4</b>	<b>10</b>	<b>1</b>	<b>5</b>	<b>10</b>	<b>1</b>	<b>4</b>	<b>4</b>

combinations: (i) a simple control wherein the satellite is entirely composed of PolyLite PLA Pro Silver Filament; (ii) a more challenging multi-material version where the satellite has 4–6 different materials; and (iii) the most complex case where satellites can have up to 11 different materials per build. The exact filament combinations are tabulated in Table 1.

### 2.3. Generating Synthetic and HIL Images

We import Bambu Lab meshes (STL/3MF) into Blender and apply a uniform rendering/capture protocol. For each satellite geometry and material, synthetic and HIL images are acquired every  $10^\circ$  about the  $z$ -axis at positive, flat, and negative tilts. The HIL protocol is repeated under all lighting conditions to mirror the synthetic setup, yielding paired masks and labels across domains.

Because Blender renders in RGB, we approximate each multispectral band by mapping its center wavelength to an RGB color. For each of the 10 TOUCAN sensor bands (rang-

ing from 431 nm to 877 nm), we apply a wavelength-to-RGB mapping function with gamma correction ( $\gamma = 0.80$ ) to compute an approximate RGB representation of each spectral band. Scene lighting is then adjusted to this RGB color to simulate illumination at the target wavelength. This yields band-specific RGB renderings that capture approximate spectral variations across materials. While this approach does not model sensor spectral responses, BRDF effects, or atmospheric transmission, it provides a controlled approximation for studying cross-domain tendencies, not absolute multi-spectral performance.

The raw HIL frames include a positioning stand; which shares the same material as our single-material models. We also acquire a single lens-capped dark frame  $D_b$  and subtract it from the corresponding raw image  $R_b$ :

$$I_b = R_b - D_b,$$

followed by clipping negative values to zero. Residual artifacts are spatially disjoint from satellite regions and had negligible effect in our segmentation runs.

## 2.4. Material Segmentation

### 2.4.1. Zero-Shot Segment Anything

We run Segment Anything (SAM) in zero-shot “auto” mode on each image to produce instance masks without prompts, using a ViT-H checkpoint with a high IoU/stability threshold. These masks are then projected to a single semantic label map via a compositing pass that paints unlabeled pixels mask-by-mask (small-to-large), requiring minimal effective coverage and preferring non-background labels on near ties. For RGB, this yields one semantic prediction per view. For multispectral, we first use PCA to pick the three bands with the largest loadings on the leading components (i.e., bands 5, 9, and 10) and then run the same SAM auto-mask to semantic conversion per selected band. The resulting label maps are then fused with a pixel-wise majority vote (non-background favored on ties) to form the final multispectral prediction. We use the top three bands for synthetic and HIL images.

### 2.4.2. Performance Metrics

For synthetic images, we evaluate macro-averaged precision, macro-averaged recall, and Dice [16]. Macro-averaged precision and recall provide class-agnostic performance assessment by computing precision and recall for each material class independently and averaging the results. This approach ensures that performance on minority material classes receives equal weight to majority classes, which is important given the varying material distributions across satellite geometries. The Dice coefficient (also known as F1-score) represents the harmonic mean of precision and recall, providing a single metric that balances both measures. Together, these

metrics provide evaluation coverage of both pixel-level accuracy and class-level performance consistency. Significance was assessed via paired t-tests at  $\alpha = 0.05$  and  $\alpha = 0.01$  with an assumed within-image correlation  $\rho = 0.9$ , acknowledging shared images but avoiding the unrealistic  $\rho = 1.0$ .

