Event-based sensor multiple hypothesis tracker for space domain awareness

Rachel Oliver
Cornell University, 124 Hoy Rd, Ithaca, NY 14850
Brian McReynolds
ETH Zürich, Winterthurerstrasse 190, 8057 Zürich, Switzerland

Dmitry Savransky
Cornell University, 124 Hoy Rd, Ithaca, NY 14850

ABSTRACT

Event-based cameras are an emerging technology that measure changes in a field of view by recording data for each pixel only when a thresholded change in log-light intensity occurs. We leverage the low volume of data produced to develop a multiple hypothesis tracker (MHT) that processes time-series data as opposed to traditional image frames with the intent to augment Space Domain Awareness (SDA) tracking operations. We implement measurement rejection based on data attributes extracted from a training set of real event-based data. Our MHT with the data reduction hypothesis test implemented identifies satellite tracks in 92.1% of the test data sets. Addition of a simple noise filter, limiting the hypothesis test evaluation and subsequent track generation to higher frequency events, increases the processing speed by 69.5%. However, the noise filter rejects more real signal information prior to the hypothesis test which results in a decrease to 76.3% of the data sets producing satellite tracks. The initial success of the MHT warrants further development of this promising technique to enable event-based sensing SDA operations.

1. INTRODUCTION

The process of maintaining awareness in space is growing in complexity with the exponential growth of trackable objects [9, 7, 16]. Increased congestion yields increased risk of conjunctions between objects in orbit. This risk is managed with satellite state updates by means of external measurements. Increasing the frequency of measurements should lower the conjunction risk; however, building additional radar sensing locations to increase measurement frequency has proven to be an expensive process [12]. Alternatively, we can explore cheaper, novel sensors, such as event-based cameras, to build a network of measurement sites to augment current Space Domain Awareness (SDA) operations [2]. In this paper, we evaluate the potential of event-based sensors to identify and track resident space objects (RSOs) with a multiple hypothesis tracker (MHT) filtered with statistics derived from real sensor data. Demonstrated success of this tracking ability will ultimately enable an event-based camera network.

In addition to their relatively low cost, event-based cameras offer additional benefits to SDA sensing. The event-based cameras differ from traditional electro-optical sensors because each pixel in the focal array operates independently and asynchronously. When the incident illumination on a pixel changes by a set threshold (corresponding roughly to a fractional change in brightness), the pixel records an event. The resulting output format is known as address event representation (AER), in which each event is identified by the pixel's location on the focal plane (x,y), the timestamp (t) with typically μsec resolution, and the polarity (p) indicating an increase or decrease in the pixel's illumination level. This sensing paradigm facilitates both high dynamic range and temporal resolution [11]. These advantages may expand the range of viewing angles between the observer, object, and the Sun and provide rate information for fast-moving objects such as those in low Earth orbit. In addition, a limited volume of data is collected due to the change detection sensing paradigm and the relatively consistent background of the night sky. The sparse output and reduced power consumption make these sensors ideal for a space-based SDA constellation where on-board computational power and downlink bandwidth is limited. These promising characteristics have inspired demonstrations of the sensors for SDA [19, 6, 15].

However, for the same reason that it is novel, this different sensing paradigm should not simply replace traditional electro-optical systems in classical recognition and tracking algorithms. Reassembling traditional frames to use with standard processing methods negates the advantage provided by limited data production and adds computational complexity. Ideally, the event-based data will be used to update satellite and star tracks in its disaggregated form. The difficulty in using these sensors for online state updates is that they are inherently noisy, making it difficult to distinguish signals from true satellites, stars, and noise in the unaltered data stream. Hypothesis testing provides a means of measurement rejection by assessing the probability that a certain measurement is a true detection. However, multiple trackable objects can be in the field of view at the same time. Therefore, each detection that satisfies the hypothesis test may not belong to the same track. The SDA community has implemented MHTs using both unfiltered and probabilistically filtered event-based data to highlight satellite features [1, 5]. The latter assumes a Poisson distribution for the likelihood of each event being a true detection without consideration of other characteristics that can be derived from the data stream.

We theorize that inference of the signal's origin can be informed by distributions of attributes in the data instead of assuming a particular distribution *a priori* for the likelihood of a detection being from a satellite. For example, a satellite signal produces a series of positive and negative polarity events as it transits a pixel, which likely differs from the event pattern produced by a star. Section 2.1 covers the development of a process to establish a prior distribution and choose relevant attributes. Then we apply a joint conditional probability filter to the batch data in Section 2.2.2 to find a threshold for the hypothesis test. Next, we outline the application of the threshold in a MHT framework in Section 2.3. Sections 3 and 4 cover the performance results and make recommendations for further development.

2. METHODOLOGY

There are three main steps to reach a functional probability based MHT. First, collected data are processed to serve as prior knowledge about the detection probability and its corresponding characteristics. Then the hypothesis test threshold is developed to reject data that are unlikely to be true detection events. The threshold is refined on full batches of data. Finally, the MHT algorithm is written to generate multiple hypotheses of tracks in the focal frame. The MHT keeps track of the probability of the alternative hypothesis associated with each pixel. This probability updates continuously as more information is collected. Once enough data reach the hypothesis threshold, a linear estimate for the next pixel is created and compared to the subsequent observations.

2.1 Data Processing

In order to calculate a probability of detection given certain data attributes, such as a list of events, prior data sets are classified into true positives and negatives of the alternative hypothesis. The alternate hypothesis, the rare case, is a satellite detection. While this sounds generally straightforward, it is impractical to manually sift through thousands of events per data set to label each as a true positive or negative. Therefore, we combine a clustering algorithm and a line drawing algorithm to group the events together and subsequently gather information about specific attributes.

The data processed and shown in this paper was collected with a Prophesee Gen3 camera with a 85mm, f/1.4, lens in Albuquerque, New Mexico from January to February 2021[1]. The mount was set to a series of fixed azimuth and altitude values. Each data set contains roughly 22 seconds of events after approximately 8 seconds are removed due to high noise rates from the background change after the mount adjustment. Both satellite and star streaks are present in the data because the sensor's aim-point is rotating with the Earth.

