

A Photorealistic Dataset and Vision-Based Algorithm for Anomaly Detection During Proximity Operations in Lunar Orbit

Selina Leveugle¹, Chang Won Lee², Svetlana Stolpner³, Chris Langley³,
Paul Grouchy³, Steven Waslander², and Jonathan Kelly¹

Abstract—NASA’s forthcoming Lunar Gateway space station, which will be uncrewed most of the time, will need to operate with an unprecedented level of autonomy. One key challenge is enabling the Canadarm3, the Gateway’s external robotic system, to detect hazards in its environment using its onboard inspection cameras. This task is complicated by the extreme and variable lighting conditions in space. In this paper, we introduce the visual anomaly detection and localization task for the space domain and establish a benchmark based on a synthetic dataset called ALLO (Anomaly Localization in Lunar Orbit). We show that state-of-the-art visual anomaly detection methods often fail in the space domain, motivating the need for new approaches. To address this, we propose MRAD (Model Reference Anomaly Detection), a statistical algorithm that leverages the known pose of the Canadarm3 and a CAD model of the Gateway to generate reference images of the expected scene appearance. Anomalies are then identified as deviations from this model-generated reference. On the ALLO dataset, MRAD surpasses state-of-the-art anomaly detection algorithms, achieving an AP score of 62.9% at the pixel level and an AUROC score of 75.0% at the image level. Given the low tolerance for risk in space operations and the lack of domain-specific data, we emphasize the need for novel, robust, and accurate anomaly detection methods to handle the challenging visual conditions found in lunar orbit and beyond.

Index Terms—Space Robotics and Automation, Data Sets for Robotic Vision, Simulation and Animation, Anomaly Detection.

I. INTRODUCTION

OVER the past two decades, the international space community has begun to focus its efforts on extending human space exploration beyond low-Earth orbit. NASA’s Artemis program aims to deploy the Lunar Gateway, the first space station in lunar orbit, that will test new technologies needed for extended deep-space missions [1]. Unlike the International Space Station (ISS), the Gateway will be required to operate autonomously and without an on-board crew for long periods. Complete autonomy is advantageous for the Canadarm3, the

Manuscript received: August 20, 2025; Revised November 15, 2025; Accepted December 21, 2025.

This paper was recommended for publication by Giuseppe Loianno upon evaluation of the Associate Editor and Reviewers’ comments.

¹Space and Terrestrial Autonomous Robotic Systems (STARS) Laboratory at the University of Toronto Institute for Aerospace Studies (UTIAS), Toronto, Canada. Email: <firstname>.<lastname>@robotics.utias.utoronto.ca

²Toronto Robotics and AI Laboratory (TRAIL) at the University of Toronto Institute for Aerospace Studies (UTIAS), Toronto, Canada. Email: <firstname>.<lastname>@robotics.utias.utoronto.ca

³MDA Space Inc. Email: <firstname>.<lastname>@mda.space
Digital Object Identifier (DOI): see top of this page.

©2026 IEEE




Fig. 1: Visualization of the rendering process used to generate an anomalous image in the ALLO dataset. (a) Positions of the Sun, Moon, and space station are determined and scene lighting is configured accordingly; (b) An anomaly is inserted into the scene; (c) The Blender Cycles engine renders the image; and (d) the corresponding three-class segmentation mask is generated.

external robotic system being developed by MDA Space for the Canadian Space Agency. It will play a multifaceted role on the Gateway, including station maintenance, inspection, and capture and berthing of visiting vehicles [2].

To enhance robotic autonomy, vision-based anomaly detection methods that can detect collision hazards, such as loose tools or debris, are of particular interest. Anomaly detection and localization refer to identifying and pinpointing image regions whose content deviates from the expected distribution of inputs. Anomaly detection in the space domain is particularly challenging because of complex lighting conditions that result from the black background coupled with harsh direct solar illumination. Furthermore, the varied camera viewpoints used during operations near the space station can cause anomalies to fall into shadow regions and blend in with the background, making detection substantially more difficult.

Although reliable automated anomaly detection would be highly valuable, existing datasets and methods do not directly address the problem of anomaly detection in the space domain. Specifically, a significant challenge in space anomaly detection is the absence of comprehensive, labelled datasets necessary for robust evaluation. Moreover, our comparative analysis in this paper demonstrates that existing approaches significantly underperform, failing to generalize effectively to this domain. To address these gaps, we introduce ALLO, a large-scale open-source anomaly detection dataset tailored for robotic

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

proximity operations in space. Using ALLO, we establish a benchmark by evaluating existing anomaly segmentation algorithms. We further present MRAD, a novel statistical algorithm specifically designed for this application that leverages camera pose information and a known CAD model of the station to perform per-image anomaly segmentation, achieving state-of-the-art performance on the ALLO benchmark. In summary, our main contributions are as follows.

- We introduce ALLO, a synthetic dataset for vision-based anomaly detection in the space domain, and release its open-source generation pipeline to facilitate extensions and future research.
- We establish an anomaly detection benchmark using ALLO, evaluating recent anomaly segmentation algorithms and highlighting their shortcomings for the space environment.
- We introduce Model Reference Anomaly Detection (MRAD), a statistical algorithm designed for the space domain, and demonstrate its superior performance over existing anomaly detection algorithms on ALLO through comprehensive evaluation and ablation studies.

The remainder of the paper is structured as follows. In Section II, we review existing anomaly detection datasets and algorithms. In Section III, we detail the ALLO dataset and its rendering pipeline. We then describe our proposed MRAD algorithm in Section IV and evaluate its performance against other state-of-the-art methods on the ALLO dataset in Section V. The ALLO dataset and all open-source code are available at <https://github.com/utiasSTARS/ALLO>.

