

## AUTOMATED EXTRACTION OF SATELLITE TRAILS FROM WIDE ANGLE CCD IMAGERY

G. Privett<sup>1</sup>, S. George<sup>1</sup>, W. Feline<sup>1</sup>, and A. Ash<sup>1</sup>

### ABSTRACT

The current UK's National Space and Security Policy states that the identification of potential on-orbit collisions and warning of re-entries over UK sovereign territory is of high importance, driving requirements for indigenous Space Situational Awareness (SSA) systems capable of delivering these products.

The UK's Defence Science and Technology Laboratory (Dstl) is examining options for the creation of a distributed network of longitudinally distributed, low cost commercial-off-the-shelf electro-optical sensors to support survey work and catalogue maintenance. To effectively exploit this, a robust data handling system is required to autonomously detect satellite trails in a manner that can handle variable target intensities, periodicities and rates of apparent motion, as well as meteors and aircraft.

Data captured during the United Kingdom/New Zealand Automated Transfer Vehicle-5 (ATV-5) de-orbit observation trial have been employed to inform the development of a prototype processing pipeline for autonomous on-site processing. The approach taken employs pre-existing and documented tools such as Astrometry.NET and DAOPHOT from the astronomical community, together with image processing and orbit determination software developed in-house by Dstl. Some preliminary results from the automated analysis of data collected from wide angle sensors are described, together with an appraisal of the limitations of the proposed system and our plans for future development

### RESUMEN

La actual Política Nacional de Espacio y Seguridad del Reino Unido establece que la identificación de posibles colisiones en órbita y la advertencia de reingresos en el territorio soberano del Reino Unido es de gran importancia, lo que impulsa los requisitos para los sistemas indígenas de Conocimiento de la Situación del Espacio (SSA) capaces de entregar estos productos.

El Laboratorio de Ciencia y Tecnología de Defensa del Reino Unido (Dstl) está examinando opciones para la creación de una red distribuida de sensores electro-ópticos comerciales, distribuidos longitudinalmente y de bajo costo, para respaldar el trabajo de inspección y el mantenimiento del catálogo. Para explotar esto de manera efectiva, se requiere un sistema robusto de manejo de datos para detectar de forma autónoma las trayectorias de los satélites de una manera que pueda manejar las intensidades variable de los objetivos, las periodicidades y las tasas de movimiento aparente, así como los meteoros y las aeronaves.

Los datos capturados durante el ensayo de observación de órbita de la Unidad de Transferencia Automática del Reino Unido / Nueva Zelanda (ATV-5) se han empleado para informar el desarrollo de un prototipo de software para el procesamiento autónomo en el sitio. El enfoque adoptado emplea herramientas preexistentes y documentadas, como Astrometry.NET y DAOPHOT de la comunidad astronómica, junto con el software de procesamiento de imágenes y determinación de órbitas desarrollado internamente por Dstl. Se describen algunos resultados preliminares del análisis automatizado de los datos recopilados de los sensores de gran angular, junto con una evaluación de las limitaciones del sistema propuesto y nuestros planes para el desarrollo futuro.

*Key Words:* astronomical instrumentation, methods and techniques — methods: data analysis — techniques: image processing

### 1. INTRODUCTION AND CONTEXT

The ever-rising tempo of satellite launches, combined with the inevitable growth in debris arising

from on-orbit collisions, has meant that space has become a relatively congested domain. Consequently, it is becoming difficult for the existing Space Situational Awareness (SSA) systems to monitor the situation fast enough to enable the safe and secure operation of national assets and infrastructure.

<sup>1</sup>Defence Science & Technology Laboratory, Porton Down, SP4 0JQ, UK.

The United Kingdom (UK) National Space and Security Policy (NSSP; Her Majesty et al. 2014) acknowledged the need to undertake work on options that might help support the creation of an enhanced UK SSA picture. To that end, work is underway at the UK Defence Science and Technology Laboratory (Dstl) to explore these options, including experimenting with the use of repurposed radar systems such as the UK's Science and Technology Facilities Council (STFC) Chilbolton Observatory - a fully steerable 25m radar normally employed for meteorological research. Unfortunately, large facilities such as Chilbolton require a considerable investment and so are less likely to form the basis of a longitudinally distributed sensor system. The radar also lack the power to detect objects significantly beyond low Earth orbit (LEO) which limits its utility for SSA which, by its nature, requires coverage across a range of orbital regimes.

To overcome this, Dstl has sought to identify electro-optical (EO) systems and components that might help form the basis of an outline for affordable systems that can be deployed widely and has studied the SuperWASP (Street et al. 2003), MAS-CARA (Talens et al. 2003), the Dragonfly camera array (Abraham et al. 2014) and other facilities (Ackermann et al. 2016). Work is being undertaken to determine how to implement a system capable of fully automated day-to-day operations, including pipelined analysis. By undertaking this process, Dstl staff will renew their knowledge of the field, gain a representative source of data together with the detailed technical experience essential to support our later interactions with potential contractors.

