Moving point source detection and localization in wide-field images

Przemek Wozniak
Los Alamos National Laboratory (LANL), wozniak@lanl.gov

Lakshman Prasad
Los Alamos National Laboratory (LANL), prasad@lanl.gov

Brendt Wohlberg
Los Alamos National Laboratory (LANL), brendt@lanl.gov

ABSTRACT

As part of the Thinking Telescopes project, a long-term Space Situational Awareness effort at the Los Alamos National Laboratory, we developed efficient algorithms for detecting and localizing unresolved moving sources in wide-field astronomical images. Our approach is based on pixel-by-pixel differencing of two images after one image has been convolved to match the point spread function of the other. The input images are obtained with the telescope tracking at the sidereal rate, and therefore stars are imaged as point sources and moving objects streak. The algorithm first identifies islands of connected pixels above a significance threshold in difference images where static sources have been removed. The next step is line detection and fitting light profiles of streaked sources using empirical models. The main objective is to optimize the accuracy of the measured streak end points while keeping the required computation reasonably fast. To this end we apply a series of progressively more sophisticated estimators until the result is no longer improved due to noise limitations. The final astrometry for sufficiently bright streaks takes into account the local point spread function for a particular location within the wide-field image. We found efficient 1D and 2D algorithms for fitting models of moving point sources (infinitely thin lines convolved with the point spread function). The 2D algorithm based on the iterative locally linearized fit parameterized with sub-pixel locations of both ends provides up to an order of magnitude improvement in accuracy over the 1D algorithm by removing the bias in the orientation of the streak. Other advantages of the new approach include: the ability to accurately measure very short streaks, much better sensitivity for faint objects, better convergence (after about two iterations), and eliminating the need for interpolation of the original image and 1D projection. These algorithms have been tested on real and simulated ground-based images of Earth satellites. They are robust and have a potential to significantly improve the utility of cost-effective space surveillance systems that rely on commodity imaging and computing hardware.

1. INTRODUCTION

Persistent monitoring of the night sky is a gold mine of information on both natural and man-made objects in space. Frequent ground-based optical imaging of the entire night sky is rapidly becoming one of the key elements of the global Space Situational Awareness. Wide-field imaging telescopes of relatively modest size can address some of the most pressing and growing challenges in this area, for example those associated with dramatically increasing density of deployed satellites and space debris [1]. The same capabilities are essential to the success of massive photometric surveys that drive progress in time-domain astrophysics and allow scientists to explore the realm of cosmic explosions such as supernovae and gamma-ray bursts [2].

Over the past two decades, the cost of commodity imaging sensors and computing hardware has been steadily decreasing, making persistent sky monitoring on a global scale possible. Explosive growth in the number of sky monitoring projects and the amount of data they deliver demands better tools to extract high quality measurements from astronomical images. The available algorithms are often inefficient and produce too many artifacts, and as a result fail to adequately support automated data processing and real-time reporting of actionable information in a data intensive setting. Very wide-field imaging poses additional challenges associated with accurately modeling the point-spread function (PSF) that varies depending on the location within the image. Another problem is that due to a tradeoff between sky coverage and resolution, wide-field imagers deliver extremely crowded images where overlapping PSF

profiles are common. Crowding poses a considerable challenge, as it hides variable sources such as explosive transients and satellite tracks that must be detected against the background clutter of static astronomical objects (stars). Specialized techniques have been developed for matching the PSF of two images so that the clutter can be subtracted by taking a direct pixel-by-pixel difference of two images of the same scene. Monitoring of Earth orbiting satellites creates additional complications around "streaking" that occurs in ground-based images with finite exposure times. When the telescope tracks at sidereal rate, the stars are imaged as PSF profiles and unresolved moving sources appear as streaks. An alternative approach is to track the moving object to maximize its signal-to-noise ratio in the image and let the stars streak. In both cases we are faced with the problem of detecting and characterizing streaks that are stretched along the direction of motion, but unresolved across the trajectory. The most important measurement to be extracted in this case is the position of the object at some reference time, for example the ends of the streak can be tied to shutter opening and closing times.