II. RELATED WORK

Although anomaly detection is an emerging topic in the space domain, there is extensive work on visual datasets for space applications and, separately, on industrial anomaly detection. In this section, we review existing image datasets for these two applications. In addition, we provide an overview of modern anomaly detection methods that will serve as baselines for our benchmark.

A. Datasets for Space Applications and Defect Inspection

1) *Space Applications*: Synthetic visual datasets are widely used in space applications due to the difficulty of acquiring real in situ imagery. Photorealistic rendering engines such as the Blender Cycles [3] engine and Unreal Engine 5 [4] are commonly employed to produce simulated space imagery. Existing datasets typically focus on specific mission scenarios. For instance, Airbus’s SurRender [5] generates images for satellite servicing and debris removal, while the Space Imaging Simulator for Proximity Operations (SISPO) [6] leverages Blender to model asteroid fly-bys.

Physical simulators have also been developed. Legentil et al. [7], for example, capture docking-port imagery under realistic dynamics and challenging lighting conditions to support satellite rendezvous and docking operations. Their dataset explicitly excludes the satellite body to ensure satellite-agnostic applicability for docking-port detection and state estimation.

In contrast, the ALLO dataset provides a comprehensive photorealistic simulation of an entire space station in lunar orbit, specifically designed for anomaly detection to mitigate collision risks. Its open-source data-generation pipeline further enables adaptation to diverse downstream tasks.

2) *Industrial Defect Inspection*: The majority of recent efforts in anomaly detection have focused on finding manufacturing defects during visual industrial inspection. The MVTec 2D anomaly detection dataset [8] is the most widely used benchmark for this application due to its pixel-level annotations and diverse range of objects [9]. However, similar to other anomaly detection datasets such as BTech [10] and Kolektor [11], MVTec images are captured from fixed view-points, under consistent lighting and with simple backgrounds. In contrast, the ALLO dataset offers a more comprehensive and diverse set of scenes, including multiple views of the ISS with varied lighting conditions and complex backgrounds.

B. Anomaly Detection Methods

Anomaly detection involves identifying abnormal samples, such as unexpected or irregular instances, that deviate from an expected distribution. Existing methods fall broadly into two categories: statistical approaches, which rely on (possibly learned) statistical models, and deep learning-based approaches, which use deep neural networks to learn a representation of normal (non-anomalous) data.

1) *Statistical Detectors*: Statistical detection methods compute anomaly scores using distance functions or statistical tests, followed by binary classification via empirical thresholds or probabilistic models [12]. Classical techniques such as one-class support vector machines (OC-SVM) [13] and support vector data descriptors (SVDD) [14] model the density or decision boundary of normal data to identify outliers. The Reed–Xiaoli detector (RXD) [15], originally developed for use with hyperspectral aerial imagery, assumes a multivariate Gaussian distribution for background pixels and detects anomalies as probabilistic outliers. However, this assumption fails for complex, non-homogeneous scenes such as those encountered in space.

Several RXD extensions relax its homogeneity requirement. The probabilistic anomaly detector (PAD) [16] separates pixels into target and background sets and compares their RXD scores, while the random selection-based anomaly detector (RSAD) [17] iteratively redefines the background set to improve robustness. Nevertheless, these methods remain sensitive to background variability and are not designed to exploit geometric priors available in space applications.

2) *Deep Learning-Based Detectors*: Most modern deep learning-based anomaly detectors are unsupervised due to the scarcity of labelled anomalies, although some self-supervised approaches introduce synthetic anomalies during training [18]. Unsupervised methods [19], [20] learn a representation of normal data and detect deviations at inference time. Many operate in feature space, using pre-trained backbones (e.g., ResNet) and either fit statistical models to extracted features [20], [21] or estimate their density with normalizing flows [22], [23]. In [24], an encoder-decoder framework is combined with anomaly scores similar to RXD to identify pixels far from

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

a Gaussian distribution. Alternatively, generative approaches reconstruct only normal images, flagging anomalies by their reconstruction error [9]. Student–teacher architectures [19], [25] follow this paradigm by distilling normal data representations from a teacher to a student, thereby failing to reproduce unseen anomalous patterns.

Although these methods achieve state-of-the-art results in structured domains such as industrial defect detection, they generalize poorly to the space domain due to high scene variability, complex lighting, and limited representative training data for learning.

III. THE ALLO PIPELINE AND DATASET

Our first contribution is a dataset for space-based anomaly detection and localization. The ALLO dataset comprises 51,409 synthetic images simulating the visual feed from arm-mounted cameras. The dataset is divided into a training set of normal (anomaly-free) images and a test set containing both normal and anomalous examples. Each image is generated with the Blender Cycles path-tracing engine, realistically simulating the lighting conditions in lunar orbit, including the blackness of deep space, direct sunlight, and inspection lighting. The image-generation pipeline, described below, is fully open source and can be adapted for related computer vision tasks, such as semantic segmentation and depth estimation.

A. Scene Modelling

We use a high-fidelity Blender model of the ISS [26] as a stand-in for the Lunar Gateway, since a detailed Gateway model is currently unavailable. Given the expected architectural and surface similarities of the Gateway to the ISS, this model provides a realistic approximation of the visual environment anticipated for the Canadarm3’s operations. The Gateway’s anticipated orbit [27] is approximated by an ellipse matching the expected perilune and apolune.