Because exhaustive ground truth is unavailable for HIL images, we report a ground-truth-free *temporal/augmentation consistency* score (mTC), which has been reported to correlate with supervised mIoU and is used as a stability proxy [17]. We define mTC as the mean IoU between a model’s base prediction and its predictions under small input perturbations, all compared in the same reference frame. Let  $x \in \mathbb{R}^{H \times W \times C}$  and  $f$  produce  $\hat{y} = f(x) \in \{0, \dots, L\}^{H \times W}$ . For  $k = 1, \dots, K$ , apply a perturbation  $\tau_k$  to  $x$ , predict  $\hat{y}'_k = f(\tau_k(x))$ , and inverse-warp  $\tilde{y}_k = \tau_k^{-1}(\hat{y}'_k)$ . We use micro-IoU (optionally omitting an ignore label  $g$ ):

$$\text{mTC} = \frac{1}{K} \sum_{k=1}^K \text{IoU}(\hat{y}, \tilde{y}_k).$$

Unless noted otherwise, we use the same three bands as synthetic  $\{5, 9, 10\}$ , convert SAM proposals to a per-band semantic map by keeping the top  $k_{\text{per\_band}} = \text{num\_classes}$  instances after IoU de-duplication (0.9), and fuse bands by majority vote in *binary* form (foreground vs. background). We then compute mTC with  $K=8$  *random integer-pixel translations* applied identically to all bands, where  $(\Delta x_k, \Delta y_k)$  are sampled uniformly from  $\{-3, -2, -1, 0, 1, 2, 3\}^2$ ; inverse warping is exact by roll back. A fixed random-number generator (RNG) seed (= 42) is used for reproducibility.

### 2.4.3. Implementation Details

All inferences were conducted on a workstation equipped with four NVIDIA Tesla V100 GPUs (32 GB each, CUDA 11.5), 251 GB of RAM, and an Intel Xeon E5-2698 v4 CPU (20 cores, 40 threads).

## 3. EXPERIMENTS

We provide a baseline performance study by adapting Segment Anything Model (SAM) [18] to multispectral data by applying it per-channel and fusing results via majority voting. The synthetic domain consists of RGB-approximated multispectral renderings, while the HIL domain contains true multispectral imagery from the TOUCAN sensor. The goal of this study is to characterize M2S2 and validate its suitability for evaluating domain gaps in spaceborne machine learning (ML) and vision applications, rather than to identify an effective multispectral segmentation approach.

From Table 2 and Figure 2, synthetic multispectral inputs show higher scores than RGB across satellites, elevations, and class counts, with the most consistent gains in macro recall and corresponding improvements in Dice. The advantage

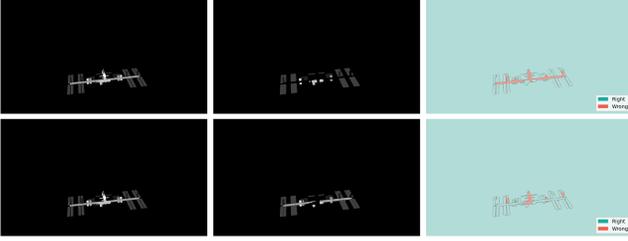


Fig. 2: Sample synthetic material segmentations, 10-material ISS, RGB (top) and multispectral (bottom) Left to right: ground truth, prediction, correctness map.

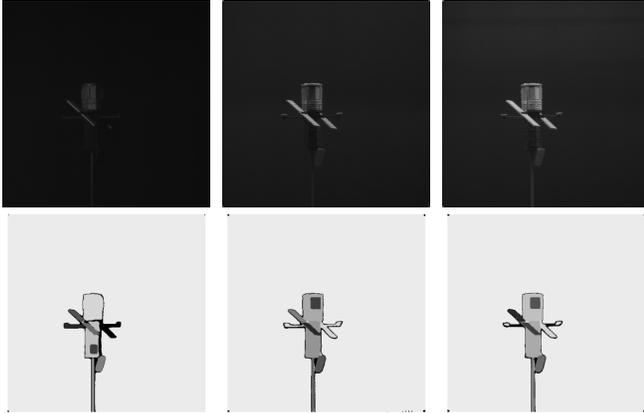


Fig. 3: Sample HIL images (top) and material segmentations (bottom), Hubble B2, E2. Left to right: directional, weak diffuse, strong diffuse.