2. THINKING TELESCOPES PROJECT AND RAPTOR TELESCOPE NETWORK

Thinking Telescopes is a long-term effort at the Los Alamos National Laboratory with the goal of automating night time sky surveillance [3]. The project has developed the RAPTOR telescope network, a system of autonomous imaging instruments with a range of capabilities, connected using autonomous real-time machine-to-machine communications. Each RAPTOR "telescope" is typically an array of telescopes with optical elements ranging from photo lenses to 0.4-m telescopes (Fig. 1) that operate in sync. The system is built almost entirely using Commercial Off-The Shelf (COTS) components (optics, sensors, computing, storage) with the exception of custom made fast-slewing mounts and a few other details.

The key idea of Thinking Telescopes is to extract interesting events on the fly as they emerge from very wide-field imaging data and immediately trigger other more sensitive telescopes on the network to collect high-fidelity follow-up data without humans in the loop. This real-time feedback between wide-area search and rapid response allows the system to interrogate transient events before they are over. All data processing tasks such as image reduction, source extraction, calibration, orbit determination, anomaly detection, etc. must keep up with data acquisition rates. Each telescope enclosure is equipped with commodity computers that perform all necessary number crunching for a given image before the next exposure is collected. Typical cadence is one snapshot every 20-30 seconds. This places relatively strict limits on the efficiency of the algorithms that can be deployed on a cost-effective computing hardware. All image processing algorithms discussed here, including streak detection and localization, were developed with the goal of supporting real-time operation of the RAPTOR telescope network.

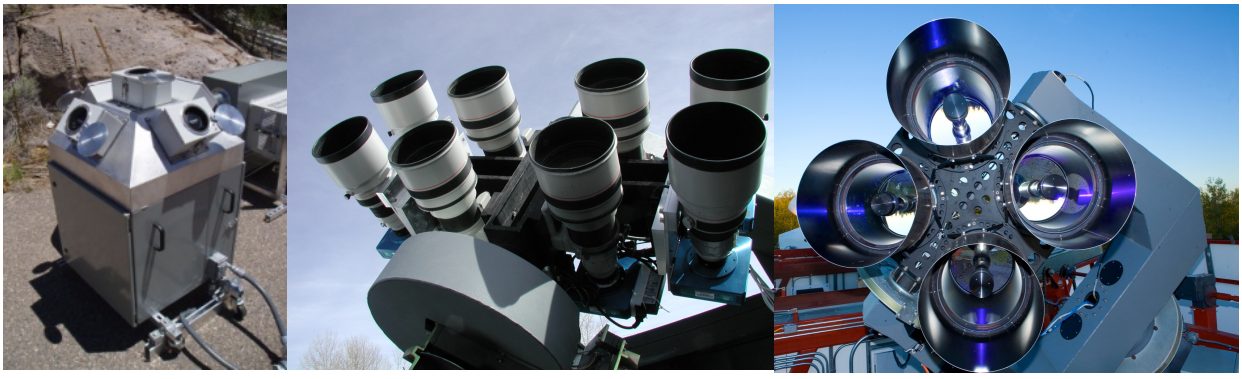

Fig. 1. Selected instruments deployed as part of the RAPTOR telescope network that illustrate the range of Thinking Telescopes capabilities: RQD2 all sky monitoring telescope shown without light baffles (left), RAPTOR-P wide-field patrol array (middle), and RAPTOR-T simultaneous multi-color imager (right). Together, these instruments provide access to a wide range of orbital regimes from LEO to GEO. The field of view covered by a single scope is 65×65 deg, 8×8 deg, and 24×24 arcmin for RQD2, RAPTOR-P, and RAPTOR-T respectively.