Of the activities underway in this area, the focus for the current paper is on the development at Dstl of an automated image analysis system and discusses, primarily, details of the data reduction and analysis pipeline being employed.

## 2. HARDWARE AND IMAGERY

Test imagery for the processing pipeline was collected during the Joint New Zealand/ United Kingdom ATV-5 trial (Ash et al. 2015) using consumer-level, off-the-shelf Canon or Nikkor lenses (of either 20mm or 24mm at f/2.8), coupled with Starlight Xpress H18 or Trius SX694 Peltier cooled Charge Coupled Device (CCD) cameras. These visible band cameras have quantum efficiencies of 51

During the trial, approximately 30GB of imagery was collected against a range of target satellites and debris. Consequently, the test dataset contained brightness trails associated with both operational

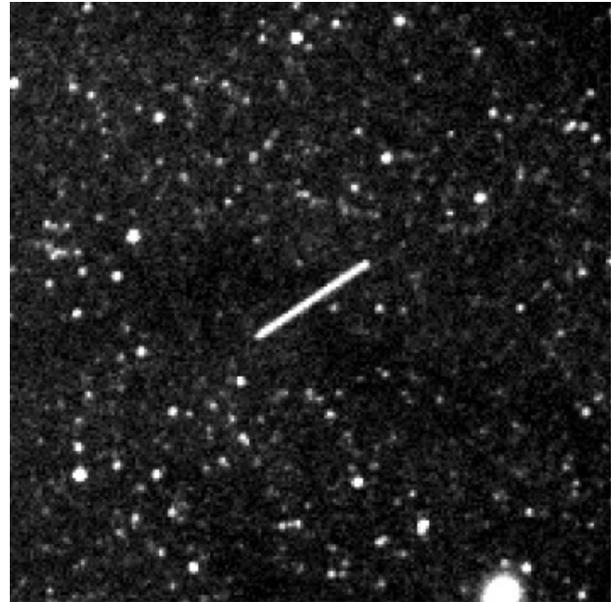


Fig. 1. Example of a section of a CCD frame showing a satellite trail.

satellites and retired/failed systems with tumbling dynamics or with surfaces generating flares in brightness; an example is presented in Figure 1. The variety of both the target set and the deployed equipment helped to ensure that the approach was adaptive to the collected data rather than stove-piped to reflect the behaviour of a single system. The variability of targets also ensured the robustness of the method was tested against a varied dataset.

## 3. METHOD

The detection, extraction and measurement of satellite trails has been the subject of considerable previous study within the community (Stöveken et al. 2005; Levesque 2009; Virtanen et al. 2016), and a considerable number of approaches to the process have been tested and adopted. As computers have become more powerful these methods have become more sophisticated and computationally intensive, but many approaches share broadly-similar components. These are:

- Calibration. Apply standard data reduction techniques to the imagery. This should include: dark subtraction, flatfielding.
- Map potential imager defects. Identify locations of defective (“dead”, “hot” or “stuck”) pixels.
- Determine the scene plate solution from the original image. Create a polynomial that is ex-

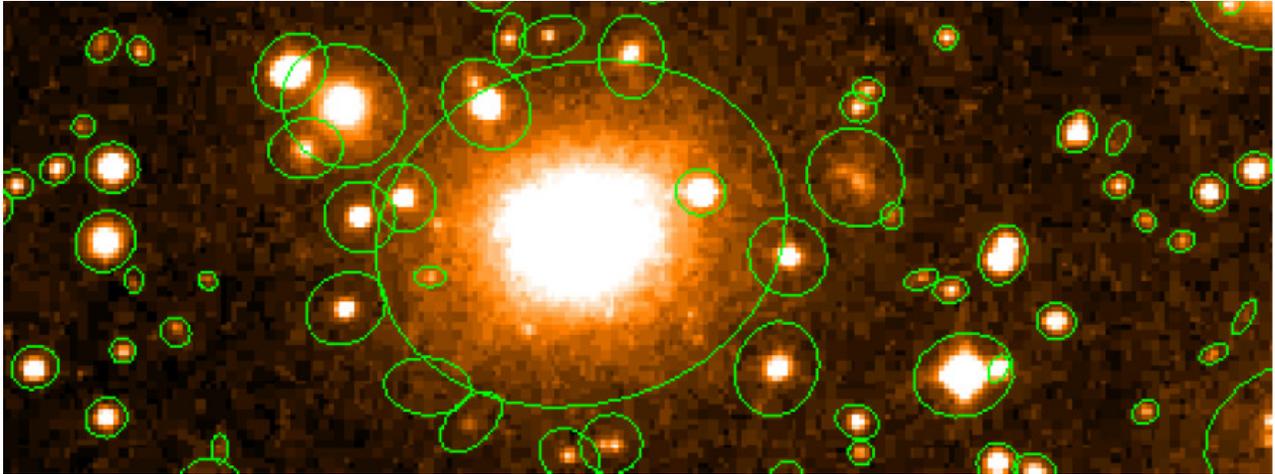


Fig. 2. Example of Source Extractor in use against an image frame containing stars and galaxies. Image: Peter Draper.

ploited to provide precise sky coordinates for each image pixel (various possible algorithms).