scales with class complexity: for Voyager, multispectral improves macro recall by +0.0464 in the 11-material case, compared with +0.0195 at 5 materials and -0.0107 at 2 materials. ISS and Hubble show the same pattern, with minimal change in the 2-class setting and clear lifts at higher class counts, notably for ISS at 6 and 11 materials. Across non-two material classes 75% of macro-recall improvements are significant at  $\alpha = 0.05$  and 55% at  $\alpha = 0.01$ . Precision shifts are modest and mixed, consistent with higher coverage of true positives without per-class calibration, while recall and Dice account for most of the significant gains. Averaged over settings, ISS benefits most (Dice +3.63%, Recall +4.53%), followed by Voyager (+1.29%, +1.23%) and Hubble (+0.87%, +1.34%). Variation across elevations within a satellite is small relative to the effects of material class count and satellite geometry, which are the primary drivers of the multispectral advantage.

From Table 2 and Figure 3, multispectral HIL prediction stability is uniformly high: mTC scores cluster around 0.9830 for across satellites, class counts, and elevations, indicating near-ceiling consistency under small perturbations. Lighting has a modest, repeatable effect, with Solo/All (diffuse) typically matching or slightly exceeding Directional by an average 0.0024/0.0023, consistent with smoother illumina-

Table 2: Per-Satellite Material Segmentation Results with SAM Baseline. Results shown for Synthetic RGB and Multispectral (Multi) and HIL Directional, Weak Diffuse, and Strong Diffuse for Voyager, ISS, and Hubble across various elevations and material combinations.

Satellite	Classes	Elevation	Macro Precision		Synthetic Macro Recall		Dice		HIL mTC		
			RGB	Multi	RGB	Multi	RGB	Multi	Dir	Weak	Strong
Voyager	2	E0	0.9497	<b>0.9614</b>	<b>0.9356</b>	0.9224	0.9420	0.9400	0.9904	0.9888	0.9904
		E1	0.9523	0.9544	0.9335	0.9385	0.9411	0.9453	0.9883	0.9864	0.9888
		E2	0.9254	<b>0.9365</b>	<b>0.8868</b>	0.8628	<b>0.9020</b>	0.8923	0.9885	0.9862	0.9883
	5	E0	0.7708	<b>0.7916</b>	0.6966	<b>0.7185</b>	0.7050	<b>0.7255</b>	0.9868	0.9868	0.9869
		E1	0.7892	0.7870	0.7506	0.7648	0.7447	0.7540	0.9877	0.9863	0.9855
		E2	0.7397	<b>0.7777</b>	0.6946	<b>0.7171</b>	0.6988	<b>0.7181</b>	0.9869	0.9853	0.9854
11	E0	0.5718	<b>0.5970</b>	0.6539	<b>0.6815</b>	0.6021	<b>0.6278</b>	0.9891	0.9875	0.9870	
	E1	0.5869	0.5912	0.6750	0.6759	0.6199	0.6201	0.9883	0.9853	0.9855	
	E2	0.5100	0.5203	0.5773	<b>0.6059</b>	0.5336	<b>0.5525</b>	0.9881	0.9883	0.9875	
ISS	2	E0	0.9079	<b>0.9121</b>	0.9613	0.9975	0.9241	<b>0.9505</b>	0.9770	0.9777	0.9788
		E1	0.9168	0.9163	0.9262	0.9396	0.9091	0.9197	0.9728	0.9773	0.9769
		E2	0.9426	0.9391	0.9323	0.9377	0.9344	0.9359	0.9631	0.9725	0.9727
	6	E3	0.9110	0.9090	0.9684	<b>0.9814</b>	0.9339	<b>0.9395</b>	-	-	-
		E0	0.7579	0.7465	0.5663	<b>0.6226</b>	0.5928	<b>0.6219</b>	0.9752	0.9785	0.9812
		E1	0.7651	<b>0.7990</b>	0.6157	<b>0.6540</b>	0.6279	<b>0.6695</b>	0.9687	0.9733	0.9756
	11	E2	0.8051	<b>0.8290</b>	0.6691	<b>0.7229</b>	0.6893	<b>0.7371</b>	0.9728	0.9762	0.9774
		E3	0.7443	0.7628	0.6487	0.6635	0.6450	0.6603	-	-	-
		E0	0.6305	0.6151	0.4800	0.4999	0.5013	0.5084	0.9732	0.9846	0.9818
11	E1	0.5594	<b>0.6136</b>	0.4382	<b>0.4862</b>	0.4578	<b>0.5025</b>	0.9610	0.9793	0.9788	
	E2	0.5814	<b>0.6326</b>	0.5117	<b>0.5791</b>	0.5174	<b>0.5797</b>	0.9676	0.9787	0.9775	
	E3	0.5832	0.5891	0.4797	0.4844	0.4857	0.4921	-	-	-	
Hubble	2	E0	0.9698	0.9683	0.9911	<b>0.9948</b>	0.9802	0.9811	0.9878	0.9871	0.9846
		E1	<b>0.9764</b>	0.9744	0.9921	<b>0.9962</b>	0.9841	<b>0.9850</b>	0.9877	0.9902	0.9861
		E2	<b>0.9744</b>	0.9732	0.9905	<b>0.9939</b>	0.9823	<b>0.9833</b>	0.9854	0.9819	0.9867
	5	E0	0.7397	0.7464	0.7578	0.7721	0.7477	0.7583	0.9867	0.9861	0.9855
		E1	0.7269	0.7388	0.7369	<b>0.7606</b>	0.7318	<b>0.7478</b>	0.9860	0.9868	0.9865
		E2	0.6950	0.7006	0.7070	0.7210	0.7004	0.7100	0.9850	0.9850	0.9862
11	E0	0.7844	0.7927	0.7995	0.8173	0.7912	0.8043	0.9867	0.9890	0.9863	
	E1	0.7322	0.7378	0.7471	<b>0.7590</b>	0.7392	0.7471	0.9817	0.9877	0.9860	
	E2	0.6978	0.6992	0.7189	0.7259	0.7070	0.7113	0.9839	0.9873	0.9852	