3. IMAGE DIFFERENCING AND PSF MODELING

Streak detection in very wide-field images is significantly complicated by the presence of a dense background of constant stars. In the extreme case of imaging near the Milky Way, every pixel contains signal from the PSF profile of some random star (Fig. 2). Away from the Galactic plane one still faces a very high probability that a fainter streak is completely dominated by moderately bright stars. Image differencing is an effective way to address this problem (Fig. 3). PSF matching may or may not be required depending on how the reference image to be subtracted is obtained. Typically, it is possible to skip PSF matching when subtracting two consecutive exposures of the same field imaged with a small telescope. In this case instrumental factors such as the telescope focus and tracking dominate the PSF and are sufficiently stable over time. However, the price for using single exposures as reference images is approximately a factor of 2 increase in photon noise variance. Stacking 10-20 images taken at different times will suppress the noise in the reference image, but the resulting PSF will almost certainly have to be matched to each program image to obtain a clean subtraction.

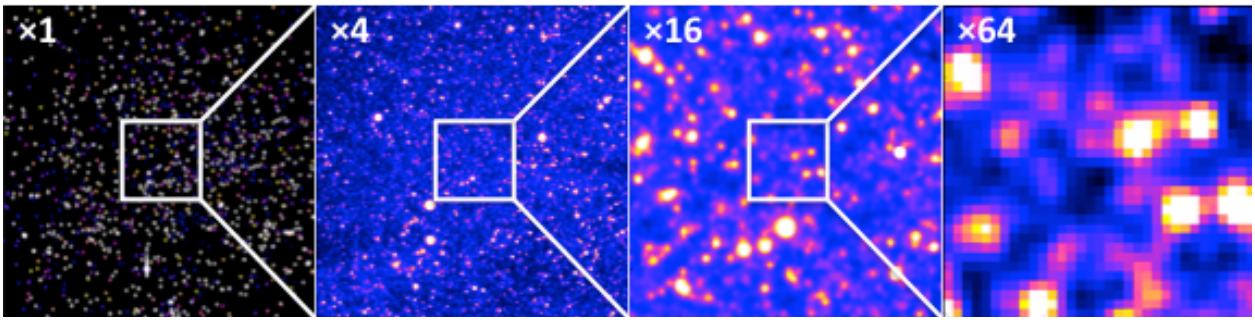

Fig. 2. Sample 10-second exposure collected by LANL's RQD2 all sky monitoring telescope illustrating challenges associated with stellar crowding. The image covers a moderately dense region of the Milky Way and has been smoothed with a Gaussian filter.