- Locate stars. Find those image sources that have a point spread function (PSF) similar to that the optical system.
- Estimate the scene background. Estimate the image brightness at location between the distinct image sources and employ that to create a 3 surface (x, y, pixel count) that approximates it.
- Remove the image scene and mask the stars. The remaining significant sources are moving targets such as meteors, planes or satellites.
- Find statistically significant clumps of pixels in the residual image. Use number of pixels and elongation of the “clump” pixels to highlight candidate trails and determine their size, orientation and start/end points.
- Create list of coordinates in RA and Dec for all trail star/end points.
- Look for consistent “movers” by comparing the positions of trails from frame to frame and assuming a “reasonable” movement model.

Clearly, there are a nearly endless array of permutations and variants possible when creating a pipeline fulfilling this approach, but for this application an approach was selected that took into account simplicity of execution, speed, robustness of execution while exploiting the following assumptions:

- Existing software packages should be exploited as appropriate – Source Extractor/Astrometry.Net, etc., (see Figure 2).
- the optical/CCD systems will be of amateur astronomy consumer quality rather than cutting edge,
- a sequence of wide-angle image frames are available, (see Figure 3).
- the targets are brighter than a Signal-to-Noise Ratio(SNR) of 5,
- the observation site is remote and unmanned,
- the connection bandwidth of the local system may be moderately constrained,
- the processing of imagery is to take place in (near) real-time,
- processing power on site is similar to that of 3-4 quad-core Intel i7-based laptops,

The focus has been on creating a demonstration system that is simple to deploy, handles the idiosyncrasies of the operated sensors effectively and employs existing (or modified) tools and techniques to reduce development time.

## 4. RESULTS

### 4.1. Data Reduction

The processing phase started with the standard data calibration techniques of dark subtraction and flat fielding (Privett et al. 2010; Howell et al. 2000). The dark and flat frame employed are both constructed through the median stacking of more than

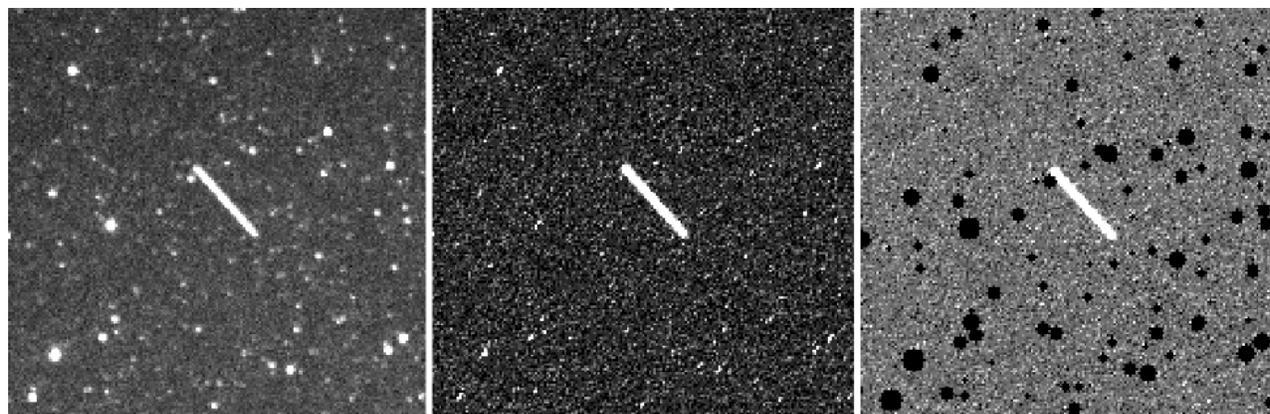


Fig. 3. An output image at 3 stages of the processing phase. After dark subtraction/flat-fielding (left), after an aligned median image stack is subtracted (middle), and after the detected stars have been blanked and the residual background modelled and removed (right.)

10 frames which eliminates the potential impact of cosmic ray strikes against the sensor during the collection of the flat or dark frames. In addition, defective pixels within the CCD array arising from manufacturing flaws were mapped: these flaws typically manifest themselves as a few very responsive (“hot”) pixels, a similar number of less responsive “cold” pixels and a small percentage of pixels that are deemed unreliable in their response; identified by comparing the statistics of all the pixels within the dark frames collected and masking these from the image. With the consumer level CCDs used in the ATV-5 trial, masking approximately 0.5

The second stage involved the detection of sources of star-like profile. Several pre-existing packages were identified for this purpose. They included: Source Extractor by Bertin and Arnout (Bertin et al. 1996), DAOPHOT by Stetson (Stetson et al. 1987) and the Starlink Project’s PISA (Draper et al. 2002). Source Extractor and PISA were designed for the detection of galaxies while DAOPHOT was designed for crowded field star detection.

