Leveraging Astronomical Surveys for Space Debris Detection and Characterization: Application to VST/OmegaCAM

Elisabeth Rachith
Laboratory of Astrophysics EPFL LASTRO, Observatoire de Sauverny, CH-1290 Versoix
elisabeth.rachith@epfl.ch

Stephan Hellmich
Laboratory of Astrophysics EPFL LASTRO, Observatoire de Sauverny, CH-1290 Versoix
stephan.hellmich@epfl.ch

Jean-Paul Kneib
Laboratory of Astrophysics EPFL LASTRO, Observatoire de Sauverny, CH-1290 Versoix
jean-paul.kneib@epfl.ch

ABSTRACT

As the growing population of space debris increasingly threatens satellite operations and orbital sustainability, new cost-effective methods for characterizing these objects are urgently needed. This study presents a pipeline that repurposes archival wide-field astronomical survey data from the VLT Survey Telescope (VST) and its OmegaCAM imager to detect and analyze satellite and debris streaks. Using a convolutional neural network (CNN) architecture enhanced with a Hough transform prior, the authors develop a machine-learning approach tailored to identify streaks across thousands of high-resolution images. The detection algorithm is trained and validated on a hybrid dataset that includes both real and realistically simulated satellite and space-debris streaks. A post-processing pipeline refines detections, performs astrometric calibration, correlates objects using Two-Line Element (TLE) catalogs, and finally extracts lightcurves, enabling the study of rotational states and spin rates of non-operational objects. Performance metrics indicate high precision (0.97) and recall (0.96) on validation sets, with good generalization to real-world data, although with some degradation due to image variability. A six-month archival analysis of OmegaCAM data from 2023 reveals a significant fraction ($\approx$20%) of detected streaks uncorrelated with known catalogs, suggesting the existence of an undercharacterized population of small Low Earth Orbit (LEO) debris. The authors introduce a repository – ODLI (Orbital Debris Lightcurve Inventory) – to store and eventually share extracted data products. This work not only demonstrates the viability of using astronomical data for Space Situational Awareness (SSA) but also lays the groundwork for expanding to other imaging archives in the future.

1. INTRODUCTION

Despite decades of monitoring and notable progress in recent years, our understanding of many key characteristics of space debris – particularly how these objects evolve over time and interact with their environment – remains limited [1]. This gap in knowledge is becoming increasingly problematic as the population of debris in Earth's orbit continues to grow, posing risks to operational satellites, ongoing space missions, and even human life. Although substantial efforts have been made to observe and catalog these objects, critical information about their physical properties and long-term dynamics is still lacking. Reducing these uncertainties is essential for the sustainability of orbital operations, making any additional insight or data extraordinarily valuable.

One promising yet underutilized avenue for gathering such information is through wide-field astronomical surveys. Primarily conducted for astrophysical research, these surveys frequently and inadvertently image space debris, which appear as streaks across their frames. Although these traces are usually considered unwanted by astronomers, they can actually reveal important details about space debris. Analyzing such archival data, however, comes with its own

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

challenges: unlike dedicated observations, these debris signatures are incidental and harder to interpret, and the sheer volume of data collected over many years can be challenging to process.

This work focuses on developing a method to efficiently identify space-debris streaks in wide-field astronomical images and extract informative features from them. While targeted observations using specialized instruments have significantly advanced debris monitoring, they are often expensive, time-consuming, and constrained by scheduling and observational limitations. In contrast, our approach exploits publicly available data from sources such as the European Southern Observatory (ESO), offering a cost-effective way to enhance existing Space Situational Awareness (SSA) efforts. Wide-field imaging systems are particularly advantageous due to their sensitivity and resolution, enabling the recovery of valuable debris characteristics that are often absent from conventional tracking databases. For example, rotational status and rotation rates – important factors for designing active debris removal missions – can be derived from analyzing lightcurves obtained through this method.

Although our study specifically focuses on data from the VLT Survey Telescope (VST) paired with the OmegaCAM [2] imager, our goal is to develop a scalable framework that can be adapted to other wide-field astronomical instruments. In doing so, we aim to demonstrate how existing astronomical datasets can be repurposed to support SSA initiatives and contribute to long-term space sustainability.

Conventional streak detection methods applied to astronomical data include approaches like the Hough transform and matched filtering. The Hough transform is particularly efficient in detecting linear features, especially when combined with pre-processing techniques such as background subtraction and edge detection to isolate streaks from stellar sources and other artifacts. Similarly, matched filtering, which compares image data with predefined streak templates, can be highly sensitive for detecting faint trails [3][4].

In recent years, machine learning and deep learning have emerged as powerful alternatives, capable of managing large-scale data and capturing complex visual patterns. Convolutional Neural Networks (CNNs) and U-Net based architectures have been broadly adopted in streak detection tasks [5], often surpassing traditional techniques in terms of accuracy and robustness. For instance, Liu et al. [6] demonstrated the effectiveness of fine-tuning domain-specific neural networks to detect satellite streaks in wide-field astronomical images. Similarly, the method proposed by Jeffries and Acuña [7] integrates feature extraction with classification to improve sensitivity while minimizing false detections. The ASTA (Automated Satellite Tracking for Astronomy) system [8] further illustrates the synergy of deep learning and computer vision, achieving impressive results in automated trail detection. Similar techniques have also been used to detect asteroid streaks in wide field-of-view images [9].

Despite these advances, several challenges persist, including reducing the rate of false positives, increasing detection capabilities for very faint streaks, and efficiently processing the ever-growing volumes of astronomical data. Combining conventional image-processing techniques with modern AI tools presents a promising strategy to overcome these limitations and advance the field.

The remainder of this paper is organized as follows: Section 2 details the methods employed in this study, including the dataset selection, detection algorithm, and subsequent steps such as photometric reduction, astrometric calibration, orbit fitting and lightcurve extraction. The same section introduces the performance metrics used to evaluate the proposed approach. Section 3 presents the results of the study, including the training and validation of the detection algorithm, its performance on real OmegaCAM data, extracted lightcurves, and the statistics of the analysis of 6 months of OmegaCAM data from 2023. Finally, Section 4 provides a discussion of the implications of the results.

2. METHODS

2.1 Dataset selection and preparation

The creation of a robust and varied dataset was a fundamental component in training and validating our machine-learning algorithm to detect streaks of satellite and space debris. For this purpose, we selected OmegaCAM, installed on the VST, due to several key advantages.

- *Public availability and extensive archive:* OmegaCAM's data is openly accessible, with an archive that spans more than a decade (since 2011), offering great opportunities to analyze the long-term evolution of space debris.
- *High sensitivity:* Coupled with the VST's 2.6-meter primary mirror, OmegaCAM's high sensitivity allows for the detection of very faint streaks that might be missed by commonly used wide-field SSA telescopes.

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

- *Superior spatial resolution:* With a 16k×16k pixels sensor and a fine pixel scale of 0.21 arcseconds per pixel, OmegaCAM enables precise reconstruction of lightcurves, an essential step toward characterizing the physical attributes of detected debris.
- *Wide field of view:* Capturing one square degree of sky per exposure, OmegaCAM can record multiple debris streaks in a single image. An example image is shown in Fig. 1.