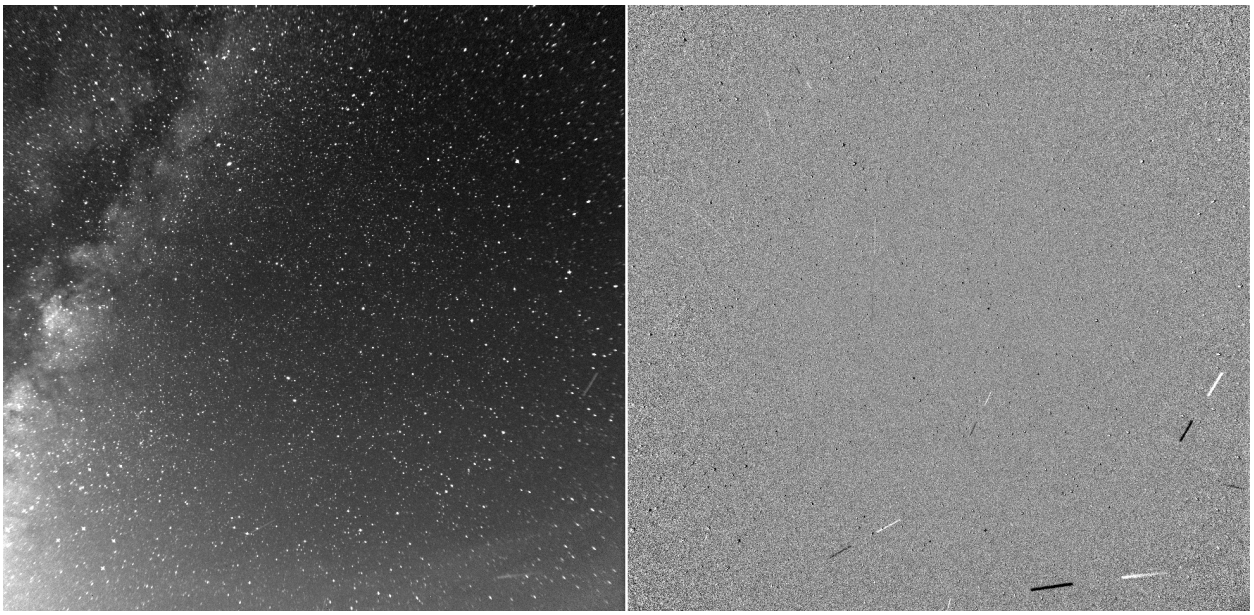

Fig. 3. Streak detection using image differencing. Very wide-field images such as those delivered by LANL's RQD2 all sky monitoring telescope are cluttered by PSF profiles of stars (left). Taking a direct pixel-by-pixel difference of two consecutive exposures of the same field removes the clutter and dramatically simplifies streak detection and characterization (right). The reference image to be subtracted must be interpolated to the pixel grid of the program image to compensate for sky motion over a fixed camera or imperfect telescope tracking.

Several effective PSF matching algorithms have been developed over the past two decades to address the challenges associated with accurate photometry and astrometry of variable sources in very dense stellar fields such as globular clusters and galaxies in the Local Group, primarily the Milky Way and Magellanic Clouds. For a review of challenges and techniques of crowded field photometry see e.g. [4]. The main assumption here, that is almost always justified, is that the scene does not change significantly between the two images to be subtracted, except for very few objects that may change their brightness and/or location. The Alard and Lupton algorithm [5, 6] and its variants [7] do not require a PSF model and intrinsically account for PSF gradients. The main idea behind this approach is to find a convolution kernel that can be applied to one image to closely approximate another image. The problem is then reduced to a global linear fit. A more recently published method called proper image subtraction [8] is closely related to cross-convolution. In this approach each image of the pair is convolved with the PSF of the other with a clever correction applied in the Fourier space to optimize the result for detection of a new point source. This, however, requires a good PSF model for each image and assumes that the spatial gradient of the PSF can be ignored. In practice, a piecewise solution with sufficiently small patches adequately handles even relatively strong PSF variations.

High-quality models of the PSF and its spatial variations are also important in streak characterization and localization discussed in the next section. For this purpose we are developing custom PSF estimation algorithms and codes that attempt to build on the strengths of well-established tools like DoPhot [9], DAOPhot [10], and SExtractor [11], and at the same time take advantage of modern image modeling approaches such as sparse representations, global blind deconvolution, and the latest regularization techniques. Our baseline approach is to obtain the initial PSF estimate using bright stars and then gradually include progressively fainter stars to account for all signal in the image (assuming it is dominated by point sources). The PSF model is a sum of the analytical core and a lookup table of corrections that can in principle fit any shape with proper regularization. The algorithm alternates between adjustments to the PSF model and the underlying scene (the number and locations/amplitudes of stars). This procedure is applied to image patches sufficiently large to provide enough information and sufficiently small to assume that the PSF is constant. Local PSF estimates are obtained by interpolating a grid of models that captures PSF variations over the entire image (Fig. 4).

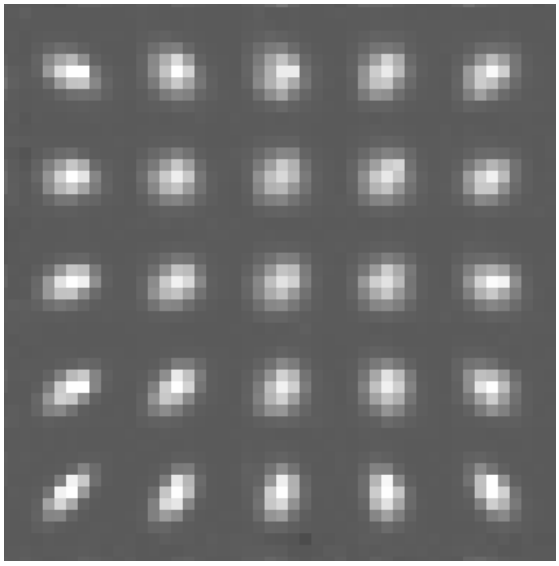

Fig. 4. PSF map for one of the very wide-field images from LANL's RQD2 all sky monitoring telescope. The shape of the PSF varies strongly depending on the location within the image.