\*Paired t-test statistical significance denoted with **Bold** for  $\alpha = 0.05$  and *italic* for  $\alpha = 0.01$ . Note: E3 is an additional elevation in our synthetic dataset.

tion yielding more stable masks. Class complexity introduces small degradations. Variation across elevations (E0–E2) is minor relative to lighting and class count. Overall, these results show the HIL pipeline produces stable segmentations in real captures, with mild sensitivity to directional light and high class cardinality that is largely mitigated by diffuse illumination.

Overall, these results suggest M2S2 is useful as a controlled testbed for evaluating domain gaps in spaceborne ML. The observed multispectral gains on synthetic data and systematic performance trends across lighting, geometry, class count, and material contrast indicate that the dataset exposes realistic challenges for adapting segmentation systems toward spaceborne settings.

## 4. CONCLUSION

We introduce M2S2, a publicly available multispectral dataset for satellite material segmentation, comprising RGB-approximated synthetic multispectral imagery and hardware-in-the-loop (HIL) true multispectral data across three satellite geometries with up to eleven material classes. Our baseline SAM evaluation demonstrates statistically significant material segmentation performance lifts for synthetic multispectral inputs over RGB. It also shows high stability in prediction for HIL images across satellite geometry, material, lighting, and elevation scenarios. This validates M2S2’s utility for characterizing sim-to-real transfer challenges in spaceborne appli-

cations. The dataset provides a standardized benchmark for developing and evaluating multispectral segmentation methods tailored to space-based scenarios. By enabling systematic study of domain adaptation techniques, M2S2 represents an important step toward autonomous material identification capabilities for future space missions that require satellite characterization and monitoring.