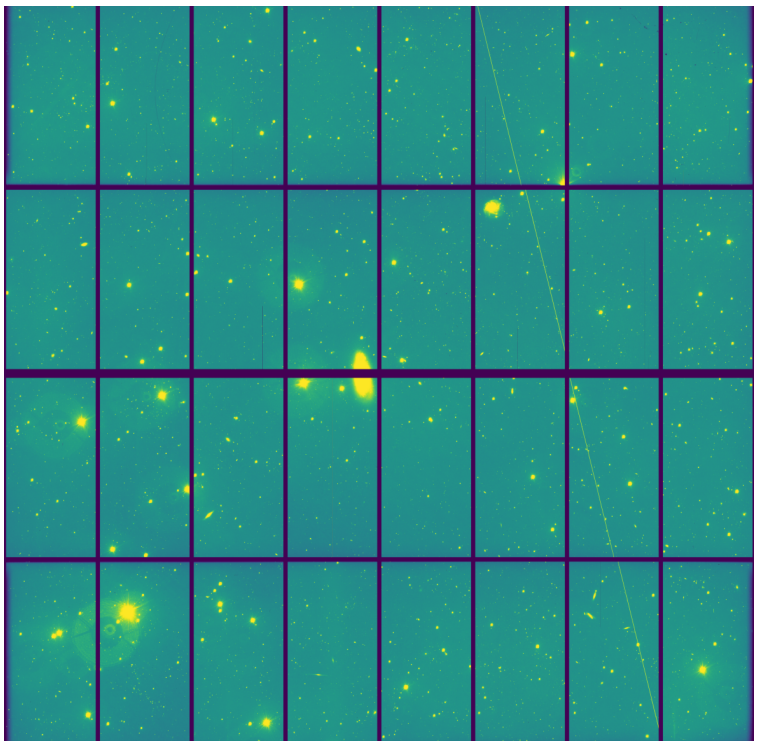


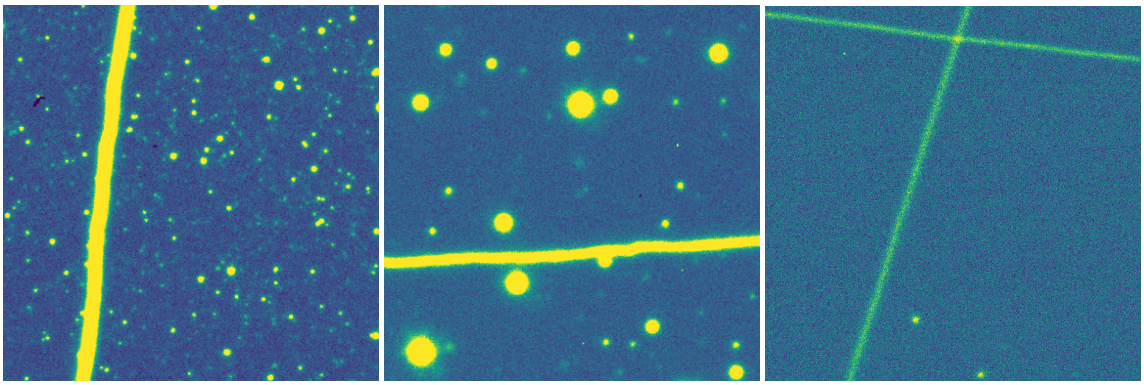
Fig. 1: OmegaCAM detector layout. A satellite streak is visible on the right side of the mosaic.

The images for the training and validation dataset were chosen from the full operational history of the OmegaCAM archive. Due to the one-year proprietary period at the time of selection, data from 2023 onward were not included. To maximize the algorithm's adaptability, we incorporated images taken with all available filters. This ensured that the model could detect streaks directly on raw frames, regardless of filter differences, without requiring additional photometric corrections. A randomly sampled set of 375 images was used to guarantee a wide range of conditions, including variations in background complexity, lunar illumination, streak morphology, stellar density, and the presence of imaging artifacts.

Due to the large dimensions of OmegaCAM's mosaic images – comprising 32 CCDs of 4096×2048 pixels each – each CCD frame was divided into smaller subcrops of 512×512 pixels. This method preserved fine structural detail and improved the model's ability to distinguish true streaks from background artifacts.

Label Studio [10] was used to manually annotate the visible streaks in the selected images. Because only 352 genuine streaks were found in the 12,000 CCD frames, synthetic streaks were introduced to augment the dataset and support a more comprehensive performance evaluation. Although this number may seem low, it reflects the reality of space debris observations during much of the archive's early coverage. The space debris population has grown considerably in recent years, which means that more recent data contain a higher frequency of debris streaks. Artificial streaks are prepared by first generating a synthetic population of space debris objects with diverse sizes and altitudes, designed to realistically represent debris types seen in actual sky images. For each object, its brightness is calculated to ensure that only realistically detectable streaks are included, and the streaks are then drawn directly onto the original images. Atmospheric effects – such as seeing and scintillation – are realistically simulated so that the artificial streaks closely match the appearance of genuine debris traces observed in astronomical images. Fig. 2 shows examples of artificial streaks, with additional information on the synthetic streak generation process available in a previous study [11].

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com



(a) Artificial streak. (b) Original streak. (c) Original and artificial streaks.

Fig. 2: Examples of artificially generated streaks. The streak in Fig. 2a shows a similar "wobble" effect as the real streak in Fig. 2b. In Fig. 2c, the horizontal streak is artificial and the vertical streak is real.

The preprocessing workflow consisted of removing prescan and overscan regions (areas of the CCD that are covered for calibration purposes and do not contain image information) from the original FITS images, adding artificial streaks and cropping the images into 512×512 pixels subframes. The resulting data were saved as 16-bit/pixel monochrome PNG files. The decision to save images as PNG files rather than using FITS directly was driven by computational efficiency. PyTorch [12] is optimized for standard image formats such as PNG, whereas FITS files resulted in slower data loading. Ensuring 16-bit depth preserved all critical information without data loss.

Of the 375 selected images, 305 were allocated for training the algorithm, corresponding to a total of 312,320 individual PNG cutouts. The remaining images were divided equally between the validation and test sets, with 35 images each, resulting in 35,840 PNG images per set.

2.2 Streaks detection

The detection of satellite and space-debris streaks in astronomical images is the first step toward the characterization and deeper study of objects in orbit. This process involves a sequence of steps and dedicated algorithms, as described below.

Machine learning algorithm The algorithm developed to detect satellite and space-debris traces is a machine learning model based on the HT-LCNN (Hough Transform Lookup-based Convolutional Neural Network), a convolutional neural network that incorporates a trainable Hough transform prior block, originally proposed by Lin et al. [13]. This integration of the Hough transform directly into the neural network enables the model to better recognize and learn geometric structures, such as lines, which are characteristic of satellite streaks in astronomical images. As the Hough transform is a widely used technique for line detection, particularly in the context of satellite trail identification, its inclusion makes the model's particularly relevant for this specific application. The original HT-LCNN framework was trained and evaluated on the Wireframe [14] and York Urban [15] datasets, both of which focus on edge detection in architectural and indoor scenes. Designed primarily for applications such as automatic device orientation, the model required adaptation to be effective in astronomical imaging. One key modification was enabling the model to process 16-bit files, which are standard in astronomy and provide higher dynamic range than the 8-bit images typically used in conventional machine learning tasks. This capability is essential for ensuring that no information is lost and for allowing the algorithm to exploit the streak profiles with the highest possible precision.

Line stitching To maintain the original image resolution and avoid the effects of harsh resizing during machine learning inference, we divide each image into 512×512 pixels subcrops. While this approach enhances detail sensitivity during streak detection, working at such a small scale is impractical for later steps like orbit fitting and light curve extraction. Therefore, after initial line detection, we reconstruct the streaks at the scale of the entire CCD tile using a structured stitching process:

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

1. *Initial clean-up:* Within each subcrop, duplicate line detections are eliminated by comparing endpoint positions and orientations. Line segments shorter than 5 pixels are also discarded.

2. *Coordinate reprojection:* Remaining line coordinates are transformed back to their corresponding positions within the full CCD tile.

3. *Hough transformation:* For each line, Hough parameters ρ (distance from the origin to the closest point on the line) and θ (angle between the X-axis and the line) are computed to represent their geometric properties.

4. *Clustering:* Lines are grouped based on similar Hough parameters, using thresholds of 50 pixels for ρ and 0.1 radians for θ, effectively joining fragments that belong to the same physical streak while minimizing incorrect matches.

5. *Streak assembly:* Within each cluster, the most distant endpoints are selected to define the new, stitched streak segment.

Because the model processes small crops, it achieves higher sensitivity to subtle streaks and endpoints but may also produce more false positives – features that look like streaks in isolation but not in the larger context. The stitching step enables us to efficiently filter out these false detections by discarding singleton segments – lines that appear in only one crop and do not cluster with any others – since these are unlikely to represent true streaks. Although this could occasionally exclude short streaks located at the very corner of an image, we accept this trade-off: genuine features will typically appear in adjacent CCDs, allowing us to reconstruct them from neighboring data, while those that cannot be recovered are typically too short to contribute meaningfully to subsequent analyses.