The positions of the Earth and Sun are determined over a 365-day period relative to the Moon using 2030 ephemeris data from the Skyfield library [28]. These bodies are positioned relative to the ISS model to replicate the background views expected by the arm’s cameras. For computational efficiency, all models are scaled down in size but positioned to preserve their apparent angular sizes from the station’s perspective, ensuring visual fidelity.

To simulate arm-mounted camera perspectives, 50 unique base poses are manually defined around the ISS. These poses are derived from possible Canadarm3 trajectories near the mounting fixtures located on the exterior of the station. To introduce realistic pose imprecision and enhance viewpoint diversity, each base pose is perturbed with Gaussian noise during rendering, using standard deviations of 1 m for position and 0.2 rad for orientation (roll, pitch, and yaw). This randomization strategy ensures comprehensive coverage of the key station regions and subtle yet critical variations in perspective.

B. Image Rendering

Different camera positions are used for the training and test sets to avoid overlap between views. The test set contains




Fig. 2: Examples of models of anomalous objects used in the ALLO dataset.

16 distinct anomalous objects representing potential collision hazards (see Fig. 2), including dislodged station equipment such as thermal blankets or cables, items lost during extravehicular activities such as tools or astronaut gloves, and objects resembling experimental equipment from prior missions such as Cubesats and small satellites.

For each anomalous scene, only one such object is inserted into the camera’s view frustum to simplify evaluation. The object is randomly positioned within the station’s vicinity and rendered at five depths: one initial random depth and additional positions displaced by ± 1 m and ± 2 m from this initial position. If the object is occluded or projected to occupy less than 0.1% of the image, it is repositioned to ensure visibility. This procedure yields a diverse set of anomaly appearances while maintaining realistic spatial consistency.

Each rendered image has a resolution of $1,920 \times 1,080$ pixels and is paired with a three-class segmentation mask delineating the background, non-anomalous foreground (e.g., station surfaces, celestial bodies), and anomalous regions. The rendering pipeline, illustrated in Fig. 1, is repeated with two random seeds to expand the dataset. The dataset composition is reported in Table I, with sample images shown in Fig. 3.

In summary, scenes were generated for each day in the ephemeris data by first positioning the Moon, Earth, and Sun according to their ephemerides. Camera poses were then selected, with indices 0–39 assigned to normal scenes and 40–49 reserved for anomalous ones. For each camera pose and anomaly configuration, the following steps were executed:

- 1) Place the camera with a Gaussian pose perturbation.
- 2) Add a nearby spotlight.
- 3) If the scene is anomalous:
 - 3.1) Insert the anomaly into the view frustum.
 - 3.2) Reposition the anomaly for visibility if necessary.
- 4) Add noise and glare.
- 5) Render the image and corresponding segmentation mask using the Blender Cycles engine.

C. Image Validation

To validate the realism of our synthetic ALLO images, we compared them with real photographs of the ISS from past ISS missions. We matched Blender camera poses to these reference images, ensuring precise geometric alignment, producing synthetic renderings that closely replicate real lighting conditions and surface textures (see Fig. 4).

Quantitatively, we used the learned perceptual image patch similarity metric (LPIPS) [29] across ten real–synthetic pairs, obtaining an average score of 0.58—comparable to state-of-the-art image synthesis models such as GILL (~ 0.70 on VIST) [30] and Stable Diffusion XL (0.88 on COCO2017) [31].




Fig. 3: Sample images from the ALLO dataset. *Top Row*: Normal (anomaly-free) images captured from five distinct camera views of the ISS. *Bottom Row*: Images containing anomalies, identified by red ellipses, that include free-floating equipment such as cables, thermal blankets, and a hand drill.




Fig. 4: Real images of the ISS (left) and their synthetic recreations (right).

We also calculated the Fréchet inception distance (FID) [32] between 319 ALLO training images and historical ISS photographs, achieving a score of 3.61. This FID score surpasses those of other recent models, such as DiGIT [33] (4.59 on ImageNet) and LC-GAN [34] (5.77 on AFHQ-v2). Together, these two metrics confirm the high quality and photorealism of our synthetic images.

Together, these qualitative and quantitative results confirm that the synthetic images effectively reproduce the visual characteristics of the target space environment.

IV. MODEL REFERENCE ANOMALY DETECTION

MRAD is an anomaly detection algorithm designed for the space domain. It detects anomalies by comparing a query image with a reference image rendered from a known camera pose and a CAD model of the station. The algorithm operates in three stages: (1) pixel-level anomaly scores are computed using an extended version of the Reed–Xiaoli detector (RXD) applied to query–reference pairs; (2) region growing clusters high-scoring pixels into candidate anomalous regions, producing a binary mask; and (3) an image is classified as anomalous

TABLE I: Composition of the ALLO dataset (camera poses and images) across the training and test sets. The training set is composed exclusively of normal images, while the test set contains both normal and anomalous images rendered from disjoint sets of camera poses.

Dataset Set	Training	Test	Total
Camera Poses	1–40	41–50	50
Normal Images	29,085	7,300	36,385
Anomalous Images	0	15,024	15,024

if the total detected area exceeds a threshold. In this work, both reference and query images are synthetically generated, and the real-to-synthetic domain gap is simulated by introducing random variations in lighting and small perturbations to the camera pose when rendering the reference images.