2.1.1 Attribute Identification

An example data set is depicted in a reassembled frame over an approximately 22 second collection in Figure 1. Polarity is expressed as blue for positive events and red for negative. From visual inspection, the streak in the upper right hand corner is oriented in a different direction than those in the middle of the field of view. This is the satellite, whereas the other lines are stars. Figure 1 provides some insight into what attributes distinguish a star from a satellite detection. The individual pixel event profiles of the stars almost uniformly end in a negative event, while some of the stimulated satellite pixels end in blue indicating a positive event. Therefore, the ordered event profile can help discriminate between the two types of detected RSOs. The figure also shows the density of the star events leaves little

[1]Provided courtesy of Dr. David Monet of the Air Force Research Laboratory Space Vehicles Directorate.

or no gaps between the pixels whereas the satellite track has gaps. The satellite may have a different flux and motion with respect to the sensor than the stars. Therefore, the rate of travel and distance between events in a track also differentiate the satellite from the star pixels. While it cannot be perceived in the two-dimensional format displayed in this figure, the duration a pixel receives information from a satellite is also a discriminating characteristic.

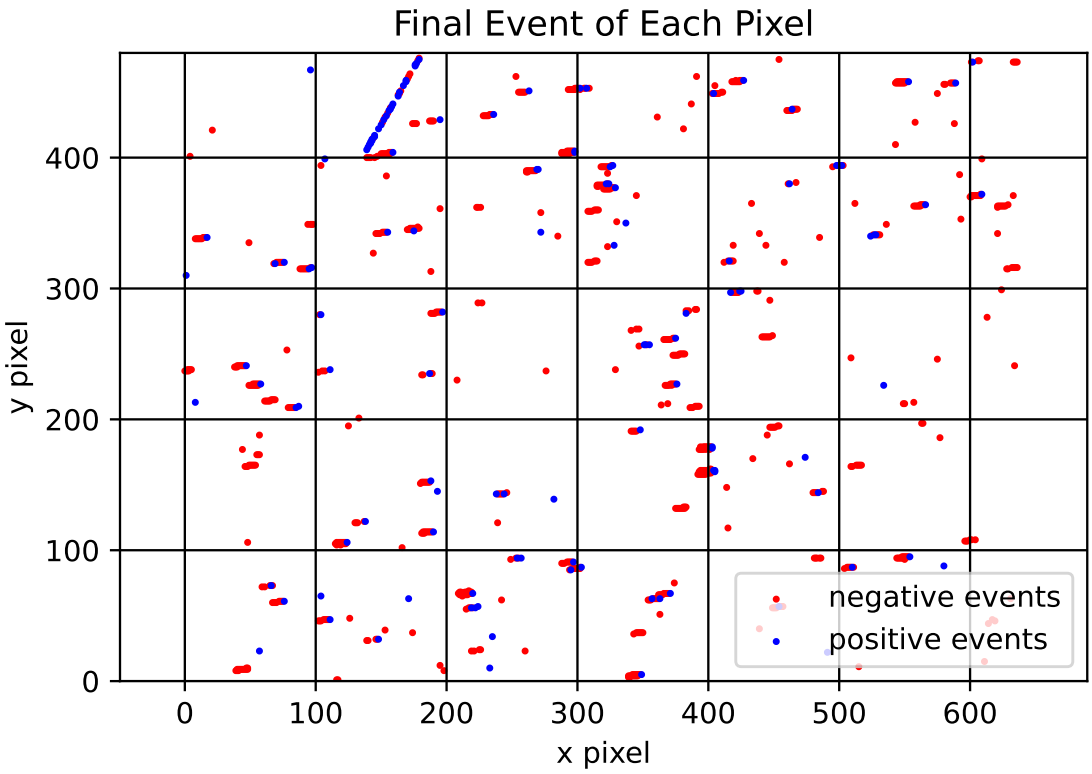

Fig. 1: Example data set reassembled into a typical frame. Positive events are blue and negative events are red. A satellite or star signal typically provides more photometric current than the background. Therefore, the first events are typically positive and the final events are typically negative. This graphic shows the last event on a pixel and therefore, is mainly red. The satellite track in the upper left-hand corner goes in a different direction from the stars.

The other events are presumed to be from hot pixels and noise. The previously derived attributes are sufficient to discriminate these types of events. Hot pixels produce temporally long strings of events that surpass the length of all other types of detected events. Noise events, on the other hand, have a lower frequency than those from true signals. Neither of these types of detected signals traverse to multiple pixels sequentially, so they are easily identified when spatially isolated. Invoking the distinguishing qualities, we create discrete conditional probability tables relating the probability of a type of detection given each attribute: the event profile, average time between events on a pixel, total time a pixel is active, average distance to the closest event, and rate that new pixels are added to the track. The creation of these tables are further described in Section 2.2.1.

2.1.2 Clustering

Each data set has many thousands of events which would be tedious to manually attribute to a type of detected signal. Instead of assigning each event as a true positive or negative, we leverage three-dimensional clustering algorithms to separate the data into groups. Once clustered, the assignment of true detections is equivalent to the number of satellites in the field of view, usually one or two per data set. The two clustering methods described below provide distinct advantages to the user when separating the data and are, therefore, offered as a choice to the user.

The first clustering method offered is the density-based spatial clustering of application with noise (DBSCAN) algorithm. This algorithm groups regions of high density information which makes it adept at grouping clusters of arbitrary shape and size. The algorithm calculates the number of events, N_e, in the data set of events, D, within a radius, r, of each event

$$N_e = \{e \in D < r\}. \tag{1}$$

If the number of events is greater than a minimum value, $N_e > N_{min}$, then it is assigned as a core point [10, 13]. The events within the radius are associated with the core point, but are not necessarily a core point unless they also meet the minimum event requirement. The result is a cluster that is locally bound by a radius, but the full cluster of dense points can take on any shape. We apply the algorithm in three dimensions by scaling the time dimension so that a single radius bounds all three dimensions appropriately.

The DBSCAN algorithm works well when the data have adequate low-density areas in all three dimensions. However, it only takes one dimension with insufficient separation to combine two unrelated signals. One common error results from the spherical bounds extending past an event being checked as a core point to future time-stamped events. For example, consider an event at the end of one cluster near another cluster in space but later in time. If the nearby, but later, cluster provides enough events within the spherical radius to meet the minimum event requirement, the non-core point will be assigned as a core point to merge the two clusters.