An alternative approach is to solve a global blind deconvolution problem, i.e. estimate both the unknown convolution kernel and the activation map representing the unblurred image (sharp impulses encoding star locations and brightness). The PSF at each location is a linear combination of "basis" kernels, also estimated from data. The algorithm performs iterative optimization similar to gradient descent. The cost function is regularized using an $L_{2,1}$ norm to promote joint sparsity (activation maps are spatially sparse, but kernel combinations at each location are not).

While our initial results based on this technique are promising, they fall short of providing an acceptable solution. Another challenge of global blind deconvolution is computational complexity that may exceed the available time budget given computational resources constrained the total cost.

4. STREAK DETECTION AND LOCALIZATION

The Thinking Telescopes image processing pipeline performs automated streak detection using an algorithm similar to the one employed in the Source Extractor code [11]. This approach balances the speed and robustness against the fidelity. The difference image is first convolved with a Gaussian kernel mesh approximating the PSF to smooth out noise fluctuations and better expose faint streaks. The convolved image is only used to establish pixel footprints of detectable objects. Actual measurements use the original pixel data. The next step is identifying islands of connected pixels above a user specified threshold relative to noise and measuring basic location and shape parameters such as the first and the second moments of the light distribution. Only "positive" streaks are detected and measured in a given difference image (Fig. 3). "Negative" streaks are captured as "positive" in another subtracted image pair. Streaks are defined as sufficiently elongated sources with the area above some number of pixels.

Once streaks are identified, the data processing pipeline proceeds to astrometric and photometric measurements. The overall location and direction of each streak within the image is obtained using a straight line fit to the light profile weighted by background subtracted pixel counts. The fit is tied to one of the four image sides so that the slope of the object track is always between 0 and 1. Depending on the signal-to-noise ratio of a particular streak detection it may or may not be beneficial to apply more sophisticated estimators, which tend to be computationally more costly and less robust on noisy data. To this end we developed a sequence of progressively more accurate streak modeling and localization algorithms that terminates when the signal in the object profile does not warrant further improvement in accuracy. For the purpose of this discussion we assume that streaks are infinitesimally thin straight line intervals of constant intensity convolved with the PSF. RAPTOR telescopes (Fig. 1) are designed to keep the length of the streak between 20 and 200 pixels for the dominant orbital regime so that single tracks are very close to straight lines. While significant variations in intensity over a few seconds are observed in a small fraction of objects, in practice the constant brightness assumption does not cause too many problems for most applications.