To test their performance, all three packages were tested against imagery containing dim, isolated and unresolved targets within normally distributed noise fields. The tests showed that they all performed very similarly for circular, broadly star-like targets and so a version of DAOPHOT, made more tolerant of non-stellar sources, was adopted as our prime detector. It is worth noting that all three detectors worked most consistently when the estimate of image standard deviation ( $\sigma$ ) was derived from the width of the histogram rather than from a direct calculation, presumably because non-Gaussian residual instrument noise exaggerates the value of  $\sigma$ . This

approach was also found to be more successful than using a statistic such as absolute deviation. The detector was used with a  $10\sigma$  threshold to determine the location of bright fiducial stars within the image, which were used to identify triangles of stars shared between the frames and thereby infer the transformation and image rotations that were required to align the images to each other. These parameters might be expected to progress linearly from frame to frame, but they were not always well fitted by a linear regression, since winds or tripod sag could (and, during the ATV-5 trial, occasionally did) affect the pointing.

The image frames were then processed in sequence and, for each frame in the sequence, a master image was created by median stacking some of the frames immediately before and after it. Each frame then had its own master frame (Frame N) subtracted from it, greatly reducing the signal from the stars present and eliminating much of the unresolved background variability arising from the Milky Way or other diffuse light sources – thereby avoiding the need to iteratively assess and estimate the image background. Given the constraints upon processing power likely to be available in a future (deployed) system, the pipeline was usually run using median stacking of 5 images i.e. frame N-2 to N+2 when exposures are 1-2 seconds and of similar duration to the download time. Some tuning would be essential for a specific operational system.

The DAOPHOT detector is run against each frame at a  $3\sigma$  threshold level to determine the location of all the stars captured within the frame. A second routine then determines the extent of the stars and replaces the pixels occupied by a star with

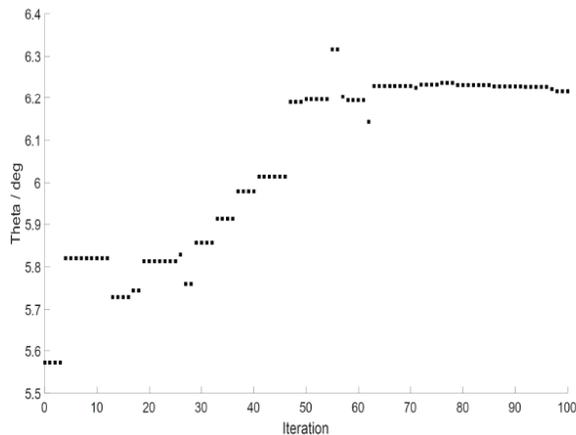


Fig. 4. Example of the downhill simplex minimisation process used to fit a trail profile to a 2-D Gaussian (PSF) function, demonstrating how the trail angle parameter varies as the iterative process is progressed.

a “magic” value which helps the subsequent processing identify when a star and satellite trail potentially overlap. In a following step, a polynomial is then fitted to each processed image and used to remove any residual background variation. This has been included to make the system more robust to the presence of cloud, haze, aurora and mists.

#### 4.2. Track Detection

By this stage, the image frames are essentially “free” of background variation and stars have been masked out and set to values that allow them to be excluded from subsequent processing. Consequently, anything consisting of a few closely-spaced (“clumped”) pixels with intensities above the image noise can be deemed to be a candidate potential detection (or may simply be a statistical outlier which must be eliminated). There are several possible approaches to target detection which may be employed.

- Existing software packages should be exploited as appropriate – Source Extractor/Astrometry.Net, etc.,
- Clumper. Use a simple “clump” detector – IDL and MATLAB have one and the Starlink CUPID package (Berry et al. 2016) has several – that deems any associated group of above threshold pixels to be a potential target. This has the virtue that potential trails are rapidly identified and subjected to tests for ellipticity, but this technique for detection tends to be less effective for trails with very low brightness levels.

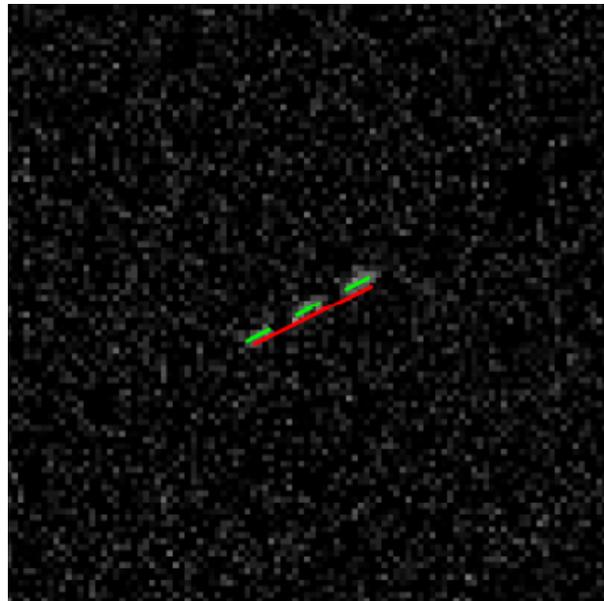


Fig. 5. Example of the downhill simplex minimisation process used to fit a trail profile to a 2-D Gaussian (PSF) function, demonstrating how the trail angle parameter varies as the iterative process is progressed.