A. Reed-Xiaoli Detector (RXD) Anomaly Score

Each pixel in the query image is first assigned an anomaly score based on the Reed–Xiaoli detector (RXD) [15], which measures a pixel’s statistical distance from the background. RXD models nominal pixels as samples from a multivariate Gaussian distribution with mean vector $\boldsymbol{\mu} \in \mathbb{R}^3$ and covariance matrix $\boldsymbol{\Sigma} \in \mathbb{R}^{3 \times 3}$. Pixels with low likelihood under $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ are considered anomalous. Given an image I with mean $\boldsymbol{\mu}_I$ and covariance $\boldsymbol{\Sigma}_I$, the RXD anomaly score for a pixel \mathbf{x} is the squared Mahalanobis distance:

$$\text{RXD}_I(\mathbf{x}) = (\mathbf{x} - \boldsymbol{\mu}_I)^T \boldsymbol{\Sigma}_I^{-1} (\mathbf{x} - \boldsymbol{\mu}_I). \quad (1)$$

While effective for relatively homogeneous scenes, RXD struggles with non-uniform backgrounds, such as space imagery, where lighting and texture vary significantly. To address this limitation, we adapt principles from the probabilistic anomaly detector (PAD) [16], which compares likelihoods under both nominal and anomalous distributions. In MRAD, the nominal (anomaly-free) distribution is derived from the reference image, while the query image provides statistics for potentially anomalous regions. Specifically, the anomaly-free distribution statistics $(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$ are computed from the reference image, and the query distribution statistics $(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1)$ are computed from the query image. By leveraging statistics from both the reference and query images, MRAD can better handle complex, varying illumination and background structure, where the pixel distribution is non-Gaussian. The MRAD anomaly score for a pixel \mathbf{x} is then defined as:

$$\text{MRAD}(\mathbf{x}) = \text{RXD}_{\text{ref}}(\mathbf{x}; \boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0) - \text{RXD}_{\text{query}}(\mathbf{x}; \boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1), \quad (2)$$

where, given nominal image alignment, $\text{RXD}_{\text{ref}}(\mathbf{x})$ is expected to be greater than or equal to $\text{RXD}_{\text{query}}(\mathbf{x})$.

To account for spatial variability in lighting and texture (e.g., the contrast between the black background of space and the station’s surface), both reference and query images are divided into 4×4 grids. Mean and covariance statistics are computed per grid, yielding localized anomaly scores that better model region-specific variations. Scores are scaled

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.




Fig. 5: Left to right: reference image from the ALLO test set with one grid cell highlighted; pixel intensity histogram of the highlighted cell; anomalous query image; predicted anomaly map.

in the range $[0, 255]$ to produce a grayscale MRAD map, where pixels with values below a predefined threshold are classified as anomaly-free and set to zero, and remaining high-scoring pixels are retained for further analysis. Fig. 5 illustrates a case where the pixel-intensity histogram for a reference grid cell is noticeably skewed and non-Gaussian, violating the Gaussian noise assumption. Nevertheless, in this instance MRAD remains robust, producing an accurate anomaly map that correctly localizes the free-floating cubesat.

B. Noise Removal

A major source of false positives is random amplification in black background regions, where small color variations can produce spuriously high anomaly scores that are not effectively filtered by simple thresholding. To mitigate this issue, we analyze the spatial distribution of anomaly scores within each grid cell before generating the final anomaly score map.

First, a preliminary check identifies grid cells that are likely to contain a large anomaly: if more than 90% of the scores in a cell exceed the minimum MRAD score, the entire cell is marked as anomalous. Next, a kernel-based method removes noise caused by random score variations. A sliding-window kernel (with a size $\sim 20\%$ of the grid dimensions) is used to compute the local standard deviation of anomaly scores; regions with a standard deviation below an empirical threshold (about 3–5%) are considered to be indicative of uniform background noise and their scores are set to zero. Finally, a Gaussian blur is applied to the noise-filtered score map to smooth local fluctuations and merge fragmented anomalous regions, improving region contiguity for subsequent detection.

C. Clustering via Region Growing

Region growing [35] is employed to cluster anomalous pixels into distinct objects based on pixel location and anomaly score. The implementation in MRAD relies on three parameters: the seed (initial pixel for expansion), the footprint (square search neighbourhood), and the tolerance (allowable score deviation for aggregation). Both the footprint and the tolerance are set empirically.

Seed pixels are found by applying mean-shift clustering [36] to the blurred score map. This non-parametric, mode-seeking algorithm detects clusters by shifting each pixel toward the region of highest density, iterating until convergence at a local maximum. The resulting cluster centroids serve as seed points for region growing.

Region growing then expands from each seed, aggregating spatially adjacent pixels with similar anomaly scores until no further candidates meet the criteria. This yields a binary anomaly map, from which the total number of anomalous pixels is computed. An image is classified as anomalous if this count exceeds a predefined threshold. The complete MRAD pipeline is outlined in Algorithm 1 and Fig. 6.

V. EXPERIMENTS

Using the ALLO benchmark, we evaluate our proposed algorithm, MRAD, against recent unsupervised anomaly detection algorithms from the industrial and medical domains. We then analyze the limitations of these competing methods and investigate the factors that contribute to the superior performance of MRAD.

A. Experimental Setup

We modified the Intel Anomalib library [37] to load, train, and evaluate learning-based algorithms on the ALLO dataset. We selected nine representative learning-based methods from the library, chosen to cover both reconstruction- and representation-based approaches and for their established performance on the MVTEC industrial benchmark.