Detection refinement The neural network produces a bitmap in which each pixel represents the likelihood of being part of a streak. Directly extracting the start and end points from this output may not yield positions that precisely align with the observed streak in the image. To enable accurate astrometric measurements, the detected streak must first be aligned with the actual observation. This is done by generating a cutout around the detection and rotating it so that the streak is approximately horizontal. The aligned streak is then divided into equal-length segments, each of which is horizontally median combined to produce cross-sectional profiles along the streak. These profiles are individually fitted to a Moffat distribution:

$$f(x) = A\left(1 + \frac{(x-x_0)^2}{\gamma^2}\right)^{-\alpha} \tag{1}$$

In this function, A defines the amplitude, x_0 marks the position of the peak, γ sets the width of the core, and α determines the shape index of the profile. We chose a Moffat fit rather than a Gaussian because the Moffat model more accurately describes the typical PSF (Point Spread Function) under observing conditions dominated by atmospheric seeing [16]. The best-fit value of x_0 for each segment gives the central position of the streak for that particular slice. These calculated centers are collectively fitted with a straight line, and the initially detected endpoints are projected onto this line to align the streak with the actual observation. The entire procedure is repeated until the slope of the fitted line is nearly zero (we set a threshold at 0.1 divided by the pixel length of the streak), ensuring proper horizontal alignment. For the segment width (along-track direction), we use 1.4 times the FWHM (Full Width at Half Maximum) of the Gaussian PSF measured from the guide stars on the auxiliary OmegaCAM CCDs. The height of the segment is initially set to allow for a maximum difference between the detected and actual streak of 3°, and subsequently reduced during the iterative process in order to avoid unnecessary large segments that potentially contain background stars, which would confuse the Moffat fitting. After alignment, the final endpoints are mapped back to the original image coordinates. In cases where a streak is very faint and has a low SNR (Signal-to-Noise Ratio), broader segments may be necessary for a robust Moffat fit. Since brightness is unknown at this stage, we progressively widen the segment until a valid solution is obtained or until the width exceeds a threshold (20 divided by the length of the streak in pixels). If no fit is found within this limit, we assume that the detected streak is too faint for reliable processing or potentially a false positive.

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

2.3 Photometric reduction

To analyze the detected space objects, we developed an image processing pipeline that extracts both astrometric and photometric information. Astrometric measurements are used to match the observed streaks with known satellites, while photometric lightcurves provide insight into the tumbling behavior of the objects.

Astrometric calibration The raw OmegaCAM images are delivered with an existing astrometric solution. However, this initial calibration is only approximate and must be refined for high-precision astrometric applications. To improve accuracy, we use a local instance of `Astrometry.net` [17] to re-calibrate the images (plate solving), resulting in significantly more accurate celestial coordinate mappings. This refined calibration is then applied to derive precise coordinates for the start and end points of the detected streaks.

Object correlation The detected streaks are then matched with known satellites and debris objects listed in publicly available catalogs, specifically `space-track.org`[1]. To do this, Two-Line Elements (TLE) closest in time to each observation are propagated to the image timestamps to estimate which objects may have crossed the detector during acquisition. These candidates are then projected into the image and compared with the detected streaks. Because TLEs can carry positional inaccuracies, candidate matches are filtered based on thresholds in both angular separation and position angle difference between prediction and observation.

Orbit fitting To better align catalog predictions with observed data, orbit refinement is performed using the batch least squares fitting method from the `Orekit` orbital dynamics library [18]. This process requires accurate timestamps for both ends of each streak. When the start and/or end of the streak is observed within the image frame, these times are determined using the exposure-time information of the image, with pixel-level corrections applied to account for the OmegaCAM shutter.

If a streak spans the full CCD and both endpoints fall outside the frame, exact entry and exit times are not directly observable. In such cases, we determine the time of minimal angular separation between the projected TLE and the observed entry/exit positions. This provides an initial timestamp for orbit fitting. We then refine the orbit by searching within a small time window around this initial guess, identifying the orbit that minimizes the Mahalanobis distance [19] between the observed and expected positions. From this fitted orbit, we derive timestamps for both endpoints.

Lightcurve extraction Once both the celestial coordinates and the observation times for the ends of the streak are known, we extract the lightcurve of the object. A series of rectangular apertures is placed along the streak, and the flux within each aperture is measured. The time an object spends within each aperture is calculated from the fitted orbit solution and used to normalize the flux values for the relative magnitude computation with field stars.

The seeing effect introduces disturbances that result in the streaks deviating from a straight line. Given OmegaCAM's fine pixel scale (0.21 arcsec/pixel) and the fact that the seeing FWHM exceeds 0.8 arcsec for half of all exposures, the seeing is strongly over-sampled, resulting in deviations that can span several pixels from a straight line fitted through the streak. To avoid incomplete aperture coverage or the use of too large apertures, we precisely center each aperture on the path of the streak. This is achieved, in a similar way as for the streak alignment previously performed, by rotating a cutout of the streak to horizontally align it and segmenting the streak. Each segment is median-combined along the horizontal axis to produce cross-sectional profiles, which are then fitted to a Moffat model. The width and height of each section are defined as 1.4 times the FWHM of the Gaussian PSF measured from the auxiliary OmegaCAM CCDs, and this value is also applied for the width of the rectangular photometry apertures. The centers of the median-combined cross-sectional profiles are used to establish the optimal positions for the measurement apertures. Any outliers in the Moffat fit parameters are identified and excluded, ensuring that only apertures based on reliable fits contribute to the photometric measurements. The height (across-track direction) of the rectangular aperture for each segment is set to 1.4 times the FWHM of the Moffat fit for that segment. For rotation, we employ `scikit-image` [20] routines, but because their spline interpolation can slightly alter flux values in the rotated image, we transform the determined aperture centers back onto the original unrotated image before performing photometry with `photutils` [21]. We observed that the interpolation involved in the rotation process introduces a sub-pixel positional uncertainty

[1] https://www.space-track.org/

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

(approximately 0.5 pixels) for the final aperture placements. To account for this, we increase the aperture height by 1 pixel to ensure accurate flux measurement.

2.4 Archive analysis pipeline and repository

Fig. 3 provides an overview of the entire processing pipeline, illustrating the flow of all the steps described previously. The final results – including an image snippet of the detected streak, orbital parameters derived from the fitted TLE, and the corresponding extracted lightcurve – are compiled and stored in our internal repository, ODLI (Orbital Debris Lightcurve Inventory). Our long-term objective is to make ODLI publicly accessible, providing a valuable resource for researchers studying satellites and space debris. An example entry from the repository is presented in Fig. 4.

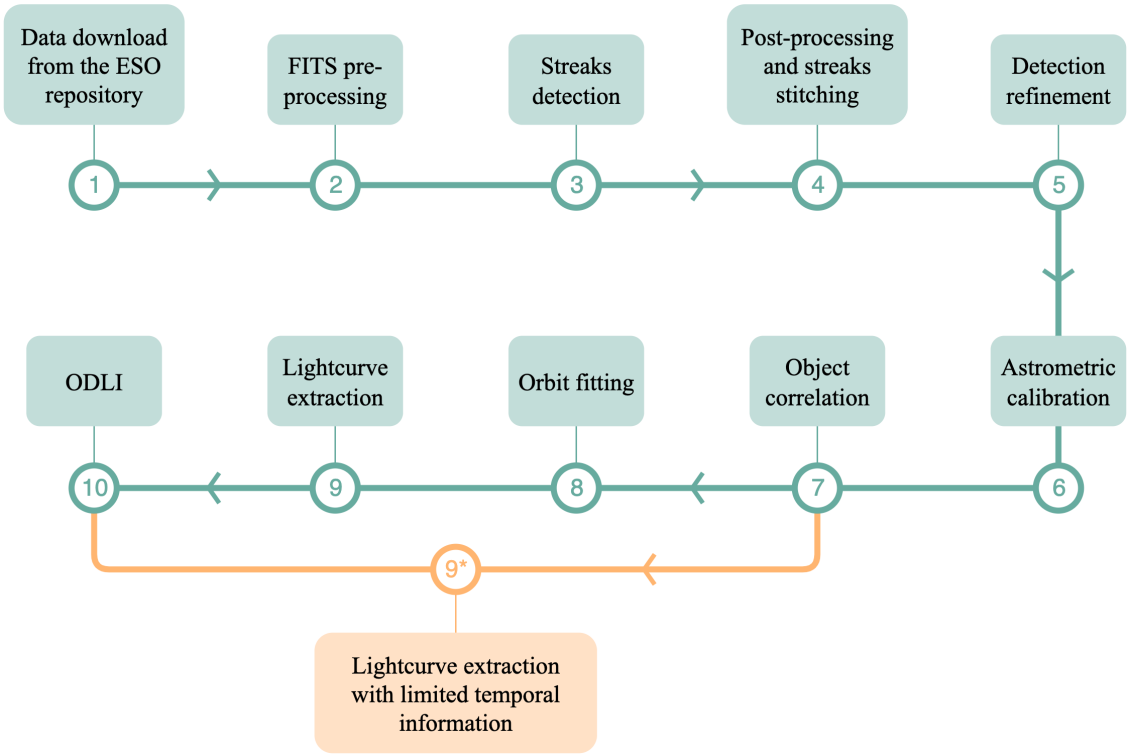


Fig. 3: Archive analysis pipeline. 1) Raw data is downloaded locally from the ESO archive. 2) Original FITS files are pre-processed to match the input format of the streak detection algorithm. 3) The trained machine learning model is applied to identify streaks. 4) Detected streaks are cleaned and stitched across CCDs. 5) Horizontal alignment is applied to refine detection accuracy. 6) `Astrometry.net` is used to improve the astrometric solution. 7) Detected streaks are matched with space-object catalogs. 8) TLE orbits are fitted to the matched object streak. 9) Photometric lightcurves are extracted with precise time stamps. 9⋆) If no catalog match is found (and relative velocity cannot be determined), lightcurves are extracted using relative time stamps only. 10) Final lightcurve and orbit data are stored in ODLI.