- (Binary) Hough. Employ a Hough (Duda et al. 1972) transform to examine the image for the presence of distinct line-type features. In such an approach, images are converted to a binary form where those pixels above a set threshold are set to 1 and all others to zero. The transform then generates a transform surface from which the location and direction of motion of trails can be identified. However, fully automating this step in a robust manner proved problematic; as the Hough transform technique is highly sensitive to the noise statistics and image artefacts, making automated thresholding difficult to implement. The technique appears to work best when performed on smaller, sub-images.
- (Greyscale) Hough. Apply a variant of the Hough approach which works against greyscale imagery. This is a more computationally intensive approach, but is also more sensitive to the presence of extended low SNR targets. As with the binary Hough procedure, the technique is best employed against small sub-images. Unfortunately, it is also has the potential to create more false maxima (alarms) within the transform surface created and, in our MATLAB implementation, displayed a tendency to provide trajectories slightly misaligned with the trail centres.

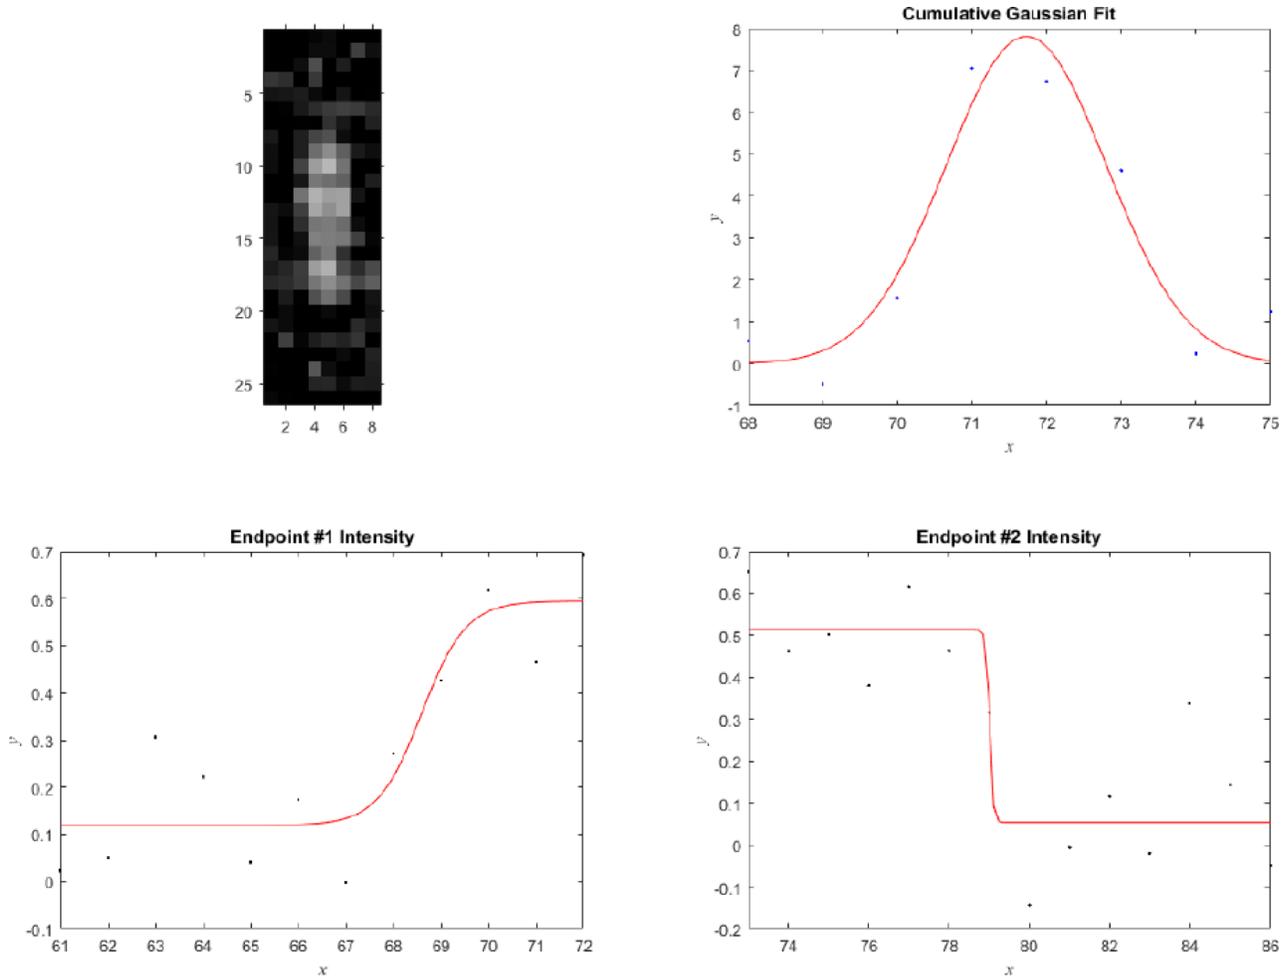


Fig. 6. (top, left) cropped region around a single satellite trail; (top, right) Gaussian profile fitted to the across-track profile of the trail; (bottom) Sigmoid functions fitted to the ends of the satellite trail. Note the differing Sigmoid fits assigned to each end; future work will look to model these more accurately based on the known optical PSF. The small number of pixel data points combined with noise remaining in the image background degrades the ability to routinely assign a robust fit.

- Correlation/matched filter. Apply a simple correlation-based detector to generate detection surface images and employ a Track Before Detect (TBD) (Salmond et al. 2001) approach to determine which of the potential detections within the scene is consistent with a moving trail. This approach is highly computationally intensive and grossly memory inefficient.

Each of these methods has its virtues and disadvantages and so, in view of our constrained hardware specification, our wide field-of-view and the knowledge that 4 satellites in the same frame can be quite a common occurrence, further development of the CPU / memory-hungry TBD method was not pur-