Normalizing flow-based methods, FastFlow [23] and UFlow [38], learn the probability distribution of normal image features and detect anomalies as samples with low likelihood. Student-teacher methods, including STFPM [25], Reverse Distillation [19] and UniNet [21] train a smaller student network to mimic the feature representations of a pre-trained teacher, flagging deviations as anomalies. UniNet introduces domain-aware feature selection, contrastive and margin losses, and a bottleneck with a weighted decision mechanism to support one model across multiple object categories without per-category training. Self-supervised and reconstruction-based methods, such as DSR [18], SuperSimpleNet [39], and Dinomaly [40],

Algorithm 1: Model-Reference Anomaly Detection

```

ScoreImage  $\leftarrow$  Zeros(SizeOf(QueryImage));
AnomalyMap  $\leftarrow$  Zeros(SizeOf(QueryImage));
foreach GridCell in QueryImage do
  foreach Pixel in GridCell do
    Score  $\leftarrow$  MRAD(Pixel); // See Eq. (2)
    if Score > MinMRADScore then
      | ScoreImage[Pixel]  $\leftarrow$  Score;
  StdDev  $\leftarrow$  StdDeviation(GridCell, Kernel);
  if Range(StdDev) < 3% then
    | Score = 0 for all pixels in GridCell;
Apply Gaussian blur to ScoreImage;
Apply mean shift clustering to ScoreImage;
Seeds  $\leftarrow$  FindClusterCentroids(ScoreImage);
foreach StartPoint in Seeds do
  AnomalyMap[StartPoint]  $\leftarrow$  1; // Anomalous
  foreach Pixel in Footprint(StartPoint) do
    if InBounds(Pixel, Tolerance) then
      | AnomalyMap[Pixel]  $\leftarrow$  1; // Anomalous
Count  $\leftarrow$  Sum(AnomalyMap);
if Count > AnomalyThreshold then
  | ImageClass  $\leftarrow$  Anomalous;
else
  | ImageClass  $\leftarrow$  Normal;

```




Fig. 6: Visualization of the steps in the MRAD algorithm. The query image (a) and reference image (b) are the algorithm inputs. The anomaly score image (c), computed using (2), highlights the anomalous thermal blanket. Noise is removed and a Gaussian blur is applied, followed by mean-shift clustering to identify region growing seed points (d). Region growing aggregates anomalous pixels to generate the final anomaly map (e), detecting the blanket while rejecting false positives. Notably, camera pose variations may introduce additional image differences that could be mistaken for anomalies if detection relied solely on pixel differences alone.

reconstruct normal images. Some approaches leverage synthetic anomalies for supervision, whereas Dinomaly uses a minimalist pure-Transformer reconstructor. Finally, CFA [20] is a feature-adaptation method that maps normal features to a shared hypersphere, detecting anomalies as features that fail to map correctly. Representation-based methods generally use fixed encoders pre-trained on ImageNet [41] and train remaining layers from scratch on the target dataset.

The proposed MRAD algorithm and the statistical detectors RXD and PAD were also evaluated on the ALLO test set with additional rendered reference images. RXD operates solely on the query image, while PAD and MRAD use both query and reference images; their anomaly scores were thresholded to produce output binary anomaly maps. Ablation studies evaluated the impact of grid size, region-growing tolerance, kernel size, and the standard deviation used for noise removal on MRAD’s detection performance. Parameters were varied around their empirically optimized values to assess robustness.

Performance was measured using image-level AUROC (I.AUROC), pixel-level AUROC (P.AUROC), pixel-level average precision (PAP), and the pixel-level false-positive rate at 95% true-positive rate (P.FPR95). On the imbalanced ALLO dataset, P.AUROC can be dominated by the negative class, obscuring localization failures. PAP, which measures the precision-recall trade-off specifically for anomalous pixels, provides a more informative localization metric in this setting [22]. Consequently, PAP was used as the primary metric for fine-grained anomaly segmentation, followed by P.FPR95.

B. ALLO Benchmark Results

The performance of existing learning-based algorithms and MRAD on the ALLO test set is summarized in Table II. Among the learning-based methods, Dinomaly achieves the highest PAP of 49.4%, followed by FastFlow (29.1%) and UFlow (20.7%), and also attains the best I.AUROC and P.AUROC on ALLO with scores of 78.6% and 91.4%, respectively. However, its performance dropped substantially relative to MVTec, with I.AUROC decreasing by 21.0 points (99.6% to 78.6%) and P.AUROC by 7.0 points (98.4% to 91.4%).

In contrast, MRAD, our learning-free method, outperforms all learning-based and learning-free baselines on ALLO (see Table III), achieving the highest PAP of 62.9%, best FPR95 of 37.5%, and the second-highest I.AUROC of 75.0%. These results demonstrate that the grid-based dual statistical framework of MRAD enables more accurate detection of fine-grained anomalies. Qualitative examples of FastFlow, STFPM and MRAD are shown in Fig. 7.

C. Discussion

1) *Limitations of Learning-Based Methods:* Analysis of our results (Section V-B) reveals the limitations of established anomaly detection methods within the space domain. Many leverage feature extractors pre-trained on ImageNet, a dataset whose visual characteristics differ significantly from ALLO. Consequently, these networks generate non-discriminative features when applied to ALLO data, compromising their effectiveness in learning the dataset’s distribution.