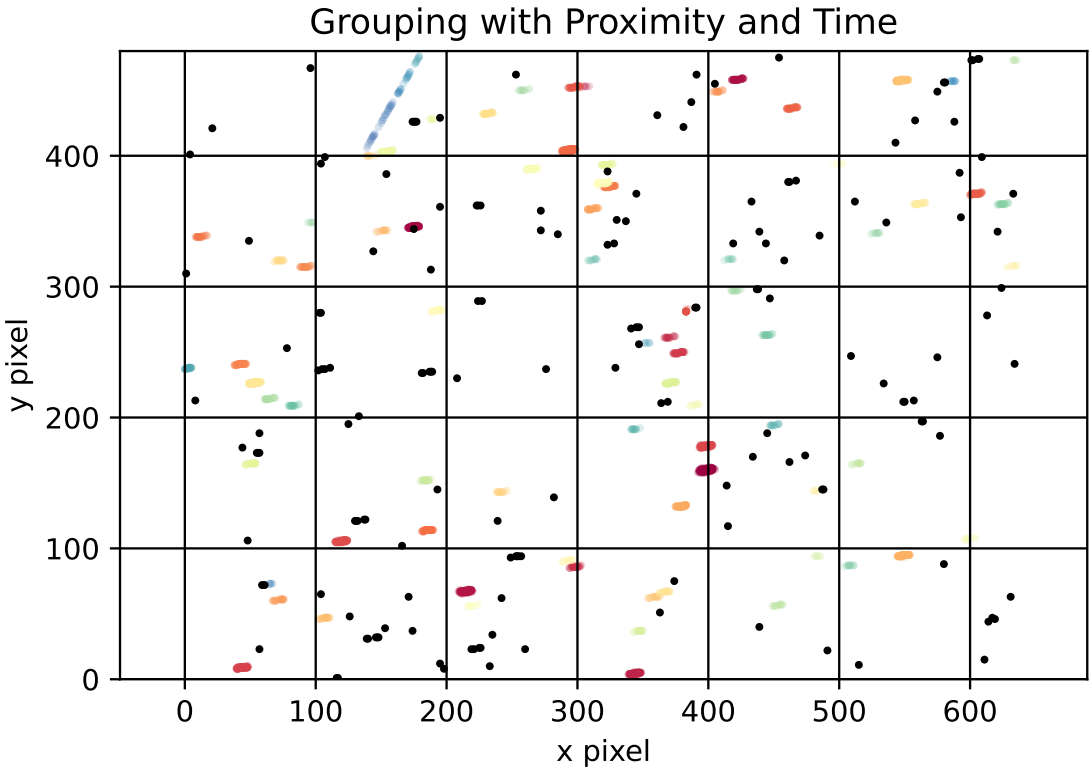

Fig. 2: The final cluster output from the clustering algorithm mapped to different colors for each cluster. The black dots are unclustered events assigned to noise. The separation in the satellite track creates multiple large clusters.

The second clustering method leverages the sequential nature of the time series data to eliminate the aforementioned error. Instead of evaluating density of events within a spherical radius, the events are clustered only considering the events that come before them in the time-series within a radius in the spatial dimension. The region considered is a cylinder as opposed to the spherical volume of the DBSCAN algorithm. In addition, each event is only clustered to the closest event within the cylindrical volume to prevent high noise data from being clustered together. The bounds must still be constrained to prevent clusters being formed without a reasonably nearby event in space and time. Compared

to DBSCAN, this method is more resilient to data sets with higher levels of noise where the regions of low density are insufficient. However, the algorithm is slower to implement due to its sequential nature.

Figure 2 depicts the cluster labels, each with a different color, produced by the clustering algorithm. In this figure, the clustering appears fairly successful with only a few noise events (shown in black) that have a slope similar to the star tracks. Despite the demonstrated success, both algorithms have spatial limitations. In some of these data sets the distance between different star tracks, noise, and satellite tracks is less than the distance between points that should be in the same cluster. This either results in multiple clusters that should be together or multiple tracks combining into one cluster. These two errors are addressed by the introduction of the Hough transform to detect lines within the data.

2.1.3 Line Formation

Due to the short collection period and small field of view, the true satellite and star detections in the analyzed data set appear linear on the focal plane projection as seen in Figure 1. We leverage the linearity of the data to automatically combine separate but co-linear clusters by way of a Hough transform which finds best fit lines in a data set. Typically a best fit line is defined using the standard equation, $y = mx + b$, where the the slope, m, is undefined for a vertical line. To avoid any singularities, the Hough transform defines lines using an angle between 0 and π, θ, and radius, r, from the origin, which translates the original equation to the form $r = xcos(\theta) + ysin(\theta)$. The points in the data set are transformed to the sinusoidal space of the θ and r plane. A voting method in this plane identifies the most likely lines [8] as selected by a threshold of the peak value.