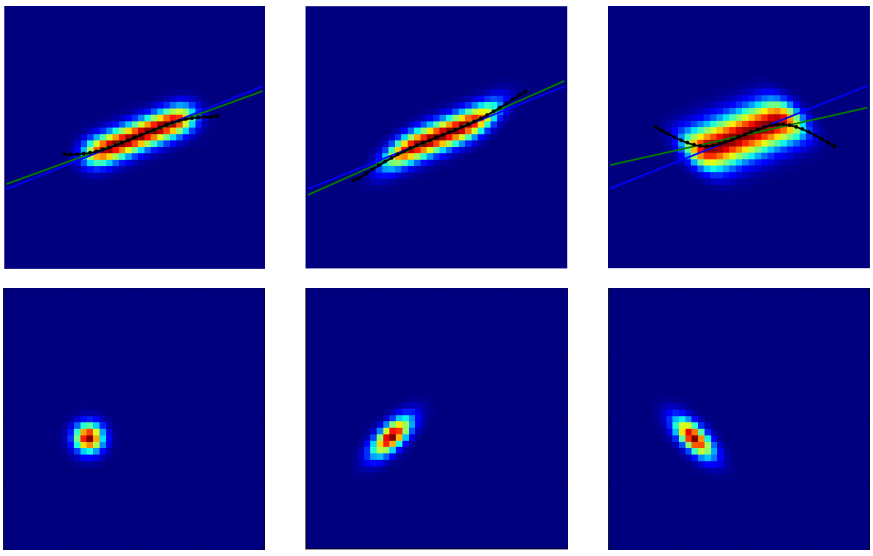

Fig. 5. Systematic errors in the measured streak orientation introduced by elongated PSFs. The bias in the estimate (top) depends on the shape and orientation of the PSF (bottom) relative to the direction of motion. Three examples are compared: perfectly circular PSF (left), elongation approximately along the streak (middle), and approximately perpendicular to the streak (right). The blue line is the true trajectory and the green line is an estimate based on a light-weighted straight line fit. The black curve shows the light centroid of each pixel column.

Streaks are localized by tying their end points to the opening and closing times of the camera shutter for a particular exposure. There are systematic effects that influence this basic relationship and may require additional corrections depending on the desired accuracy and the specifics of the imaging system. However, the main task of localizing the end points of the streak in image pixel coordinates remains the same. The first estimate is obtained from the intersection of the straight line along the track with the lowest and highest pixel row/column that overlaps with the object footprint above the detection threshold. Except for marginally detected streaks, this can be further improved by integrating the 2D light profile across the track to get the corresponding 1D profile along the track. Ideally, a detailed model fit of this 1D profile would account for the effective 1D PSF along the track to yield an accurate location of the streak. In practice, fitting a Gaussian error function *erf(x)* or a simple estimate of the half intensity point relative to the flat section of the streak is sufficient for faint streaks. For bright, well detected streaks a full 2D model fit is a better approach.

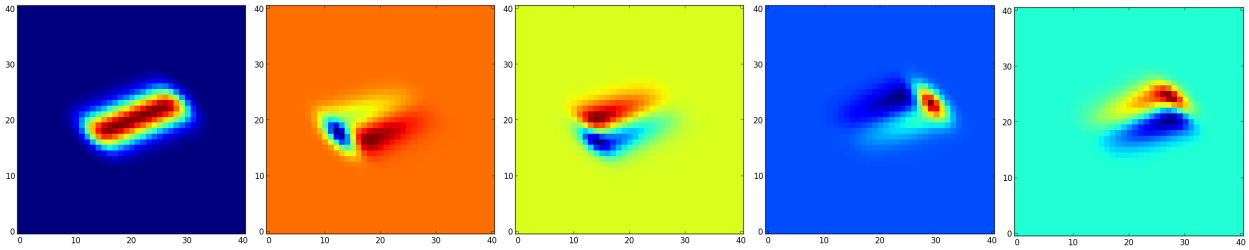

Fig. 6. Basis vectors of the linearized fit used to iteratively estimate streak end points. From left to right: 1) streak profile $P(x, y)$ given as a straight line interval convolved with the PSF, 2) dP/dx_1, 3) dP/dy_1, 4) dP/dx_2, 5) dP/dy_2.

The main problem with 1D modeling of streaks in the presence of strong PSF gradients and very elongated PSFs is the difficulty of obtaining an unbiased estimate of the orientation of the track. The resulting systematic error in astrometry depends on the relative orientation of the streak and the PSF (Fig. 5). This bias can be avoided by using full 2D streak models based on the local PSF profile estimated at the location of the streak. For streaks detected at high signal-to-noise ratio, our image processing pipeline implements an iterative, locally linearized fit parameterized with sub-pixel locations of both end points. From 1D algorithms we have the initial values of the end point locations (x_1, y_1) and (x_2, y_2) to compute the current best estimate of the streak profile $P(x, y)$. The intensity of each pixel within the streak footprint can be approximated using a linear expansion in end point offsets $(\Delta x_1, \Delta y_1, \Delta x_2, \Delta y_2)$:

$$I(x, y) = A \times \{ P(x, y) + \Delta x_1 \times dP/dx_1 + \Delta y_1 \times dP/dy_1 + \Delta x_2 \times dP/dx_2 + \Delta y_2 \times dP/dy_2 \}.$$

The offsets and the amplitude A are then estimated by fitting the streak image $I(x, y)$ with a linear combination of five model images, the current best streak profile $P(x, y)$ and its four spatial derivatives (Fig. 6). This basic step can be repeated to compute very accurate streak locations. Typically, just two or three iterations are sufficient to converge on a solution that does not improve with additional repetitions.