sued. It was also decided that the benefits and behaviour of the greyscale Hough approach needed more substantial examination before it could be vigorously implemented. Consequently, the simple Clumper and the binary Hough transform methods were adopted as the basis of our pipeline configuration. These yield a trajectory estimate and the end points of each track within a frame. Both techniques are relatively simple to understand and to code, leaving automation the outstanding challenge. The detection process operates on individual image frames independently and on a cropped area of the total image frame, with no a priori information about either the number of targets in the scene or the expected position of a known target in the subsequent frame.

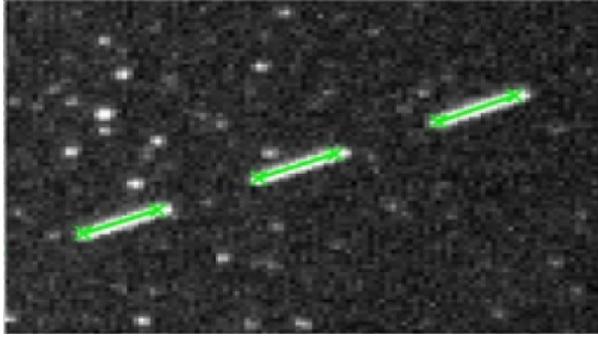


Fig. 7. Example 3-streak (6 shutter operations during CCD exposure) trail showing its limits as deduced by the trail cross section refinement method.

#### 4.3. Positional Refinement

To enable accurate orbit determination, the start and stop locations for each satellite trail must be derived while taking into account the convolution of the unresolved satellite with the PSF of the optical imaging system, since these are the only target positions for which there is well defined time stamp; i.e. the start and stop times of the exposure. The times at points along the trails can only be estimated.

It was found that, while the Hough and Clumper approaches efficiently identified candidate target trails and their start and stop locations, their accuracy was limited to approximately 1 pixel resolution in the direct output of the detector. For much of the imagery used, this corresponds to an accuracy of approximately 74 arc seconds which is equivalent to a 2.4km positional error for a typical LEO satellite. To improve upon this, a refinement step was needed to achieve greater (sub-pixel) accuracy. Two methods of end point refinement were explored using the Hough/Clumper outputs as initial estimates of the position and orientation of the trail:

In the first method, a trail profile (orientated to the trajectory identified by the Hough/Clumper method and of similar length and intensity) is convolved with a 2-D Gaussian function representative of the PSF of the optical system to create a “template” image. The parameters of the template image – PSF, trail length, intensity, start point position and orientation – are then iteratively fitted to the observed trail by chi-squared minimisation. The minimisation algorithm employed was a multidimensional downhill simplex minimisation derived from the Nelder-Mead “AMOEBa” technique (Nelder et al. 1965). In general, the algorithm converged within 100-150 iterations; an example of parameter convergence is shown in Figure 4.