2.5 Performance metrics

In order to evaluate the results of the machine learning algorithm on the validation and test sets, several metrics commonly adopted in the field are used and presented below.

Precision Precision is an indication of how many cases detected as positive are actually true positives. It tells us how good the algorithm is at making correct predictions when something is detected, and is calculated as follows:

$$\text{Precision} = \frac{TP\ (\text{True Positive})}{TP + FP\ (\text{False Positive})} \tag{2}$$

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

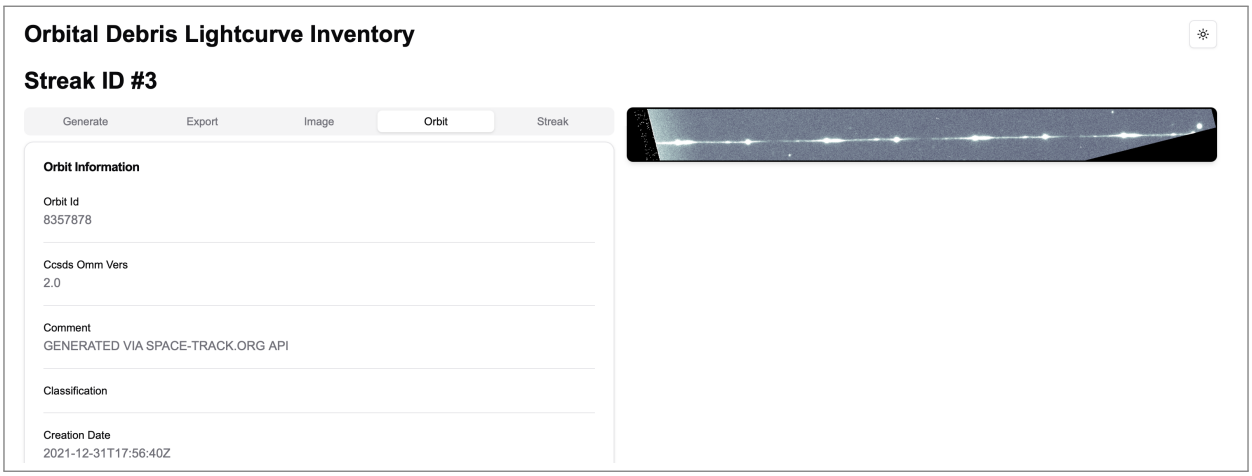


Fig. 4: Sample entry in ODLI.

Recall Recall gives an indication of how many positive detections, whether true or false, have been made by the model. It indicates the ability of the model to make positive detections. The formulation is as follows:

$$Recall = \frac{TP}{TP + FN \text{ (False Negative)}} \tag{3}$$

F1-score The F1-score is the harmonic mean of precision and recall, offering a balanced measure of a model's performance by equally weighting both metrics. It is especially useful in situations where achieving perfect precision and recall is unrealistic. By combining these two metrics, the F1-score provides a comprehensive evaluation of a model's ability to make accurate predictions while effectively managing trade-offs between false positives and false negatives. The formulation is as follows:

$$F1 = 2 \cdot \frac{Precision \cdot Recall}{Precision + Recall} \tag{4}$$

Precision-Recall curve and Average Precision The precision-recall curve is a graphical tool used to evaluate the performance of detection algorithms, particularly when working with imbalanced data where true positives are relatively rare. In our case, the detection algorithm assigns a confidence score to each candidate detection, indicating how likely it is that the identified feature corresponds to a real streak. The precision-recall curve plots precision against recall as the confidence threshold is varied. This allows us to visualize how the trade-off between detecting more streaks (recall) and maintaining accuracy (precision) changes depending on how strict or lenient we are with the confidence threshold.

Average precision (AP) summarizes the information encapsulated by the precision-recall curve into a single score. It is calculated as the area under the precision-recall curve and represents the model's ability to maintain high precision at increasing levels of recall. A higher AP indicates that the model consistently achieves high precision and high recall.

3. RESULTS

3.1 Training of the detection algorithm

The detection algorithm was trained using the Adam optimizer with an initial learning rate of 1×10^{-4}, a value determined by a dedicated learning-rate search described in detail in a previous study [11]. To mitigate potential convergence issues, we adopted the AMSGrad variant of Adam and applied L2 regularization with a strength of 2×10^{-5}. Training was carried out over 4 epochs, taking advantage of the large size of the data set, with the learning

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

rate scheduled to update every 2 epochs. The progression of training and validation loss throughout the process is presented in Fig. 5.

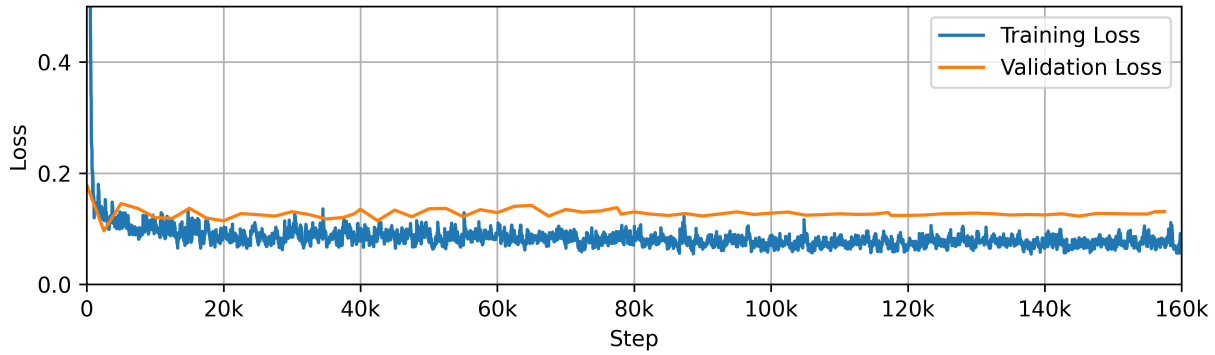
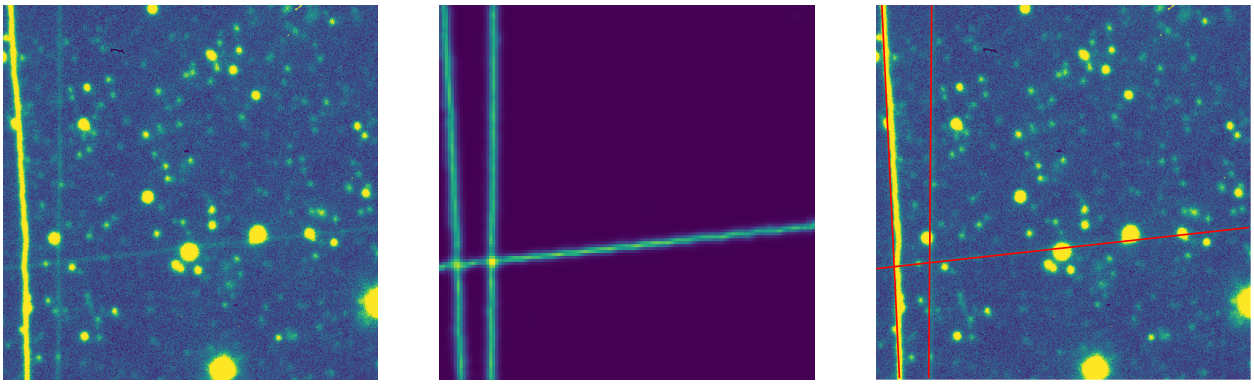


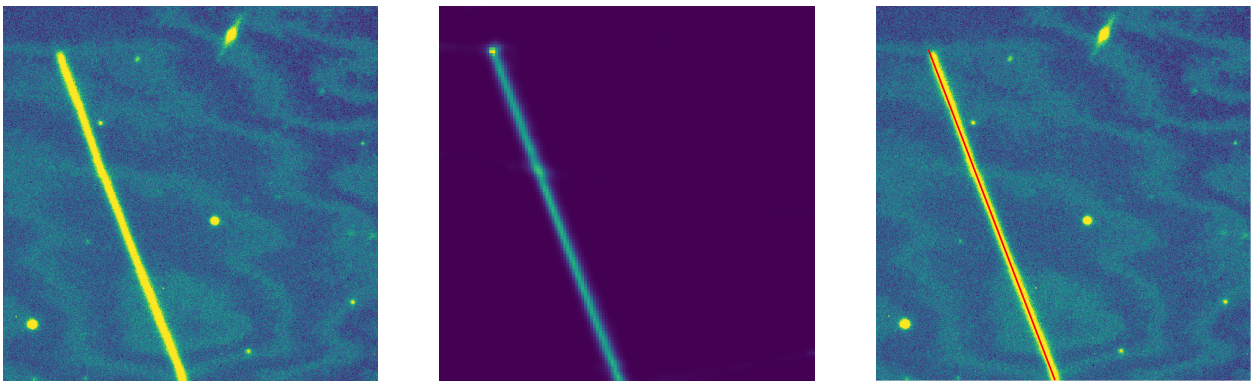
Fig. 5: Training and validation loss.