In addition, learning-based anomaly detection methods are often limited by their assumptions on the distribution of normal data. Approaches based on unimodality [23] fail when anomalies closely resemble normal samples, or when the normal data distribution is multi-modal. Many algorithms also presume image consistency—fixed lighting and viewpoints typical of industrial inspection tasks. On datasets like ALLO, where lighting and viewing conditions vary and misalignments are common, anomaly scores become unreliable. Feature-reconstruction models such as Dinomaly [40], leverage vision foundation models to learn rich features and achieve the strongest learning-based results on ALLO. Likewise, normalizing flow methods (FastFlow [23], UFlow [38]) better model complex distributions and reduce false positives. Nonetheless, even these advanced methods struggle to detect anomalies that closely resemble the space station, particularly under variable illumination or perspective.

Ultimately, learned methods struggle on ALLO because they attempt to model a single global distribution of non-anomalous features across all views and lighting. In contrast, MRAD analyzes each query image individually via direct reference comparison, eliminating the need for global modeling and enabling robust anomaly localization regardless of lighting or viewing variation.

2) *Statistical (Learning-Free) Anomaly Detection:* Further analysis of our results demonstrates how a non-learning algorithm can be used for anomaly detection and which components most strongly affect its performance. The value of a reference image is evident in the performance gap between RXD and PAD: computing background statistics from reference rather than query images increases the PAP of RXD by 39.5%. By comparing RXD scores between the query and reference images, MRAD attenuates the impact of lighting variations and relaxes the Gaussian pixel-distribution assumption of RXD. Combined with noise removal and region growing to suppress false positives and cluster anomalous pixels, MRAD yields a 5.8% increase in PAP and an 18.7% improvement in I.AUROC over PAD.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.TABLE II: Performance of state-of-the-art anomaly detection algorithms on the ALLO test set. For each metric, the three top-performing methods are highlighted in **red**, **orange** and **yellow**, respectively.

Metric	DSR [18]	Rev. Dist. [19]	SuperSimpleNet [39]	CFA [20]	Dinomaly [40]	UniNet [21]	STFPM [25]	UFlow [38]	FastFlow [23]	MRAD (ours)
I.AUROC \uparrow	55.6 \pm 2.0	57.4 \pm 4.4	49.6 \pm 0.5	51.3 \pm 0.5	78.6 \pm 0.5	56.8 \pm 2.9	61.8 \pm 2.4	60.7 \pm 3.2	65.9 \pm 0.6	75.0
P.AUROC \uparrow	69.5 \pm 6.5	75.7 \pm 0.7	79.5 \pm 9.5	84.8 \pm 7.7	91.4 \pm 0.1	85.9 \pm 2.2	87.8 \pm 1.7	85.8 \pm 1.6	90.4 \pm 0.5	71.8
P.AP \uparrow	6.9 \pm 0.7	7.0 \pm 2.5	9.3 \pm 1.0	11.0 \pm 2.8	49.4 \pm 1.0	12.3 \pm 2.4	12.5 \pm 1.3	20.7 \pm 1.6	29.1 \pm 1.2	62.9
P.FPR95 \downarrow	95.4 \pm 2.5	95.0 \pm 1.7	64.4 \pm 18.3	58.7 \pm 29.1	47.7 \pm 0.8	51.1 \pm 13.6	43.0 \pm 7.0	66.7 \pm 10.6	41.7 \pm 1.4	37.5

Fig. 7: Examples of predicted anomaly masks for two images from the ALLO test set by two learning-based algorithms, FastFlow and STFPM, and by MRAD. Ground truth masks are provided for comparison. *Top Row*: The anomalous object, a thermal blanket, is located in the upper left of the query image. *Bottom Row*: The anomalous object, a power cable, is located in the upper centre of the query image.TABLE III: Performance of MRAD with ablations and of non-learning algorithms on the ALLO test set. The optimized settings are a 4×4 grid, a tolerance of 50 MRAD score, kernel size of 20%, and a SD of 5.

Method	Ablation	I. AUROC \uparrow	P. AUROC \uparrow	P. AP \uparrow
RXD		50.0	55.1	17.6
PAD		56.3	70.1	57.1
MRAD—Grid Size	8x8	75.7	73.6	60.0
	6x6	76.3	71.8	61.3
	3x3	72.1	68.1	61.3
	2x2	65.6	65.1	59.2
	No grids	57.9	58.9	54.7
MRAD—Tolerance	25%	74.9	69.4	60.7
	50%	74.9	68.3	61.0
	150%	67.5	76.6	58.9
	200%	68.5	67.3	54.9
MRAD—Kernel Size	5%	74.9	70.8	61.7
	10%	74.9	70.0	62.1
	20%	75.0	71.8	62.9
	30%	74.9	70.3	60.1
MRAD—Noise SD	SD 2	73.0	70.7	62.8
	SD 5	75.0	70.6	62.9
	SD 10	67.5	77.3	59.4
Optimized	–	75.0	71.8	62.9

The ablation studies identify grid size and region-growing tolerance as the dominant parameters. Increasing the grid size beyond the optimized 4×4 configuration degrades both pixel-level and image-level performance, as large grids favor global context at the expense of fine-grained anomalies. In contrast, smaller grids exploit the reference image more effectively by emphasizing local deviations. Region growing tolerance governs how anomaly scores are aggregated: lower tolerances cause a modest decrease in P.AP with little effect on I.AUROC, while higher tolerances rapidly degrade performance by incorrectly including non-anomalous pixels. For noise handling, kernel size has a larger influence than the Gaussian standard deviation; larger kernels risk removing truly anomalous pixels, whereas smaller kernels fail to suppress background fluctuations. The optimized configuration (a 4×4 grid, tolerance of 50

MRAD score, kernel size of 20%, and standard deviation of 5) provides a practical balance between localization sensitivity and robustness.