5. PERFORMANCE TESTING

The performance of our 2D streak fitting algorithm was tested using Monte Carlo simulations of images that span the relevant range of randomized streak parameters: length, orientation, sub-pixel location, PSF width and elongation, sky background, and signal-to-noise ratio in the brightest pixel. Our simulations include photon noise that follows Poisson statistics. Examples of images used to recover ground truth streak locations are shown in Fig. 7. Fig. 8 shows the r.m.s. scatter (per coordinate) of a single estimated end point around its true position as a function of the signal-to-noise ratio of the brightest pixel (averaged over all remaining parameters). At moderate to high signal-to-noise the accuracy of the full 2D algorithm is almost an order of magnitude better when compared to a 1D estimate using a half-max point along the track. Measurements as good as 0.05 pix are possible in the low noise limit. While the gains are somewhat less dramatic at low signal-to-noise ratio, streaks that are hardly discernible by eye (Fig. 7) can still be localized with sub-pixel accuracy. A major factor contributing to this improvement is the elimination of the bias from straight line fits to the light distribution of the streak.

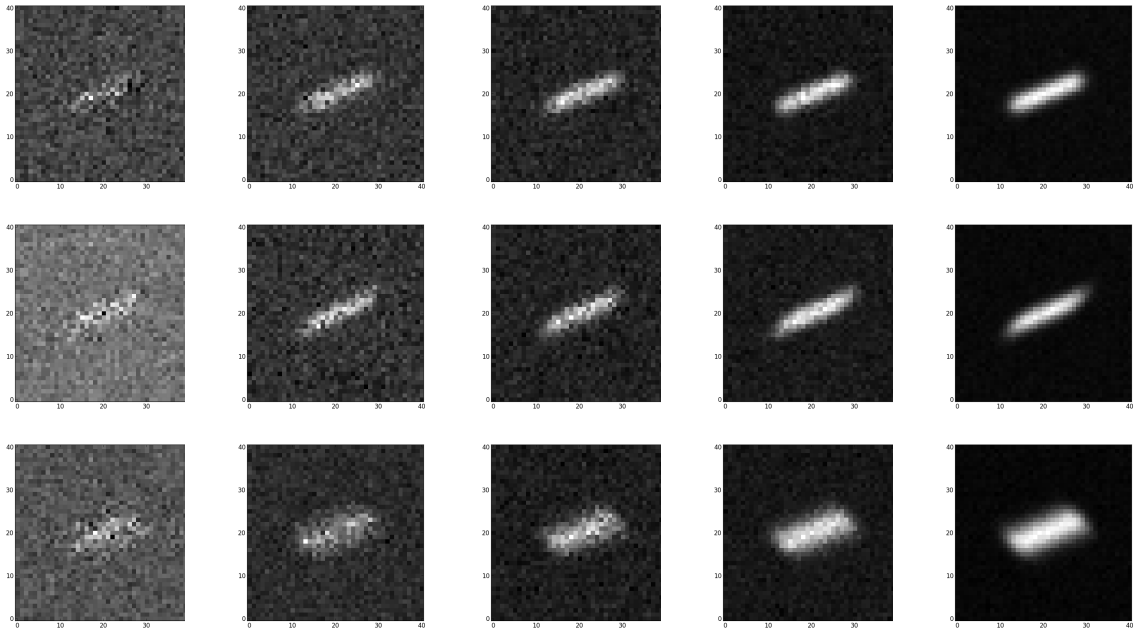

Fig. 7. Examples of simulated streak images used to evaluate the astrometric accuracy of the measured streak end points. Rows correspond to different PSF shapes and orientations with respect to the point source trajectory (compare to Fig. 4). The signal-to-noise ratio per pixel at the peak varies from left to right: SNR = 1, 3, 5, 10, 30.