Fig. 6 to 8 present several examples of accurately detected streaks, illustrating the typical output produced by the algorithm.



(a) Original image. (b) Detection heatmap. (c) Detected lines.

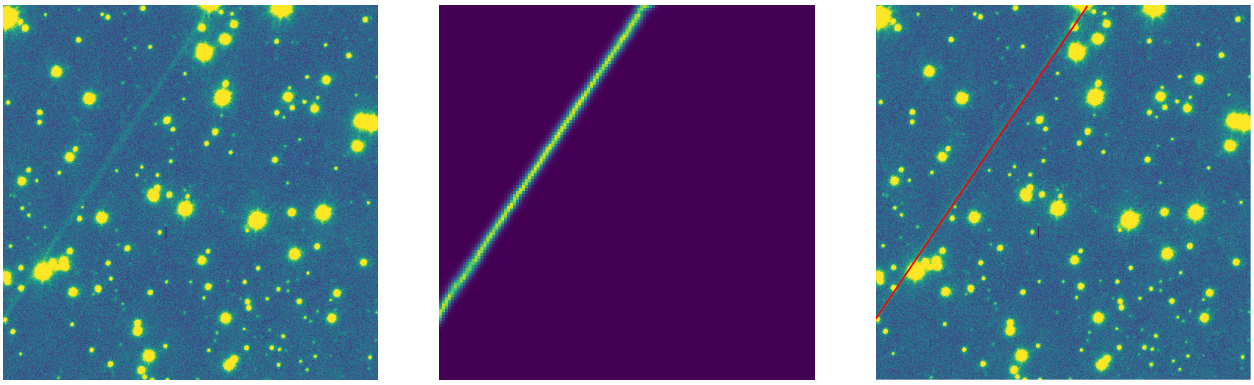
Fig. 6: Detection of 3 individual objects in a single image. The brightest streak is one of a slow object, with visible characteristic wobble.



(a) Original image. (b) Detection heatmap. (c) Detected lines.

Fig. 7: Detection of an object with the end of the streak within the image boundaries.

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com



(a) Original image. (b) Detection heatmap. (c) Detected lines.

Fig. 8: Detection of an object with low SNR (SNR 3.3) in a crowded star field.

3.2 Performance analysis

To evaluate the performance of the detection algorithm, we used the validation dataset comprising 35,840 PNG images. During the assessment, we identified one mosaic with tracking issues, where all stars appeared as elongated trails rather than point sources. To prevent this anomaly from skewing the results, we excluded this image from the analysis. The initial detections performed on the 512×512 pixels sub-crops yielded the following metrics:

Table 1: Performance of the detection algorithm on the validation dataset, crop-wise.

Precision	Recall	F1-score
0.944	0.939	0.947

Although these baseline scores are promising, it is evident that some false positives and false negatives remain. As discussed in Section 2.2, we address this issue by stitching detected streaks across the entire CCD. This process eliminates many false positives by discarding detections present in only one crop and also helps recover some false negatives, as the stitching can reconnect small segments that might otherwise be missed, especially in image corners where short lines are harder to detect. After applying this CCD-wise stitching, our performance metrics improve to:

Table 2: Performance of the detection algorithm on the validation dataset, after CCD-wise stitching.

Precision	Recall	F1-score
0.971	0.961	0.966

These enhanced scores demonstrate a marked improvement, underscoring the effectiveness of our stitching approach.

The precision-recall curve for the validation set (evaluated crop-wise) illustrated in Fig. 9 further highlights the robustness of the algorithm, with both precision and recall remaining consistently high and well-balanced. There is no observable bias toward over- or under-detection. On the validation set, the algorithm achieves an impressive average precision of 0.985. We have investigated possible detection threshold values to optimize performance; however, in light of the shape of the precision–recall curve and in order to avoid discarding any potential true positives, we chose not to apply a fixed threshold and instead retained all candidate detections. We acknowledge that this may result in a higher number of false positives, but given the exploratory nature of the study and our ongoing evaluation of the algorithm's capabilities across archival data, we adopted this intentionally conservative approach to maximize completeness.

In terms of detectability with respect to SNR, Fig. 10 illustrates the distribution of true detections and false negatives (i.e. missed streaks) for the artificial streaks – where the SNR is known – plotted against streak length. More than 95% of the streaks with SNR values greater than 4 were successfully detected by the algorithm. For streaks with SNR between 2 and 4, detection performance was highly dependent on background conditions, such as image noise and field complexity, but still achieved a detection rate exceeding 90%.

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

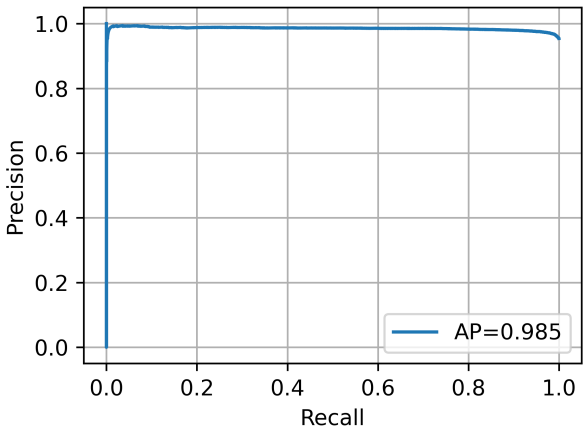


Fig. 9: Precision-Recall curve.

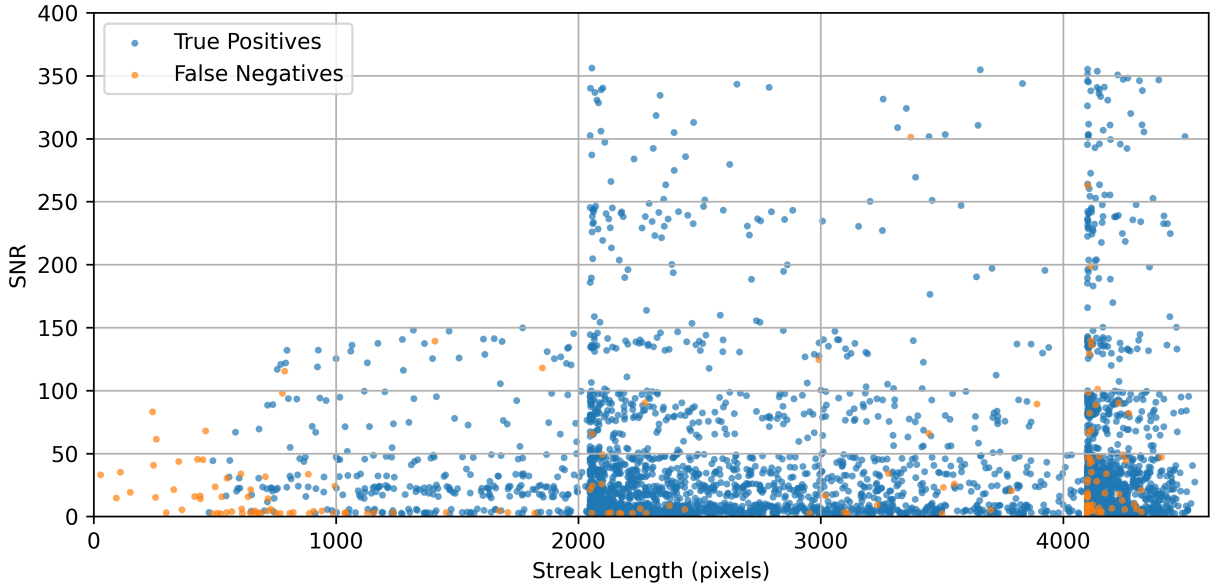


Fig. 10: Detection performance for artificial streaks in the validation dataset. The cluster of false negatives in the lower-left region of the graph results from the line stitching process, which discards true positive detections shorter than the sub-crop size. Additional false negative clusters near the full image dimensions (2048 pixels and 4100 pixels) are primarily caused by streaks positioned too close to the image edges, making them harder for the algorithm to detect reliably.