These results illustrate how MRAD’s simple statistical formulation and few tunable parameters yield interpretable behavior and straightforward adaptation to different operational scenarios. Unlike deep learning approaches, which typically behave as black boxes and often require retraining, MRAD’s explicit use of background statistics produces predictable responses to changes in illumination and scene content. Moreover, this statistical approach avoids the data dependence of deep learning: acquiring large on-orbit datasets is prohibitively expensive, and training on synthetic imagery introduces a sim-to-real gap that can undermine on-orbit performance. MRAD sidesteps these issues by generating a reference image at inference time from a known CAD model and pose estimate.

This design also introduces limitations. Because MRAD does not learn from data, its performance is constrained by the fidelity of the CAD model, the quality of the rendered reference image, and the validity of its statistical assumptions.

Finally, while the present analysis focuses on free-floating anomalies, since MRAD operates on a per-image basis with a rendered reference, the same framework could naturally be extended to inspection tasks such as detecting structural damage or surface degradation.

VI. CONCLUSION

In this paper, we introduced ALLO, the first open-source anomaly detection dataset with synthetic lunar-orbit imagery, and established a benchmark for space-based anomaly detection by evaluating state-of-the-art methods. We also proposed MRAD, a statistical algorithm that leverages known camera poses and a CAD model to render anomaly-free reference images and detect statistical deviations. On ALLO, MRAD achieves state-of-the-art performance, outperforming existing methods in pixel-level AP, FPR95, and image-level AUROC.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

While MRAD achieves strong results, its reliance on accurate CAD models and runtime rendering introduces practical constraints. Learning-based methods, though potentially more computationally efficient, currently fail to generalize to the varied lighting and geometric conditions of the space domain. A promising direction for future work is to develop novel learning strategies, where models are trained to discriminate between normal and anomalous synthetic examples. By explicitly teaching a model what constitutes an anomaly in this specific domain, we could enhance its discriminative power and robustness to the extreme lighting and viewpoint variability, enabling reliable anomaly detection for future autonomous space operations.