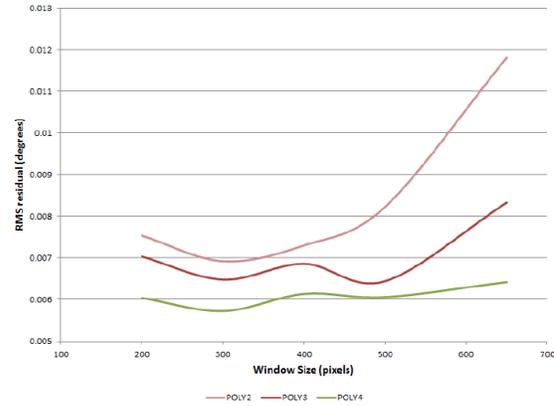


Fig. 8. The error residuals (degrees) as a function of size of sub-image employed, for 2nd-, 3rd- and 4th-order WCS polynomial fits. Results were generated for varying window sizes at 4 different locations of the original image; the mean RMS angle error across all 4 locations is displayed here.

The second refinement method examines the trail cross-section at various locations along its length and estimates the optical PSF of the system which, in turn, allows the trail orientation to be more precisely determined. This approach makes the assumption that the image PSF is Gaussian, which holds true when the trail is far enough away from the edge of the image frame. The variation in pixel intensity along the trail’s central axis is then estimated and an attempt made to define the shape of the intensity roll-offs at the ends of each trail, thereby allowing an estimate of the trail start and stop positions to be made (see Figure 5).

Both methods have proved to be effective in refining the initial estimates of the trail, but are sensitive to rapid variations in satellite brightness and may struggle with low SNR trails and where the intensity profile varies quickly – an example being the CanX-7 solar sail test satellite where, during a recent observation, it was seen to vary in brightness by nearly 3 magnitudes within a second. Further work is needed to make both systems more robust.

#### 4.4. Astrometry

Following the application of the endpoint-refinement method, the system yields the satellite position in x, y pixel space which must be converted to an RA-DEC angle pair. The Linux package Astrometry.NET (Lang et al. 2010), funded by the US National Science Foundation and others, was utilised to perform blind-solving astrometric calibration for our image frames. The Astrometry.NET plate solver employs pattern-matching techniques against star

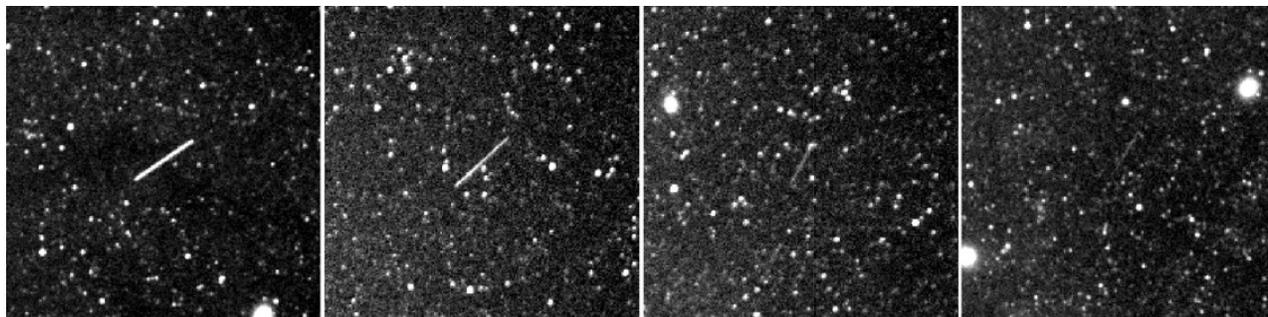


Fig. 9. The error residuals (degrees) as a function of size of sub-image employed, for 2nd-, 3rd- and 4th-order WCS polynomial fits. Results were generated for varying window sizes at 4 different locations of the original image; the mean RMS angle error across all 4 locations is displayed here.

catalogues, comparing triangles formed from stars within the scene with triangles generated using the catalogue positions. Importantly, it does not require any a priori information on the image scale and has proved very robust.

Our current instantiation of the pipeline calls the Astrometry.NET web server via its Application Program Interface (API) and supplies small, trail-centred image subsets for solution. Astrometry.NET then blind-solves the image to derive a World Coordinate System (WCS) polynomial relating pixel space coordinates to celestial coordinates. Our studies found that the accuracy of the final result depends upon a number of factors including the size of the image sample supplied, its location on the image frame, the order and type of the polynomial order employed to store the WCS transformation. An example is presented in Figure 8.

As might be expected given the wide field-of-view of the imagery, the astrometric accuracy improved as the sub-image size was reduced and appeared to be near optimal for sub-images of between 300-500 pixels wide. Additionally, varying the order of the polynomial showed that the errors were smallest when using a 4th-order polynomial, with minimal improvement (and over tuning) observed at higher orders. Tests of the accuracy against catalogued star positions indicated that errors associated with the solution were approximately  $0.007^\circ$  (around 0.3 pixels) and comparable with the error associated with the uncertainty in system time. The timing error of the system is associated with a Windows-based laptop clock locked to GPS using the NMEATIME2 software (by VisualGPS LLC.) rather than via a GPS receiver's 1pps output, so can likely be improved upon.