3.3 Generalization to real-world data

While validation results provide valuable insight, the true performance of a detection algorithm can only be accurately assessed on real-world data. For this reason, we set aside a dedicated test set consisting of 35 OmegaCAM mosaics – matching the size of the validation set, as described in Section 2.1 – but this time containing only real streaks, with no artificial data included. After applying the CCD-wise line stitching procedure, the algorithm achieved the following performance on this test set:

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

Table 3: Performance of the detection algorithm on the test dataset, CCD-wise.

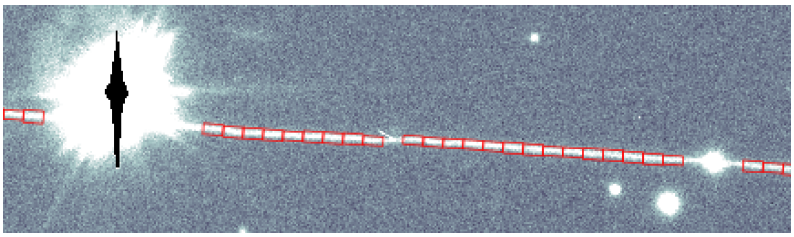
Precision	Recall	F1-score
1.0	0.744	0.853

Although these results, particularly the high precision, are encouraging, they must be interpreted with caution. The test set contained only 39 real streaks, which limits the statistical robustness of the performance metrics. To supplement this and gain a more reliable understanding of the behavior of the model in real world conditions, we ran the detection pipeline on all OmegaCAM data from January 2023, a total of 2,866 full-mosaic images.

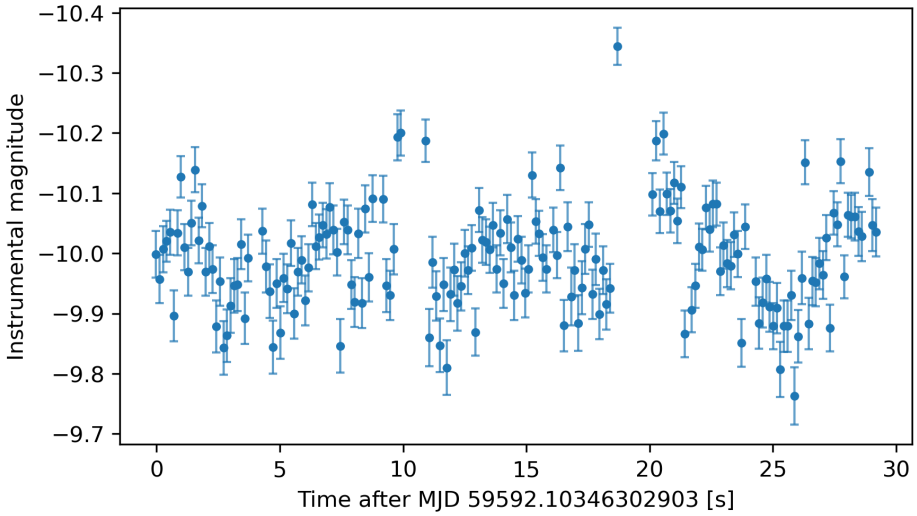
Due to the scale of this dataset, full annotation was not feasible. Instead, we focused on estimating the false positive rate by manually reviewing the detections and filtering out erroneous results. Based on more than 3,000 detected streaks, this large-scale run yielded a precision of 0.63. The most frequent sources of false positives were identified as diffraction spikes from bright stars, saturation bleed from overexposed sources, and fringe patterns from imperfect filters. These types of false positives are common and have already been addressed through various strategies, including sub-crop analysis and the line stitching procedure. However, the drop in precision in this broader dataset reflects the immense diversity of OmegaCAM images. Despite deliberate efforts to construct a representative and diverse training set, it remains extremely difficult to capture the full range of features and artifacts that may appear in real observational data.

3.4 Lightcurve extraction

Once the streaks are detected, we apply our lightcurve extraction pipeline. Fig. 11 illustrates the placement of rectangular apertures along a satellite streak, as well as the rejection of contaminated apertures through outlier removal based on the distribution of Moffat model parameters. The resulting raw lightcurve for the entire streak is also shown.



(a) Placement of the apertures along the streak.



(b) Resulting raw lightcurve.

Fig. 11: Aperture photometry for lightcurve extraction.

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

The object was identified as a spent rocket body (NORAD ID 28704) close to GEO (Geostationary orbit). It was detected in seven tiles of a single exposure, resulting in seven individual lightcurves extracted from that image. By combining all lightcurves and applying the FALC algorithm [22] for frequency analysis, implemented in WinFALC [23], we identified the dominant rotation period of approximately 7 s. Fig. 12 shows all lightcurves folded to this period. This example demonstrates that when the observational timespan significantly exceeds the rotation period (just over 200 s versus 7 s in this case), a reliable period determination is achievable even when the photometric uncertainty approaches the lightcurve amplitude. More detailed illustrations of the streak alignment, orbit fitting process, and lightcurve extraction method can be found in a prior study [24].

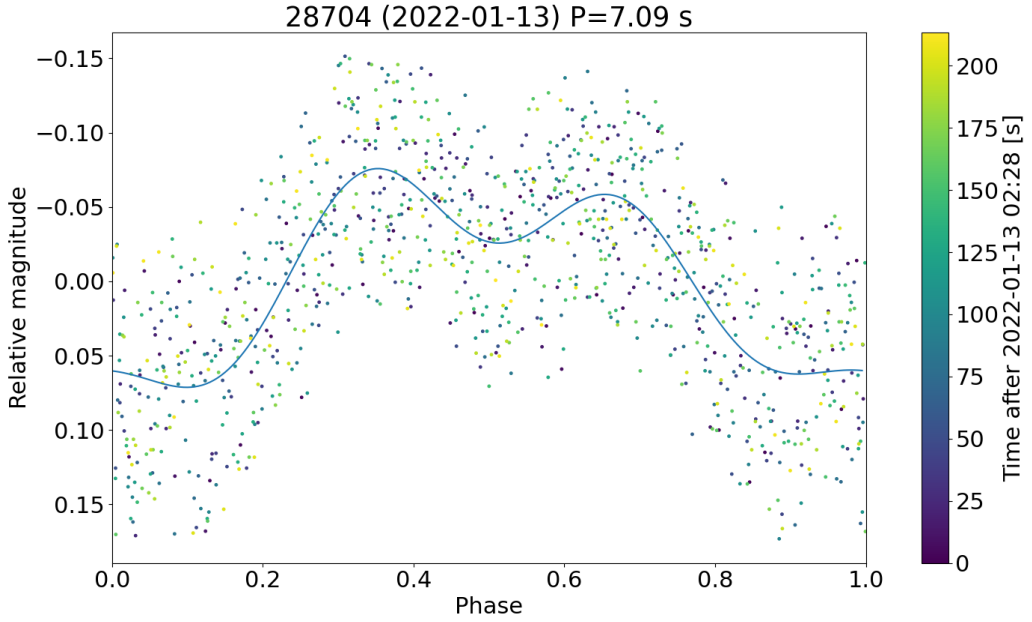


Fig. 12: All lightcurves from seven streaks caused by the same object (NORAD ID 28704) in a single exposure, folded to the shortest identified period of 7.09 s, color coded by the actual observation time of each individual measurement.

3.5 Results of the analysis of 6 months of archival data

Now that the various analysis methods have been thoroughly evaluated, we can move on to the most compelling part of this study: the analysis of archival data. We applied our detection algorithm to the first six months of OmegaCAM data from 2023, and the results are presented below. Fig. 13 shows examples of detected streaks, showcasing the remarkable diversity present in the OmegaCAM archive.

January 2023 As discussed in Section 3.3, we conducted a manual review of all detections for the month of January, classifying each detection as a true positive or a false positive. This manual validation step allowed us to better understand the proportion of true detections that could not be matched to any known object in the `space-track` catalog. The results were as follows:

Table 4: Identification of January 2023 detections through correlation with the `space-track` catalog.

TP detections	Correlated with cataloged objects	Uncorrelated
1912	1439	473

These figures indicate that approximately 25% of the verified detections could not be matched to any known satellite or debris object. After accounting for possible miscorrelations – for example, those due to very short streaks or FP detections strongly resembling genuine streaks that we might have missed – we estimate that around 20% of the true streaks found in the OmegaCAM archival data correspond to uncataloged or unidentified objects.