## 5. REFERENCES

- [1] John Kielkopf, Frank O. Clark, and Wellesley Pereira, “Remote sensing of satellite activity through optical and infrared temporal differential spectrophotometry informed by analysis of noise,” in *Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS)*, 2022.
- [2] Yuri Matsushita, Ryohei Arakawa, Yasuhiro Yoshimura, and Toshiya Hanada, “Light curve analysis and attitude estimation of space objects focusing on glint,” in *First International Orbital Debris Conference*, 2019.
- [3] Imene R. Goumiri, Amanda L. Muyskens, Benjamin W. Priest, and Robert E. Armstrong, “Light curve forecasting and anomaly detection using scalable, anisotropic, and heteroscedastic gaussian process models,” in *Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS)*, 2023.
- [4] Stephen Hobbs, Leonard Felicetti, George Brydon, Toby Harris, Lounis Chernak, and Umair Soori1, “Towards in-orbit hyperspectral imaging of space debris,” in *2nd NEO and Debris Detection Conference*, 2023, vol. 2.
- [5] Tanner Campbell, Adam Battle, Dan Gray, Om Chabra, Scott Tucker, Vishnu Reddy, and Roberto Furfaro, “Stingray sensor system for persistent survey of the geobelt,” *Sensors*, vol. 24, no. 8, pp. 2596, 2024.
- [6] Mark R. Ackermann, Rex R. Kiziah, Peter C. Zimmer, John T. McGraw, and David D. Cox, “A systematic examination of ground-based and space-based approaches to optical detection and tracking of satellites,” in *31st Space Symposium*, 2015.
- [7] Keith Morris, Chris Rice, and Eugene Little, “Relative cost and performance comparison of geospace situational awareness architectures,” in *Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS)*, 2014.
- [8] J. Zachary Gazak, Ian McQuaid, Ryan Swindle, Matthew Phelps, and Justin Fletcher, “Spec-tranet: Learned recognition of artificial satellites from high contrast spectroscopic imagery,” *CoRR*, vol. abs/2201.03614, 2022.
- [9] Kimmy Chang, J. Zachary Gazak, and Justin Fletcher, “Satellite material segmentation using multispectral imaging for space domain awareness,” in *SPIE Optics + Optoelectronics*, 2025.
- [10] Miguel Velez-Reyes and Jiarui Yi, “Hyperspectral unmixing for remote sensing of unresolved objects,” in *Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS)*, 2020.
- [11] Massimiliano Vasile, Lewis Walker, Andrew Campbell, Simão Marto, Paul Murray, Stephen Marshall, and Vasili Savitski, “Space object identification and classification from hyperspectral material analysis,” *Scientific Reports*, vol. 14, 2024.
- [12] Tae Ha Park, Marcus Märten, Gurvan Lécuyer, Dario Izzo, and Simone D’Amico, “SPEED+: next generation dataset for spacecraft pose estimation across domain gap,” *CoRR*, vol. abs/2110.03101, 2021.
- [13] Ru Li, Jia Liu, Guanghui Liu, Shengping Zhang, Bing Zeng, and Shuaicheng Liu, “Spectralnerf: Physically based spectral rendering with neural radiance field,” *Association for the Advancement of Artificial Intelligence*, 2024.
- [14] Jiabao Li, Yuqi Li, Ciliang Sun, Chong Wang, and Jinhui Xiang, “Spec-nerf: Multi-spectral neural radiance fields,” 2023.
- [15] Matthew Phelps, Thomas Swindle, J. Zachary Gazak, and Justin Fletcher, “Improving spectral-based estimation of space object orientation,” in *Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS)*, 2022.
- [16] O. Rainio, J. Teuvo, and R. Klén, “Evaluation metrics and statistical tests for machine learning,” *Scientific Reports*, vol. 14, no. 6086, 2024.
- [17] Serin Varghese, Yasin Bayzidi, Andreas Bär, Nikhil Kapoor, Sounak Lahiri, Jan David Schneider, Nico Schmidt, Peter Schlicht, Fabian Hüger, and Tim Fingscheidt, “Unsupervised temporal consistency metric for video segmentation in highly-automated driving,” in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2020, pp. 1369–1378.
- [18] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick, “Segment anything,” 2023.