During the ATV-5 trial, Dstl collaborated with colleagues in the New Zealand Defence Technology Agency (DTA), who are examining similar techniques (in both hardware and software) for extrac-

tion of satellite trails from EO imagery (Skuljan et al. 2015). The DTA procedure currently uses a manual process to detect the target trails, but implements a specialised, in-house astrometric solver named StarView (Skuljan et al. 2016). Based on qualitative visual analysis, then techniques appear to derive similar results in terms of trail endpoint positioning and astrometric referencing. However, this remains an area of future work to compare and contrast differing techniques and to validate performances against alternative methods.

## 5. CONCLUSION

IDL and MATLAB code has been developed at Dstl which provides a demonstration pipeline for the automatic processing of satellite trails from fixed pointing or sidereally tracked sensors. The camera lens hardware used for this demonstration lacked the sensitivity to capture most CubeSats (Ackermann et al. 2003), but the software would work equally well with a more capable system employing a system of greater light grasp, such as a Celestron Rowe-Ackermann tube assembly.

We have demonstrated the application of an algorithm that was relatively easy to code (and maintain), that was not computationally intensive and which, when used with relatively modest hardware, could be optimised to provide near real-time data creation. The system is amenable to improvement and enhancement, including the implementation of a “slow time” processing phase to be employed during the observatory daytime, when more intensive processing may be applied.

**Acknowledgments:** Crown copyright (2018), Dstl. This material is licensed under the terms of the UK Open Government Licence except where otherwise stated. To view this licence, visit <http://www.nationalarchives.gov.uk/doc/open-government-licence/version/3> or

write to the Information Policy Team, The National Archives, Kew, London TW9 4DU, or email: [psi@nationalarchives.gsi.gov.uk](mailto:psi@nationalarchives.gsi.gov.uk)”.

## REFERENCES

- Abraham, R. G. & van Dokkum P. G. 2014, *PASP*, 126, 55
- Ackermann, M. R., Cox, D. D., McGraw, J. T., Zimmer, P. C. 2016, *Lens and Camera Arrays for Sky Surveys and Space Surveillance*, Proceedings AMOSTECH SSA Conference, Maui 2016, (SAND2016-8077)
- Ackermann, M. R., McGraw, J. T., Martin, J. B., Zimmer, P. C. 2003, *Blind Search for Micro Satellites in LEO: Optical Signatures and Search Strategies*, Proceedings AMOSTECH 2003
- Ash, A., Donnelly R. P., *The critical Role of Experimentation to Further SSA Understanding*, Proceedings AMOSTECH, Maui, 2015
- Berry D. S., *CUPID A 3D Clump Identification Package*, SUN 255, Starlink Project, UK STFC, 2016
- Bertin, E. & Arnouts, S. 1996, *A&AS*, 117, 393
- Draper P. W., Eaton N., *PISA Position Intensity and Shape Analysis*, SUN 109, Starlink Project, UK STFC, 2002
- Duda, R.O. & Hart, P. E. 1972, *Use of the Hough Transformation to Detect Lines and Curves in Pictures*, *Communications of the ACM*, 15, 11
- Her Majesty’s Government, 2014 United Kingdom National Space Security Policy, Her Majesty’s Stationery Office, URN:UKSA/13, 1292, April 2014
- Howell, S. B. 2000, *Handbook of CCD Astronomy*, (Cambridge, MA: CUP)
- Lang, D., Hogg, D. W., Mierle, K., Blanton, M., & Roweis, S. 2010, *AJ*, 139, 1782
- Levesque M. P., *Automatic Reacquisition of Satellite Positions by Detecting Their Expected Streaks in Astronomical Images*, Proceedings AMOSTECH, Maui, 2009
- Nelder, J. A. & Mead, R. 1965, *A Simplex Method for Function Minimization*, *Computer Journal*, 7, 308
- Privett, G. 2010, *Creating and Enhancing Digital Astro Images*, (Springer)
- Salmond, D.J. & Birch, H. 2001, *A Particle Filter for Track-Before-Detect*, Proceedings of American Control Conference, 25, 3755
- Skuljan, J. 2015, *Trial Firefly – Building a Space Situational Awareness Capability at the Defence Technology Agency*, DTA Report, 402
- Skuljan, J. & Kay, J. 2016, Proceedings of Advanced Maui Optical and Space Surveillance Technologies, ed. S. Ryan (Maui Economic Development Board), 26
- Stetson, P. B. 1987, *PASP*, 99, 191
- Stöveken, E., Schildknecht, T., *Algorithms for the Optical Detection of Space Debris Objects*, Proceedings of the European Conference on Space Debris, 2005
- Street, R. A., Pollacco, D. L., Fitzsimmons, A., et al. 2003, *ASPC*, 294, 405
- Talens, G. J. J., Spronck, J. F. P., Lesage, A. L., et al. 2017, *A&A*, 601, 11
- Virtanen, J., Poikonen, J., Sääntti, T., et al. 2016, *AdSpR*, 570, 1607