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

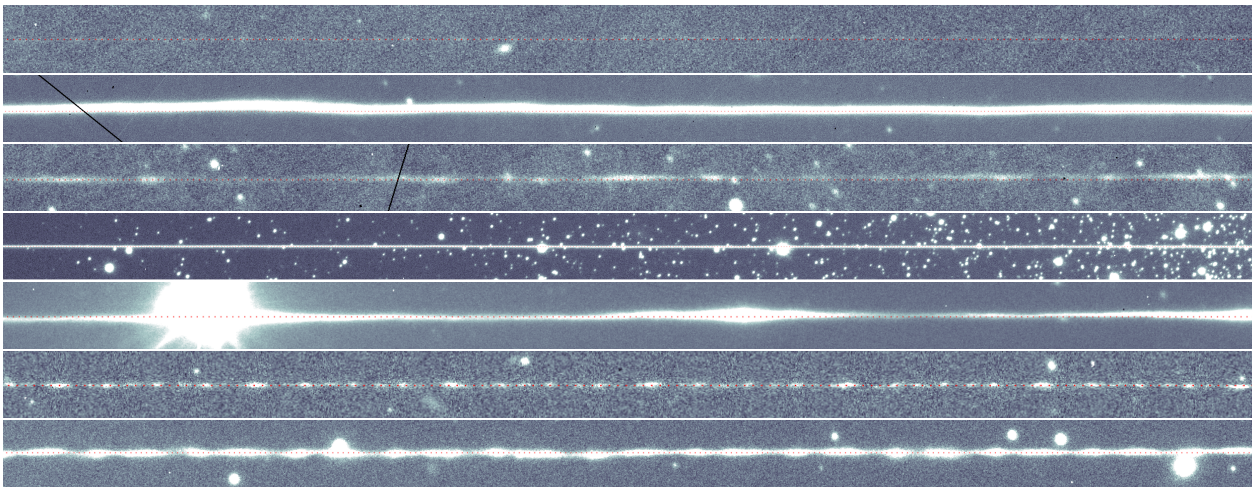


Fig. 13: Sample detections from the OmegaCAM archive. All streaks were correlated with known objects.

February to June 2023 For the subsequent months, from February through June, we did not manually validate the detections to distinguish true positives from false positives. Therefore, in this portion of the analysis, we focus solely on correlated detections, matched with known objects in the `space-track` catalog.

Table 5 presents an overview of the key results of the detection and correlation analysis of satellite streaks. During the observation period, a total of 1,415 unique objects were identified. Of the 750 detected payloads, 190 were attributed to Starlink satellites and 120 to OneWeb satellites. As anticipated, the majority of detections correspond to objects in LEO (Low Earth Orbit). However, a non-negligible fraction of the detections was associated with MEO (Medium Earth Orbit) and GEO objects. The farthest object to be observed was a Breeze-M rocket body, at an altitude of 66,057 km at the time of observation. In terms of physical characteristics, the population of detected objects was dominated by satellites with large radar cross-sections (RCS), consistent with the inherent difficulty of detecting smaller objects. Although direct size information is not available for cataloged objects, we estimate that the smallest detected objects in LEO have dimensions on the order of a few centimeters, based on their RCS values that are smaller than 0.1 m^2. The temporal range of detections is broad: the oldest object was a Vanguard satellite launched in 1959, whereas the most recent detection corresponds to a rocket body launched in May 2023 and observed less than two weeks post-launch.

Fig. 14 illustrates the pointing directions of all images acquired with OmegaCAM between January and June 2023, plotted in altitude-azimuth polar coordinates. The figure highlights all images containing at least one satellite streak. The distribution reveals that streak contamination occurred across nearly all surveyed regions of the sky. During the six-month interval, 9.9% of the acquired images were affected by at least one streak.

Table 5: Summary from the OmegaCAM archive analysis from January to June 2023.

Number of images (mosaic)	19045
Number of streaks (mosaic-wise)	2254
Unique objects	1415
Object types	Payloads: 750 | Rocket bodies: 391 | Debris: 274
Orbit types[a]	LEO: 585 | MEO: 247 | GEO: 108 | HEO: 19 | Other: 456
RCS sizes[b]	Large: 1021 | Medium: 293 | Small: 66 | Unknown: 35

[a] LEO (Low Earth Orbit): perigee and apogee < 2000 km
MEO (Medium Earth Orbit): 2000 km $\leq$ perigee < 35586 km and 2000 km $\leq$ apogee < 35586 km
GEO (Geostationary Orbit): perigee ≥ 35000 km, apogee ≤ 36500 km, inclination $\leq 10°$
HEO (Highly Eccentric Orbit): perigee < 2000 km and apogee > 36500 km

[b] Large: RCS ≥ 1 m^2
Medium: 0.1 $m^2 \leq$ RCS < 1 m^2
Small: RCS < 0.1 m^2

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

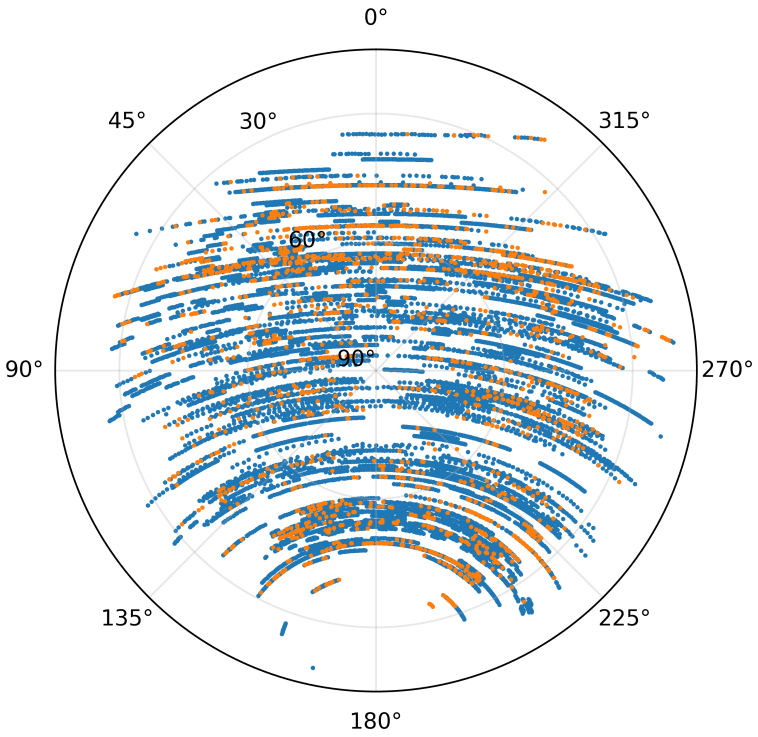


Fig. 14: Telescope pointing directions for OmegaCAM images acquired between January and June 2023. Orange points indicate images containing at least one detected and correlated streak. The concentric circles represent the altitude in degrees, while the radial axis shows the azimuth in degrees.

4. DISCUSSION

The results of this study are highly promising. Our detection algorithm has shown the ability to identify a wide range of streaks, including very faint ones. However, the less favorable results regarding the algorithm's generalization to real-world data illuminate an important challenge when working with large datasets: even with careful selection of training and validation datasets – incorporating diverse backgrounds, moon illumination conditions, seeing variations, and artifacts – by selecting images throughout the archive and from all filters, the images in an archive such as OmegaCAM, which spans more than 14 years of data acquisition, exhibit such extensive variation that we cannot incorporate all conditions the algorithm will encounter when processing the data. One potential solution for creating a more diverse training and validation dataset would be to train on portions of mosaics rather than complete ones, using perhaps a few CCDs from many more images. This approach would provide greater image variety while maintaining a reasonable number of training images, representing an important trade-off since simply incorporating more data would make training extremely slow and computationally intensive. Another possible approach would be to use the correctly detected and correlated streaks as the basis for a new dataset, as this would include far more examples than our initial set. However, caution is needed, since this method could reinforce any biases already present in the algorithm.

As demonstrated in our analysis of the January 2023 data, the current false-positive rate is too high, and reducing it is a critical priority. Reducing false positives is crucial because we have identified many genuine streaks that cannot be correlated with existing catalogs, and we must separate these from false positives to conduct statistically significant studies. An approach we have begun exploring involves using a classifier to distinguish false positives from true positives after the detection algorithm completes its analysis. We have tested a classifier consisting of a simple convolutional neural network that performs binary classification (true positive or false positive) using our manually sorted January 2023 data results as the training dataset. Although more work is certainly needed, we have already achieved accuracy scores that exceed 80%. As mentioned above, simply expanding the training dataset would be computationally intensive and resource-consuming, and dataset annotation is extremely tedious. Additionally, we

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

have observed that false positives can be easily categorized into common types: diffraction spikes from bright stars, pixel bleeding from bright stars, and filter fringes. Because of this classification, we believe that focusing on post-processing will yield better results and prove more efficient than modifying the base dataset itself.