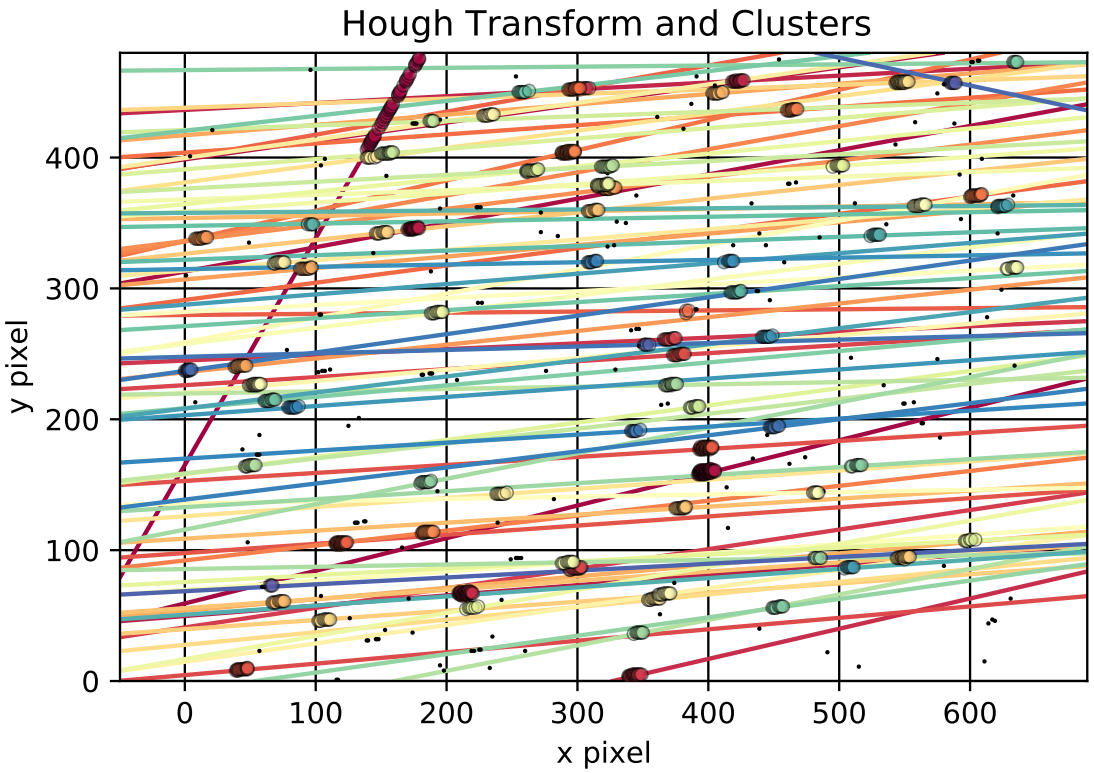

Fig. 3: The cluster output is adjusted with a Hough Transform. Co-linear, but separate clusters of the satellite track in the upper left-hand corner are combined without merging with the two clusters in the vicinity of the location x = 150 and y= 400.

We are not the first to implement the Hough transform technique on event-based data. The transform is applied to extract the linear trend in angles in a related effort. Consistent angles aid in star track identification and, subsequently,

relative rotation is determined by changes in those angles. [3] proves linear trends in event-based data are extractable through the Hough transform. However, we utilize the Hough transform with more caution. Due to the poor signal to noise ratio in some of the data sets, Hough transform lines likely pass through more than true detections. Without additional assessment, these techniques may produce spatially spread clusters with limited regard to the temporal dimension.

Instead, we implement the Hough transform by defining a maximum distance and corresponding maximum time difference along the lines between clusters. This maximum distance is more relaxed than the clustering radius, but reserved enough to minimize the addition of noise events to the final cluster. If two co-linear clusters are within the maximum distance and time, they are combined into one cluster. For example, the two satellite signals in Figure 2 are combined into one cluster in Figure 3. After the co-linear combinations are complete, we assign each cluster to the classes of satellite, star, noise, or hot pixel.

2.2 Hypothesis Test

Once the data is sorted into detection classifications, we develop a hypothesis test to reject unlikely satellite measurements based on the attributes of the processed data sets. We utilize 70% of the processed data as the prior knowledge to develop the hypothesis test. The remaining 30% is used to evaluate the MHT. We then convert the training data into discrete conditional probability tables to apply in the evaluation of the joint conditional probability. Finally, we choose a threshold on the joint conditional probability by testing the filtering on a batch data set.

2.2.1 Probability Tables

Each data attribute discussed in Section 2.1.1 provides information to infer the possibility the data comes from a satellite. In order to leverage multiple data attributes, we employ Bayes theorem, which defines the conditional probability of A given B as

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}. \tag{2}$$

The probability of B can be evaluated as

$$P(B) = P(B|A)P(A) + P(B|\sim A)P(\sim A) \tag{3}$$

where $\sim A$ is the possibility other than A. This can be further expanded for multiple pieces of information using

$$P(A|B \wedge C) = \frac{P(B|A)P(C|A)P(A)}{P(B|A)P(C|A)P(A) + P(B|\sim A)P(C|\sim A)P(\sim A)} \tag{4}$$

with the assumption that the information in B and C are independent [18]. Therefore, we can calculate the probability that the detection is a satellite, the alternative hypothesis, with the liberal assumption that our chosen data characteristics are independent.

In order to implement the joint conditional probability, we must first generate conditional probability tables from the clustered data classifications. Given that each pixel is assigned a class, counts of each class are discretely organized by the values of the attributes. The continuous attributes, time and distance, are separated into discrete bins to enable this organization. After binning, each class of detected events in an attribute is summed and the binned totals in each class are divided by the total of that class. The resulting table represents the probability of each attribute's profile or binned values for a given class which is the conditional probability $P(B|A)$. For example, Table 3 in Appendix B reports the conditional probability of the four classes of detections given the pixel profile derived from the data set analyzed in this paper.

2.2.2 Batch Hypothesis Filter

The next step in the hypothesis test development is to choose an appropriate decision threshold. We inform our choice of a threshold by employing the joint conditional probability filter on the clustered data sets. Employed on this batch form of the data, the filter calculates the probabilities with the greatest amount of information available in each data set. This method provides a quick way to estimate final performance of the measurement rejection. Lower hypothesis test

thresholds that isolate the satellite track in the batch method increase the chance that the online probability calculation will reach the targeted threshold before the full 30 seconds of data are available. We accept the lower threshold may also let true negatives, star tracks, noise, or hot pixels, register as satellites in the early data sets. We select the threshold value from inspection of the batch joint conditional hypothesis filter plots that balance the rejection of true negative measurements and retention of the true positive measurements.

Figure 4a depicts an implementation of the batch data filter that only considers the conditional probability due to the event profile. Because events separate from the satellite track have the same polarity profiles as some of the satellite data, the probability of the true positives and true negatives overlap. The threshold depicted is the maximum for any data to pass the hypothesis test. Therefore, it is not possible to completely remove star detections with the conditional probability hypothesis test that considers only the event profile on a pixel.

Consideration of multiple data attributes reduces the probability overlap, hence the improvement from Figures 4a to 4b. However, each attribute's classifications are not unique between the true positives and negatives and, therefore, some overlap still exists. As a result, the joint conditional probability in Figure 4b does not have perfect performance in terms of null hypothesis measurement rejection.