REFERENCES

- [1] J. C. Crusan, R. M. Smith, D. A. Craig, J. M. Caram, J. Guidi, M. Gates, J. M. Krezel, and N. B. Herrmann, "Deep space gateway concept: Extending human presence into cislunar space," in *Proc. IEEE Aerosp. Conf.*, Mar. 2018, pp. 1–10.
- [2] (2020, June) About Canadarm3. <https://www.asc-csa.gc.ca/eng/canadarm3/about.asp>. Last updated: 2025-02-25.
- [3] Blender: The free and open source 3D creation software. <https://www.blender.org/>.
- [4] Unreal Engine 5. <https://www.unrealengine.com/en-US/unreal-engine-5>.
- [5] R. Brochard, J. Lebreton, C. Robin, K. Kanani, G. Jonniaux, A. Masson, N. Despré, and A. Berjaoui, "Scientific image rendering for space scenes with the SurRender software," in *Intl. Astronautical Conf. (IAC)*, Bremen, Germany, 2018.
- [6] M. Pajusalu, I. Iakubivskiy, G. J. Schwarzkopf, O. Knuutila, T. Väisänen, M. Bühner, H. Teras, G. L. Bonhomme, M. F. Palos, J. Praks, and A. Slavinskis, "SISPO: Space imaging simulator for proximity operations," *PLOS ONE*, vol. 17, no. 3, Mar. 2022.
- [7] C. L. Gentil, J. Naylor, N. Munasinghe, J. Mehama, B. Dai, M. Asavkin, D. G. Dansereau, and T. Vidal-Calleja, "Mixing data-driven and geometric models for satellite docking port state estimation using an RGB or event camera," in *Intl. Conf. Robotics Automat. (ICRA)*, 2025.
- [8] P. Bergmann, M. Fauser, D. Sattlegger, and C. Steger, "MVTec AD—A comprehensive real-world dataset for unsupervised anomaly detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2019, pp. 9584–9592.
- [9] Y. Zheng, X. Wang, Y. Qi, W. Li, and L. Wu, "Benchmarking unsupervised anomaly detection and localization," *arXiv preprint arXiv:2205.14852*, 2022.
- [10] MyNew, "Btech dataset," <https://universe.roboflow.com/mynew/btech>, Feb. 2024, visited on 2024-09-11.
- [11] D. Tabernik, S. Šela, J. Skvarč, and D. Škočaj, "Segmentation-Based Deep-Learning Approach for Surface-Defect Detection," *J. Intell. Manuf.*, May 2019.
- [12] L. Ruff, J. R. Kauffmann, R. A. Vandermeulen, G. Montavon, W. Samek, M. Kloft, T. G. Dietterich, and K.-R. Müller, "A unifying review of deep and shallow anomaly detection," *Proc. IEEE*, vol. 109, no. 5, pp. 756–795, May 2021.
- [13] L. M. Manevitz and M. Yousef, "One-class SVMs for document classification," *J. Mach. Learn. Res.*, vol. 2, pp. 139–154, 2001.
- [14] H. Hoffmann, "Kernel PCA for novelty detection," *Pattern Recognit.*, vol. 40, no. 3, pp. 863–874, Mar. 2007.
- [15] I. Reed and X. Yu, "Adaptive multiple-band CFAR detection of an optical pattern with unknown spectral distribution," *IEEE Trans. Acoustics, Speech, Signal Process.*, vol. 38, no. 10, pp. 1760–1770, 1990.
- [16] L. Gao, Q. Guo, A. J. Plaza, J. Li, and B. Zhang, "Probabilistic anomaly detector for remotely sensed hyperspectral data," *J. Applied Remote Sens.*, vol. 8, no. 1, p. 083538, Nov. 2014.
- [17] B. Du and L. Zhang, "Random-selection-based anomaly detector for hyperspectral imagery," *IEEE Trans. Geosci. Remote Sens.*, vol. 49, no. 5, pp. 1578–1589, May 2011.
- [18] V. Zavrtnik, M. Kristan, and D. Škočaj, "DSR—A dual subspace re-projection network for surface anomaly detection," in *Computer Vision – ECCV 2022*, ser. Lecture Notes in Computer Science, S. Avidan, G. Brostow, M. Cissé, G. M. Farinella, and T. Hassner, Eds. Springer Nature Switzerland, 2022, pp. 539–554.
- [19] H. Deng and X. Li, "Anomaly detection via reverse distillation from one-class embedding," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, June 2022, pp. 9727–9736.
- [20] S. Lee, S. Lee, and B. C. Song, "CFA: Coupled-hypersphere-based feature adaptation for target-oriented anomaly localization," *IEEE Access*, vol. 10, pp. 78 446–78 454, 2022.
- [21] S. Wei, J. Jiang, and X. Xu, "Uninet: A contrastive learning-guided unified framework with feature selection for anomaly detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2025, pp. 9994–10003.
- [22] D. Gudovskiy, S. Ishizaka, and K. Kozuka, "CFLOW-AD: Real-time unsupervised anomaly detection with localization via conditional normalizing flows," in *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis.*, Jan. 2022, pp. 1819–1828.
- [23] J. Yu, Y. Zheng, X. Wang, W. Li, Y. Wu, R. Zhao, and L. Wu, "FastFlow: Unsupervised anomaly detection and localization via 2D normalizing flows," *arXiv preprint arXiv:2111.07677*, 2021.
- [24] M. Sodano, F. Magistri, L. Nunes, J. Behley, and C. Stachniss, "Open-world semantic segmentation including class similarity," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2024.
- [25] G. Wang, S. Han, E. Ding, and D. Huang, "Student-teacher feature pyramid matching for anomaly detection," in *Proc. British Machine Vis. Conf.*, Mar. 2021.
- [26] International space station 3D model - NASA science. <https://science.nasa.gov/resource/international-space-station-3d-model/>.
- [27] D. E. Lee, "White paper: Gateway destination orbit model: A continuous 15 year NRHO reference trajectory," NTRS Author Affiliations: NASA Johnson Space Center NTRS Report/Patent Number: JSC-E-DAA-TN72594 NTRS Document ID: 20190030294 NTRS Research Center: Johnson Space Center (JSC).
- [28] B. Rhodes, "Skyfield: High precision research-grade positions for planets and Earth satellites generator," Astrophysics Source Code Library, record ascl:1907.024, Jul. 2019.
- [29] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 586–595.
- [30] J. Y. Koh, D. Fried, and R. R. Salakhutdinov, "Generating images with multimodal language models," *Adv. Neural Inf. Process. Syst.*, vol. 36, pp. 21 487–21 506, 2023.
- [31] D. Podell, Z. English, K. Lacey, A. Blattmann, T. Dockhorn, J. Müller, J. Penna, and R. Rombach, "SDXL: Improving latent diffusion models for high-resolution image synthesis," in *The 12th Intl. Conf. Learn. Represent.*, 2023.
- [32] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, "Gans trained by a two time-scale update rule converge to a local nash equilibrium," *Adv. Neural Inf. Process. Syst.*, vol. 30, 2017.
- [33] Y. Zhu, B. Li, H. Zhang, X. Li, L. Xu, and L. Bing, "Stabilize the latent space for image autoregressive modeling: A unified perspective," *Adv. Neural Inf. Process. Syst.*, vol. 37, pp. 28 636–28 661, 2024.
- [34] S. Lee, M. Kim, Y. Chae, and B. Stenger, "Linearly controllable GAN: Unsupervised feature categorization and decomposition for image generation and manipulation," in *European Conf. Comput. Vis.* Springer, 2024, pp. 229–245.
- [35] R. Adams and L. Bischof, "Seeded region growing," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 16, no. 6, pp. 641–647, 1994.
- [36] Y. Cheng, "Mean shift, mode seeking, and clustering," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 17, no. 8, pp. 790–799, 1995.
- [37] S. Akcay, D. Ameln, A. Vaidya, B. Lakshmanan, N. Ahuja, and U. Genc, "Anomalib: A deep learning library for anomaly detection," in *Proc. IEEE Intl. Conf. Image Process.*, Oct. 2022, pp. 1706–1710.
- [38] M. Taitanian, A. Pardo, and P. Muse, "U-flow: A u-shaped normalizing flow for anomaly detection with unsupervised threshold," *J. Mathematical Imaging Vis.*, vol. 66, no. 4, pp. 678–696, Aug. 2024.
- [39] B. Rolih, M. Fučka, and D. Škočaj, "SuperSimpleNet: Unifying unsupervised and supervised learning for fast and reliable surface defect detection," in *Int. Conf. Pattern Recognit. (ICPR)*. Springer, 2025, pp. 47–65.
- [40] J. Guo, S. Lu, W. Zhang, F. Chen, H. Li, and H. Liao, "Dinomaly: The less is more philosophy in multi-class unsupervised anomaly detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, June 2025, pp. 20 405–20 415.
- [41] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Adv. Neural Inf. Process. Syst.*, vol. 25. Curran Associates, Inc., 2012.