Our estimate that a significant fraction (20% of the streaks) in the OmegaCAM archive cannot be correlated with existing satellites and space-debris catalogs provides an exceptional opportunity to study this untracked population. We believe most untracked streaks originate from small debris in LEO, a particularly concerning category of debris. Obtaining information about this population, including debris density, rotational state, and rates, would be extremely valuable and necessary. Since these objects cannot be linked to existing catalogs, we lack velocity information, which prevents us from obtaining the temporal data required to construct lightcurves. Although we can still recover the shape of the lightcurve (which tells us if an object is rotating), we cannot determine the actual frequency of rotation or combine the lightcurves of a single object observed at different epochs (i.e. detected in different exposures). To address this limitation, we are developing methods that extract approximate speed information from the visible "wobble" of streaks seen in images. Since this wobble depends on the object's velocity across the detector and the seeing, we hope to recover a rough speed estimate that will allow us to estimate its altitude (assuming a circular orbit). Initial tests of a machine-learning algorithm for this purpose have been performed on artificial streaks and show promising results, though the lack of real data has been a limiting factor. With real data analysis currently underway, we anticipate being able to test and improve our algorithms soon.

On the topic of lightcurve extraction, though our current results are promising, we are continuing to improve the fitting of TLEs to our detected streaks, particularly for streaks crossing entire detectors, where we inherently lack precise observation time since no streak endpoints were recorded. Currently, a prototype for the whole process, starting from the raw detection to the lightcurve, is available. In order to make sure all kinds of different streaks are handled correctly, extensive testing and fine tuning of the parameters are necessary. Period determination currently still involves manual intervention and needs to be automated in order to process the large number of lightcurves.

Once our in-house repository, ODLI, is populated with detections from the OmegaCAM archive and becomes operational, the community will have access to valuable data to perform further analysis. In parallel, we are actively working toward adapting our methods to other telescope archives, such as the DECam archive and, in the near future, data from the Vera C. Rubin Observatory. We are also collaborating closely with the IAU Centre for the Protection of the Dark and Quiet Sky from Satellite Constellation Interference (IAU CPS) to develop compatible pipelines and connect our database to theirs. In this era of increasing risks related to space debris and enormous challenges in tracking everything in orbit above us, open data exchange is not just beneficial – it's essential.

REFERENCES

[1] Scientific Committee on the Peaceful Uses of Outer Space and Technical Subcommittee. IADC report on the status of the space debris environment. Technical report, Inter-Agency Space Debris Coordination Committee, 2025.

[2] K. Kuijken. OmegaCAM: ESO's newest imager. *The Messenger*, 146:8, 2011.

[3] P. Hickson. A fast algorithm for the detection of faint orbital debris tracks in optical images. *Advances in Space Research*, 62(11):3078–3085, 2018.

[4] G. Nir, B. Zackay, and E. O. Ofek. Optimal and efficient streak detection in astronomical images. *The Astronomical Journal*, 156(5):229, 2018.

[5] G. Yang, Y. Xianlong, X. Yao, Z. Zhengxu, Y. Xu, and D. Chenggang. Enhanced YOLOv8-based method for space debris detection using cross-scale feature fusion. *Discover Applied Sciences*, 2025.

[6] L. Liu, S. Rongyu, and Z. He. Identification of megaconstellations in wide-field astronomical images with machine learning. *The Astronomical Journal*, 169(4):205, 2025.

[7] C. Jeffries and R. Acuña. Detection of streaks in astronomical images using machine learning. *Journal of Artificial Intelligence and Technology*, 4(1):1–8, 2023.

[8] F. Stoppa, P. J. Groot, R. Stuik, P. Vreeswijk, S. Bloemen, D. L. A. Pieterse, and P. A. Woudt. Automated detection of satellite trails in ground-based observations using U-Net and Hough transform. *A&A*, 692:A199, 2024.

[9] B. Y. Irureta-Goyena, E. Rachith, S. Hellmich, J.-P. Kneib, B. Altieri, C. Lemon, T. Saifollahi, O. Hainaut, W. Freudling, F. Dux, M. Micheli, F. Ocaña, P. Ramírez-Moreta, F. Courbin, L. Conversi, M. Millon, G. Ver-

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

does Kleijn, and M. Salzmann. A method for asteroid detection using convolutional neural networks on VST images. *Astronomy & Astrophysics*, 694:A49, 2025.

[10] M. Tkachenko, M. Malyuk, A. Holmanyuk, and N. Liubimov. Label Studio: Data labeling software, 2020-2022. Open source software available from https://github.com/heartexlabs/label-studio.

[11] E. Rachith, S. Hellmich, and J.-P. Kneib. Detection and characterization of space debris using VST/OmegaCAM archival data. In *9th European Conference on Space Debris*, volume 9. ESA Space Debris Office, 2025.

[12] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.

[13] Y. Lin, S. L. Pintea, and J. C. van Gemert. Deep Hough-transform line priors. In *Computer Vision – ECCV 2020*, pages 323–340. Springer International Publishing, 2020.

[14] K. Huang, Y. Wang, Z. Zhou, T. Ding, S. Gao, and Y. Ma. Learning to parse wireframes in images of man-made environments. In *Computer Vision – ECCV 2018*, pages 626–635, 2018.

[15] P. Denis, J. H. Elder, and F. J. Estrada. Efficient edge-based methods for estimating Manhattan frames in urban imagery. In *Computer Vision – ECCV 2008*, pages 197–210. Springer Berlin Heidelberg, 2008.

[16] I. Trujillo, J. A. L. Aguerri, J. Cepa, and C. M. Gutiérrez. The effects of seeing on Sérsic profiles – II. the Moffat PSF. *Monthly Notices of the Royal Astronomical Society*, 328(3):977–985, 2001.

[17] D. Lang, D. W. Hogg, K. Mierle, M. Blanton, and S. Roweis. Astrometry.net: blind astrometric calibration of arbitrary astronomical images. *The Astronomical Journal*, 139(5):1782–1800, 2010. Open source software available from https://astrometry.net.

[18] B. Cazabonne and P. Cefola. Towards accurate orbit determination using semi-analytical satellite theory. In *31st AAS/AIAA Space Flight Mechanics Meeting*, 2021.

[19] A. Vananti and T. Schildknecht. Metric for object association in the Keplerian orbits space. In *71st International Astronautical Congress (IAC)*, 2020.

[20] S. van der Walt, J. L. Schrönberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Gouillart, T. Yu, and the scikit-image contributors. scikit-image: image processing in Python. *PeerJ*, 2:e453, 6 2014.

[21] L. Bradley, B. Sipőcz, T. Robitaille, E. Tollerud, Z. Vinícius, C. Deil, K. Barbary, T. J. Wilson, I. Busko, A. Donath, H. M. Günther, M. Cara, P. L. Lim, S. Meßlinger, Z. Burnett, S. Conseil, M. Droettboom, A. Bostroem, E. M. Bray, L. A. Bratholm, W. Jamieson, A. Ginsburg, G. Barentsen, M. Craig, S. Pascual, S. Rathi, M. Perrin, and B. M. Morris. astropy/photutils: 2.2.0, February 2025.

[22] A. W. Harris, J. W. Young, E. Bowell, L. J. Martin, R. L. Millis, M. Poutanen, F. Scaltriti, V. Zappalà, H. J. Schober, H. Debehogne, and K. W. Zeigler. Photoelectric observations of asteroids 3, 24, 60, 261, and 863. *Icarus*, 77(1):171–186, 1989.

[23] S. Mottola, W. D. Sears, A. Erikson, A. W. Harris, J. W. Young, G. Hahn, M. Dahlgren, B. E.A. Mueller, B. Owen, G.-H. Ricardo, J. Licandro, M. A. Barucci, C. Angeli, G. Neukum, C.-I. Lagerkvist, and J. F. Lahulla. The slow rotation of 253 Mathilde. *Planetary and Space Science*, 43(12):1609–1613, 1995.

[24] S. Hellmich, E. Rachith, and J.-P. Kneib. Detecting space objects in the ESO VST astronomical data archive and extracting their astrometry and photometry. In *11th European Conference for AeroSpace Sciences EUCASS*, 2025.

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com