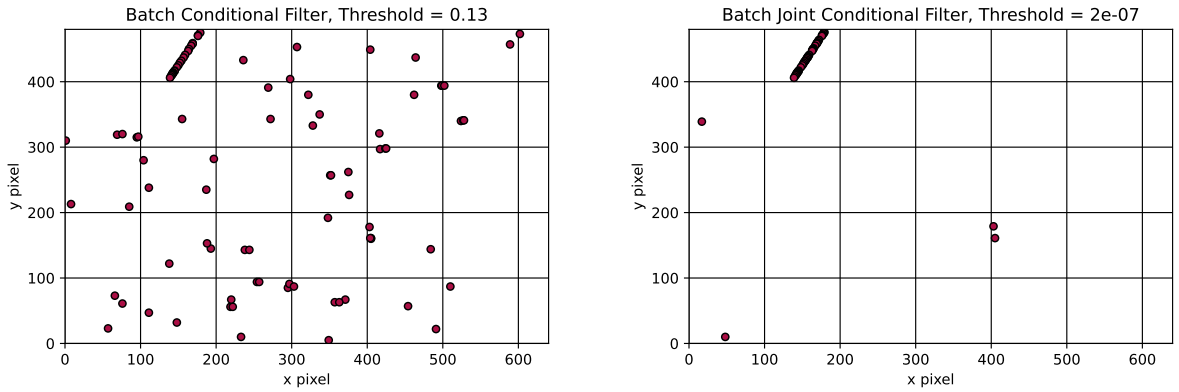

(a) Threshold = 0.13

(b) Threshold = 2e-7

Fig. 4: a) Conditional probability based only on the event profile. This is the maximum threshold that allows data past the hypothesis test. b) joint conditional probability considering other data factors reduces false positives in the data association.

2.3 Multiple Hypothesis Tracker

The foundational idea of MHTs is to evaluate multiple hypotheses in tandem until enough data are available to infer which hypotheses are true positives and which are true negatives [14]. A traditional MHT has four steps when a new dataset is available: form new clusters, form new sets of hypotheses, reduce the number of hypotheses by elimination or combination, and segregate confirmed tracks that have a probability of unity [17]. There are several ways to organize MHT hypotheses in a multi-target and multi-measurement process. Measurement-oriented hypotheses list every possible target for each measurement. Target-oriented hypotheses list every possible measurement for each target [17]. Alternatively, a Bayesian method employs Gaussian mixtures to build probability density functions of the joint distribution of the targets being tracked [4].

To simplify the MHT construction for this paper, we take a target-oriented hypotheses approach, but not every measurement in the field of view is applied to each target hypothesis. As part of the clustering portion of the MHT process, Algorithm 1 applies an initial Euclidean distance threshold to the most recent track pixels before assigning a measurement to the closest hypothesis. In more complex algorithms, the measurements should be assigned to all targets within a radius and only removed once the measurement has probability of unity association with one of the tracks. Our approach in Algorithm 1 reduces the need to aggressively prune hypotheses, but also degrades the algorithm's robustness by not representing all possible data associations.

Algorithm 1 Track Creation

Require: $dist_{max}$
  while $t \leq t_{max}$ **do**
    $pixel = (x[t], y[t])$
    if $t == 0$ **then**
      H_1 {New Hypothesis}
    else
      for $pixel_{previous}$ **do**
        $dist = pixel - pixel_{previous}$
      end for
      if $min(dist) \leq dist_{max}$ **then**
        $H_{min(dist)}$ {Add to Hypothesis}
      else
        H_{new} {New Hypothesis}
      end if
    end if
  end while

Algorithm 2 Probability Update

Require: $prob_{minimum}$
  for $pixel_{H_{current}}$ **do**
    $prob_{current}$ {Run Joint-Cond Prob}
    if $prob_{current} \geq prob_{minimum}$ **then**
      $count = count + 1$ {Threshold Count}
    end if
  end for
  if $count \geq count_{previous}$ **then**
    for $pixel_{new}$ **do**
      $track_{update} = pixel_{new}$ {Update Track}
      $error_{estimate} = pixel_{new} - est_{previous}$ {Estimate Error}
      $y = m * x + b$ {Fit Data to Line}
      $est_{new} = m * x_{diff} + b$ {Estimate Next}
      $count_{previous} = count$ {Update Count}
    end for
  end if

After the clustering is complete, the track hypothesis with new data enters Algorithm 2. In Algorithm 2 we implement the joint conditional probability based hypothesis test developed on the batch information. We evaluate each pixel in the track during the update step because the spatial attributes of the track are modified with the new information. A running count of pixels satisfying the hypothesis test is maintained. A track update is triggered when the count increases in value. The pixels above the threshold are considered part of the possible track and are added to a track list. Essentially, the hypothesis test applies measurement rejection to the incoming data until the confidence in the information is high enough to form or add to a track hypothesis. From the pixels in the track list, an estimate of the next pixel is obtained by regressing a line and using the average Manhattan distance in the x direction to estimate the next y value. We prune the MHT after the probability update for tracks that have not received updating information within a chosen time frame and all of their pixels are below the probability threshold. Pruning reduces the size of the maintained set of track hypotheses and is intended to remove tracks built of true negative detections.

We implement two versions of the MHT outlined above, one with and one without an additional noise filter prior to the clustering step. The simple noise filter leverages the higher frequency of events on the detector from satellite and star signals by comparing the previous event timestamp on a pixel to the current information. If an event occurs before the maximum time threshold from the previous event, it is passed to Algorithm 1. Otherwise the event's timestamp is saved and the algorithm moves onto the next event. The first event on a pixel to pass the filter also retroactively passes the previous event on that pixel to ensure its inclusion in the hypothesis test.

2.3.1 Mutliple Hypothesis Tracker Evaluation

We evaluate our MHT's ability to remove star, noise, and hot pixel information while simultaneously identifying a satellite track by calculating the online hypothesis true positive and true negative rates, the sensitivity and specificity, respectively. We cluster and associate the test data with a detection class using the process outlined in Section 2.1. Then, we compare the assigned detection class against the online assignment. We record a true positive, TP, for each event assigned to the alternative hypothesis and identified as a satellite detection in the clustered data. Likewise, true negatives, TN, correspond to the online null hypothesis and a star, noise, or hot pixel in the clustered data. We calculate the sensitivity, TPR, as

$$TPR = \frac{TP}{P} = \frac{TP}{TP + FN} \tag{5}$$

where the total positive assignments, P, to the alternative hypothesis is comprised of the total true positives, TP, and total false negatives, FN, at each simulation step. Correspondingly, we calculate the specificity, TNR, as

$$TNR = \frac{TN}{N} = \frac{TN}{TN + FP} \tag{6}$$

where the total assignments to the null hypothesis, N, is comprised of the total true negatives, TN, and total false positives, FP.