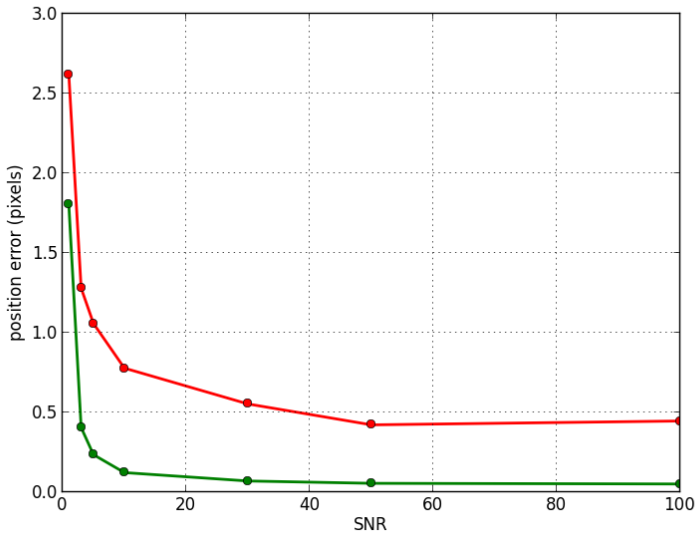

Fig. 8. Simulations of the astrometric accuracy from detailed modeling of streak profiles (green) compared to a simple estimate based on half-max points along the trace (red). The standard deviation (per coordinate) of the measured pixel position for a single streak end point around its true position is shown as a function of the signal-to-noise ratio in the peak pixel.

Unlike the computation of a 1D profile along the track, the 2D algorithm does not require any interpolation of the input data. Another advantage is a straightforward implementation that naturally supports modeling of arbitrarily short streaks. The algorithm is very robust, as long as we can provide a good PSF model tailored to the location within the wide-field image and a reasonable first guess for the extent of the streak. Convolving a line interval with the local PSF to estimate the streak profile requires a large number of PSF evaluations and therefore can be computationally expensive. However, this increased complexity is acceptable given that the number of detectable streaks rarely exceeds a dozen in a typical sky surveillance image.

6. REFERENCES

1. Shell, J.R., "Optimizing orbital debris monitoring with optical telescopes", Proc. of the 2010 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS), Wailea, Maui, Hawaii, 2010
2. Wozniak, P., "Robotic and Survey Telescopes", Planets, Stars, and Stellar Systems, Vol. 1, p. 45, 2013
3. Vestrand, T. et al., "Autonomous global sky surveillance with real-time robotic follow-up: Night Sky Awareness through Thinking Telescopes Technology", Proc. of the 2008 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS), Wailea, Maui, Hawaii, 2008
4. Wozniak, P., "Crowded Field Photometry and Difference Imaging", Proc. of the Manchester Microlensing Conference: The 12th International Conference and ANGLES Microlensing Workshop, 2008
5. Alard, C., & Lupton, R.H., "A Method for Optimal Image Subtraction", Astrophysical Journal, Vol. 503, p. 325-331, 1998
6. Alard, C., "Image subtraction using a space-varying kernel", Astronomy & Astrophysics, Vol. 144, p. 363-370, 2000
7. Bramich, D.M., "A new algorithm for difference image analysis", Monthly Notices of the Royal Astronomical Society, Vol. 386, p. L77–L81, 2008
8. Zackay, B., Ofek, E.O., & Gal-Yam, A., "Proper Image Subtraction—Optimal Transient Detection, Photometry, and Hypothesis Testing", Astrophysical Journal, Vol. 830, p. 1, 2016
9. Schechter, P.L., Mateo, M., & Saha, A., "DOPHOT, a CCD photometry program: Description and tests", Publications of the Astronomical Society of the Pacific, Vol. 105, p. 1342-1353, 1993
10. Stetson, P.B., "DAOPHOT - A computer program for crowded-field stellar photometry", Publications of the Astronomical Society of the Pacific, Vol. 99, p. 191-222, 1987
11. Bertin, E. & Arnouts, S., "SExtractor: Software for source extraction", Astronomy & Astrophysics, Vol. 117, p. 393–404, 1996