3. RESULTS

The multiple hypothesis tracker, with and without the initial noise filter, proves successful at rejecting many star detections and their resulting tracks for some data sets. For example, the final time for one data set of the MHT is depicted in Figures 5a and 5b for the non-filtered and filtered versions respectively. In these figures, the estimated track is depicted in blue and the thresholded events are plotted in red. Comparing against the clustered data in Figure 3, the MHT without the additional noise filter covers the full spatial extent of the original track cluster. The noise filtered MHT, on the other hand, is missing portions of of the track. The missing events are a direct consequence of the noise filter. These events did not have a corresponding second event and were never passed to the hypothesis tracker. Despite the apparent downside to the noise filter, the rejection of single events also reduces the hypothesis tracks comprised of star signals. This overall reduction still allows the satellite track to be identified from more active pixels, but with fewer false track hypothesis from stars.

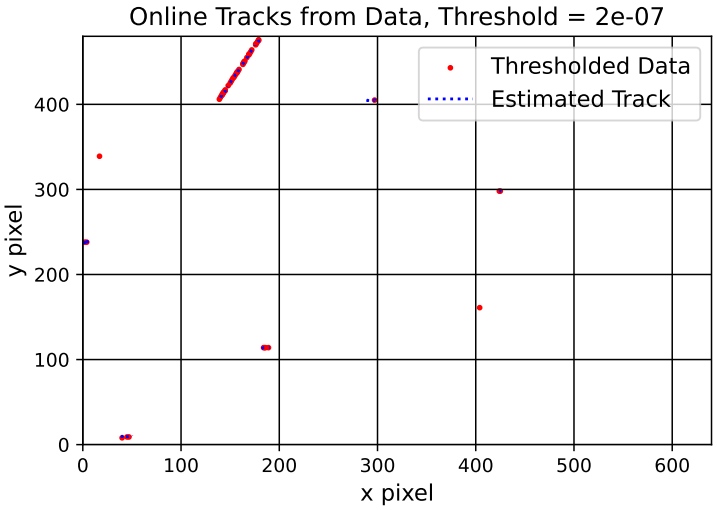

(a) MHT without filter

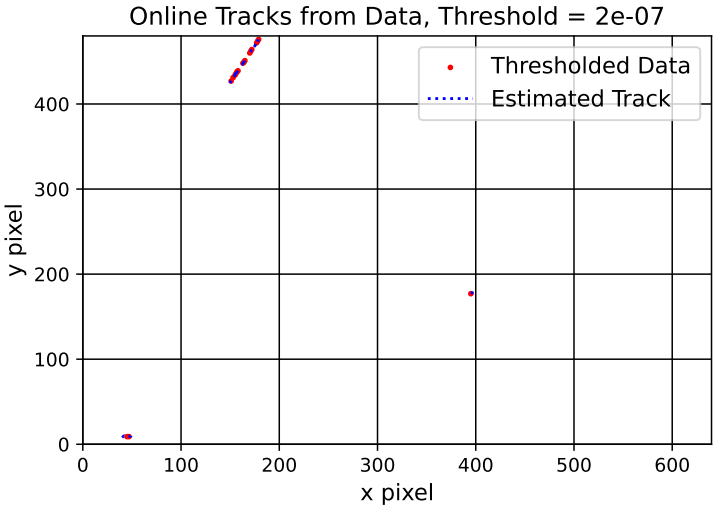

(b) MHT with noise filter

Fig. 5: a) The MHT at the final time without the initial filter removing lower frequency events has seven star tracks. b) The MHT at the final time with the initial filter has two star tracks, but portions of the satellite track events were also removed by the filter.

Gaps in the satellite track hypothesis, identified by discontinuities in the estimated track, are another artifact of the current MHT implementation, visible in both panels of Figure 5. Similar spatial and temporal boundaries are applied in the MHT clustering and the offline clustering, described in Section 2.1.2, to prevent combining distinct non-satellite clusters. As a result, the final track is not cohesive despite all sections of the track being identified in the non-filtered MHT. One possible solution is to relax the clustering bounds in the region surrounding the linear estimate from the last pixel added to the track hypothesis, but retain tighter bounds for tracks without events above the hypothesis test threshold.

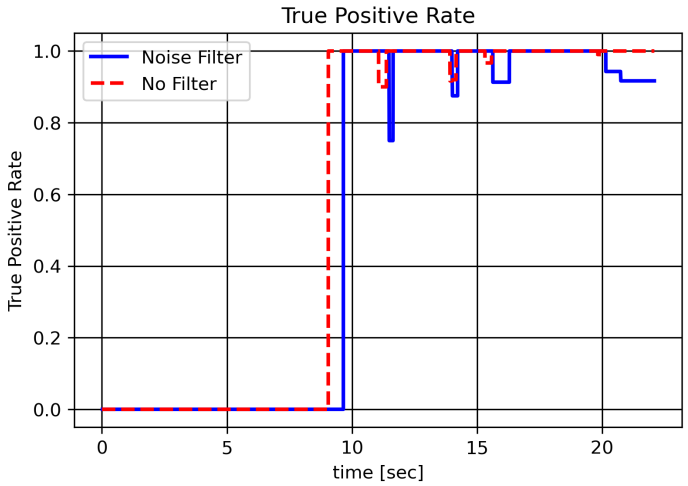

Fig. 6: The sensitivity of the hypothesis test updated over the course of the MHT runs.

We derive more insight on the hypothesis track production through examination of the true positive and true negative ratios. The sensitivity over the course of the runs of the MHT, depicted in Figure 6, reflect the possible time required to accrue enough information at the track level to evaluate a set of events as the alternative hypothesis. Signals with higher frequencies and more events will reach the threshold more quickly. While poor signal to noise ratios may require longer observations to have success with the hypothesis test. The subsequent drops in the true positive rate after the initial rise are due to the addition of more true satellite events to the hypothesis test. These events require additional information prior to association with the alternative hypothesis, especially if clustered separately from the rest of the track. The drops in the sensitivity are greater for the filtered MHT because fewer events, including those associated with the satellite designation, reach the hypothesis test stage of the algorithm.

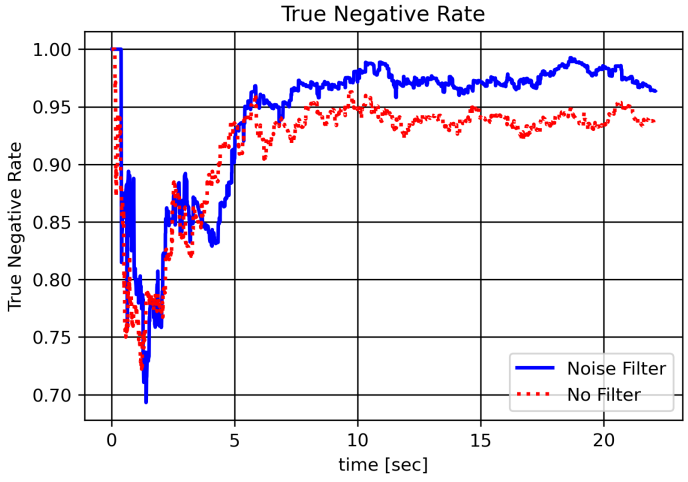

Fig. 7: The specificity of the hypothesis test updated over the course of the MHT runs.

Figure 7 shows the specificity over the course of the MHT runs. The initial drop in the true negative rate demonstrates how the initial clusters of events derived from star signals have similar attributes to events derived from satellite signals. As the star signals persist, their distinguishing profile of events and consistent temporal dynamics—discussed in Section 2.1.1—drop the joint conditional probability below the selected threshold. The true negative rate stabilizes with additional events having a minimal effect on the ratio. We attribute the improved true negative rate in the noise filtered MHT to the removal of events that do contribute to the classic star signal characteristics. The remaining data has a higher rate of evaluating below the hypothesis threshold and, therefore, fewer erroneous satellite tracks are produced.

The noise filtered version of the MHT demonstrates its value when applied to noisier data sets as demonstrated by the reduction of track hypothesis between Figures 8a and 8b. As long as the satellite signal produces pixel profiles with greater than one event, it survives the filter. The result is a reduction of track hypotheses continuing to be evaluated and maintained by the multiple hypothesis tracker and an increase in computational speed.

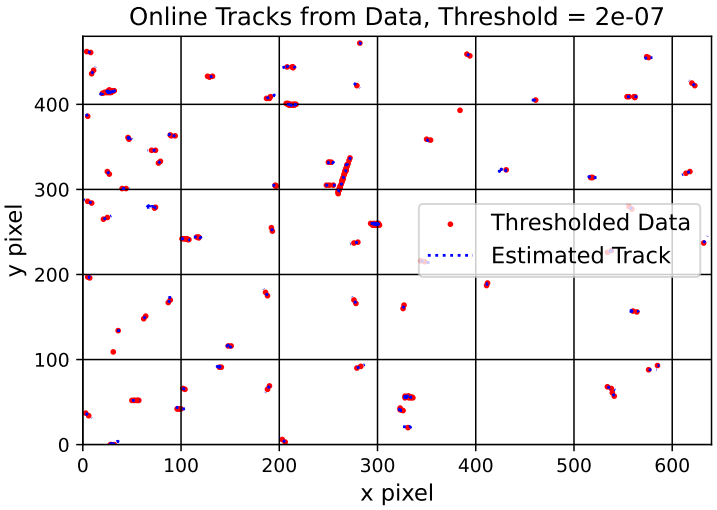

(a) Noisy data MHT without filter

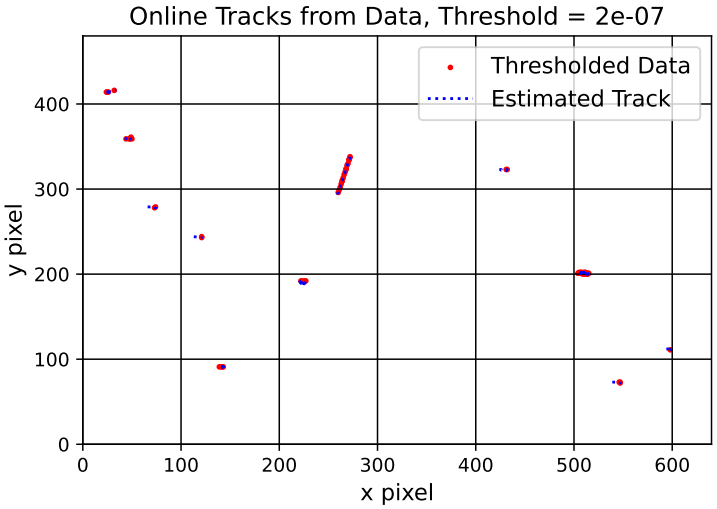

(b) Noisy data MHT with filter

Fig. 8: a) The noisier data set produces more track hypotheses which will have to be subsequently evaluated. b) When the signal to noise ratio of the satellite is favorable, the preliminary noise filter effectively reduces the track hypotheses.

4. CONCLUSIONS

We construct a MHT to process event-based time series data into satellite track hypotheses through implementation of a measurement rejecting joint conditional hypothesis test. We derive the hypothesis test from attributes in the data that we obtain by clustering then attributing to satellite, star, or noise detections. The MHT performs consistently on the majority of the tested data sets. The unfiltered version identifies 92.1% of the satellite tracks during the course of the algorithm. The noise filtered MHT is less successful with a rate of 76.3% of satellite track identifications. We attribute the noise filtered version's decreased success rate on the data sets where large portions of the satellite data has a single event profile. The remaining events are too sparse to assemble a track hypothesis. However, with an appropriate satellite signal to noise ratio, the noise filter provides a reduction to additional track hypotheses and an average 69.5% decrease in the computation time when run on the same hardware. Our goal is to progress the MHT towards online data association and track hypothesis maintenance. Therefore, reducing the computation time is a priority for future refinements. Future work will also consider more sophisticated noise reduction techniques to improve performance with a filter when the signal to noise ratio is unfavorable.

The success of the MHT is also dependent on the hypothesis test sensitivity and specificity. The average final true positive and negative ratios are summarized in Table 1. The hypothesis test is particularly effective at rejecting the null hypothesis measurements with a specificity above 0.9 for both implementations of the MHT. While the average sensitivity is under 0.5, the MHT still had enough information to generate most of the satellite track hypotheses. Ultimately, while our MHT serves as an effective demonstration of applying event-based time series data to SDA, there is significant room for further improvements. We plan to implement alternative data association methods such as support vector machines and maximum likelihood estimators in the MHT as they may yield improved sensitivity and possibly improve processing times. We also intend to utilize track statistics to evaluate the generated track hypotheses and inform more intricate pruning of the hypotheses generated.

Metric	MHT Without Noise Filter	MHT with Noise Filter
Yielded Satellite Tracks	92.1%	76.3%
Average Data Association TPR	0.473	0.393
Average Data Association TNR	0.901	0.972
Average Run Time [sec]	119.7	36.5

Table 1: Noise filtered and non-filtered MHT performance statistics.

5. REFERENCES

[1] S. Afshar, A. Nicholson, A. Schaik, and G. Cohen. Event-based object detection and tracking for space situational awareness. In *Proc. of the 21st Int. Conf. on Information Fusion*. University of Cambridge, IEEE Xplore, 2018.

[2] D. Aiken. Western secures funding for australian-first multi-sensor space observatory. https://www.westernsydney.edu.au/newscentre/news_centre/story_archive/2020/western_secures_funding_for_australian-first_multi-sensor_space_observatory, 2020.

[3] S. Bagchi and T. Chin. Event-based star tracking via multiresolution progressive hough transform. In *Proc. of the IEEE Workshop on Applications of Computer Vision*. IEEE, 2020.

[4] S. Blackman. Multiple hypothesis tracking for multiple target tracking. *IEEE A&E Systems Magazine*, 19(1), January 2004.

[5] B. Cheung, M. Rutten, and G. Cohen. Probabilistic multi hypothesis tracker for an event based sensor. *IEEE Sensors Journal*, 20:15117–15132, December 2020.

[6] G. Cohen, S. Afshar, A. Schaik, A. Wabnitz, T. Bessell, M. Rutten, and B. Morreale. Event-based sensing for space situational awareness. In *Proc. of the 18th Advanced Maui Optical and Space Surveillance Technologies*. Maui Economic Development Board, September 2017.

[7] DoD. United states in space: National security strategy and national space policy. Technical report, Department of Defense, 2011.

[8] R. Duda and P. Hart. Use of the hough transformation to detect lines and curves in pictures. *Communications of the ACM*, 15(1), January 1972.

[9] ESA. Esa's annual space environment report. Technical Report April, ESA Space Debris Office, Darmstadt, Germany, 2022.

[10] M. Ester, H. Kriegel, J. Sander, and X. Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proc. of the Int. Conf. on Knowledge Discovery and Data Mining*. AIAA, 1996.

[11] G. Gallego, T. Delbruck, G. Orchard, C. Bartolozzi, B. Taba, A. Censi, S. Leutenegger, A. Davison, J. Conradt, K. Daniilidis, and D. Scaramuzza. Event-based vision: a survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.

[12] GAO. Gao-15-342sp assessments of selected weapon programs. Technical report, Government Accountability Office, March 2015.

[13] Kamran Khan, Saif Ur Rehman, Kamran Aziz, Simon Fong, and S. Sarasvady. Dbscan: Past, present and future. In *Proc. of the Fifth Int. Conf. on the Applications of Digital Information and Web Technologies (ICADIWT 2014)*, pages 232–238. IEEE, 2014.

[14] C. Kim, F. Li, A. Ciptadi, and J. Rehg. Multiple hypothesis tracking revisited. In *IEEE International Conference on Computer Vision (ICCV)*, 2015.

[15] P. MchMahon-Crabtree and D. Monet. Commercial-off-the-shelf event-based cameras for space surveillance applications. *Applied Optics*, 60(25):G144–G153, September 2021.

[16] Office of the Secretary of Defense. Defense space strategy summary. Technical Report June 2020, Department of Defense, 2020.

[17] D. Reid. An algorithm for tracking multiple targets. *IEEE Transactions on Automatic Control*, 24(6), December 1979.

[18] J. Weisberg. Multiple conditions. https://jonathanweisberg.org/vip/multiple-conditions.html, 2021.

[19] M. Zolnowski, R. Reszelewski, D. Moeys, T. Delbruck, and K. Kaminski. Observational evaluation of event cameras performance in optical space surveillance. In *Proc. of the 1st NEO and Debris Detection Conference*. ESA Space Safety Programme Office, 2019.

A. CLASS PROBABILITY TABLE

Hot Pixel	Noise	Star	Satellite
0.005	0.464	0.506	0.024

Table 2: Probability of a class of detection occurring in the training data set.

B. CONDITIONAL PROBABILITY TABLES

Profile	Hot Pixel	Noise	Star	Satellite
(0)	0.85	0.02	0.07	0.03
(1)	0.05	0.02	0.15	0.25
(1,0)	0.02	0.004	0.04	0.05
(1,1)	0.01	0.02	0.09	0.14
(1,1,1)	0.002	0.02	0.06	0.08
(1,1,1,1)	0.001	0.02	0.04	0.05
(1,1,1,1,1)	0.001	0.02	0.03	0.03
(1,1,1,1,1,1)	0.001	0.02	0.02	0.02
(1,1,0)	0.003	0.001	0.03	0.05
(1,1,0,0)	0.0008	0.0002	0.02	0.02
(1,1,1,0)	0.0005	0.0003	0.02	0.03
(1,1,1,0,0)	0.0003	0.00008	0.02	0.03
(1,1,1,1,0)	0.0002	0.0002	0.01	0.02
(1,1,1,1,0,0)	0.0002	0.0002	0.01	0.02
(1,1,1,1,0,0,0)	0.0001	0.0001	0.009	0.01
(1,1,1,1,1,0)	0.0001	0.00008	0.008	0.01
(1,1,1,1,1,0,0)	0.0001	0.00008	0.008	0.01
(1,1,1,1,1,0,0,0)	0.0001	0.00008	0.008	0.01

Table 3: Profiles that have probabilities greater than 0.01 given a satellite detection has occurred. 1 is a positive threshold change. 0 is a negative threshold change