A Novel Method for Satellite Maneuver Prediction

Charlotte Shabarekh
Aptima, Inc.
Jordan Kent-Bryant
Aptima, Inc.
Gene Keselman
LEO Projects, LLC
Andonis Mitidis
Aptima, Inc.

ABSTRACT

A space operations tradecraft consisting of *detect-track-characterize-catalog* is insufficient for maintaining Space Situational Awareness (SSA) as space becomes increasingly congested and contested. In this paper, we apply analytical methodology from the Geospatial-Intelligence (GEOINT) community to a key challenge in SSA: *predicting where and when a satellite may maneuver in the future*. We developed a machine learning approach to probabilistically characterize Patterns of Life (PoL) for geosynchronous (GEO) satellites. PoL are repeatable, predictable behaviors that an object exhibits within a context and is driven by spatio-temporal, relational, environmental and physical constraints. An example of PoL are station-keeping maneuvers in GEO which become generally predictable as the satellite re-positions itself to account for orbital perturbations.

In an earlier publication, we demonstrated the ability to probabilistically predict maneuvers of the Galaxy 15 (NORAD ID: 28884) satellite with high confidence eight days in advance of the actual maneuver. Additionally, we were able to detect deviations from expected PoL within hours of the predicted maneuver [6]. This was done with a custom unsupervised machine learning algorithm, the Interval Similarity Model (ISM), which learns repeating intervals of maneuver patterns from unlabeled historical observations and then predicts future maneuvers. In this paper, we introduce a supervised machine learning algorithm that works in conjunction with the ISM to produce a probabilistic distribution of when future maneuvers will occur. The supervised approach uses a Support Vector Machine (SVM) to process the orbit state whereas the ISM processes the temporal intervals between maneuvers and the physics-based characteristics of the maneuvers. This multiple model approach capitalizes on the mathematical strengths of each respective algorithm while incorporating multiple features and inputs. Initial findings indicate that the combined approach can predict 70% of maneuver times within 3 days of a true maneuver time and 22% of maneuver times within 24 hours of a maneuver. We have also been able to detect deviations from expected maneuver patterns up to a week in advance.

1. PROBLEM

Space Situational Awareness (SSA) tradecraft needs to evolve to keep pace with the increasing number of objects in space. The rate of new satellites in recent years have spiked to almost exponential growth partially due to commercial enterprises who have decreased the cost of launches and made small satellites (small sats) readily available. The LA Times reports that in 2010 there were just 25 small sats launched into space but it is estimated that over 200 will be launched in 2016 [4]. While many of the older small sats did not have maneuvering capabilities, there is a push to add propulsion to the new models for collision avoidance. As such, within the next five years, we can not only expect the total number of objects (both small sats and standard satellites) in space to increase greatly, but we can expect the total number of *maneuverable* objects to increase as well.

From an SSA perspective, having maneuvering capabilities on satellites is important because it can help to reduce conjunctions and collisions. However, from and Indications and Warnings (I&W) perspective, maneuverable objects can pose a challenge because a satellite's track cannot be reliably propagated using orbital dynamics and physics alone. Dynamic events, such as an unanticipated maneuver, could result in a broken track and possibly the introduction of an Uncorrelated Target (UCT). Timeliness is critical for collision avoidance but unexpected maneuvers can result in a time-consuming process of data association to regain the satellite's chain of custody. Therefore, I&W needs to be able to predict when a satellite might maneuver with enough advanced notice to execute a course of action (COA).

2. ACTIVITY ANALYIS APPROACH

Predicting when a satellite will maneuver requires establishing a baseline understanding of how a satellite behaves under specific environmental and spatio-temporal constraints. Within the geo-spatial intelligence (GEOINT) community, there has been a push to move away from analyzing *where* an object is and instead centering analysis about *what* the object is doing [3]. This *activity analysis* approach hinges on the premise that an objects locations and behaviors are constrained by activity that it is performing. Activities have a sequence of events and expected behaviors that are associated with it, which we call Patterns of Life (PoL) [2]. Once an object's activity is inferred, the PoL that they are performing at the current time and future times can be probabilistically estimated. For instance, satellites will perform station-keeping maneuvers to offset gravitational pull and other atmospheric conditions to maintain a fixed position over earth or keep from re-entering the atmosphere. These station-keeping maneuvers are typically repeatable and predictable for a satellite, time of year and location. Although there is a degree of variability of how PoL are executed, they still can serve to narrow the field of possible times that a satellite would most likely be performing maneuvers. In this paper, we continue development on our probabilistic satellite maneuver prediction technology that automatically learns a satellite's PoL to establish baseline or normal maneuver patterns. Then, using these PoL we predict the next maneuver and quickly identify when deviations from this pattern will most likely occur. Deviations are not necessarily threats, but they are unexpected behaviors which can be flagged for space operators to analyze as an early stage of a threat warning and assessment (TWA) system. It is our hypothesis that an activity analysis approach can discover anomalous maneuvers sooner to increase the amount of time possible for a Course of Action (COA) assessment.

3. MODEL DETAILS

In an earlier publication [6], we introduced the Interval Similarity Model (ISM), as an approach to learn PoL. The ISM is an unsupervised machine learning algorithm which effectively clusters temporal intervals based on periods of maneuvers and non-maneuvers. The ISM lends itself well to the challenge of activity analysis because it models the context over a duration of time and re-estimates its models as new context becomes available. It abstracts away from the object being modeled and rather focuses on the observational behavior being executed. However, even with online re-estimation of models, it can take a little while to pick up on new PoL leading to false negative predictions and decreased recall. Additionally, all unsupervised machine learning methods can learn bias in the data. Although we have not experienced this phenomenon on the test data, we are aware of the high degree of noise and bias in operational datasets. Therefore, in this paper, we introduce a widely-used supervised machine learning model, the Support Vector Machine (SVM) [5], to use in conjunction with the ISM. We initially gravitated away from supervised machine learning approaches because ephemeris data and a history of known patterns is not always available, but is required to train a supervised algorithm. However, our dataset did have ephemeris so we decided to capitalize on the powerful classification capabilities that the SVM provides as part of our multiple-model approach (Figure 1). We refer to this multiple model approach as the ISM+SVM combined model and the single model approach as just the ISM model. The SVM processes the orbit state to whereas the ISM processes the temporal intervals between maneuvers and the physics-based characteristics of the maneuvers. The probabilistic estimates of produced by both models are combined into a single probabilistic distribution function (PDF) that predicts the likelihood of maneuvers occurring at time. Details of each model follow in the subsections below.

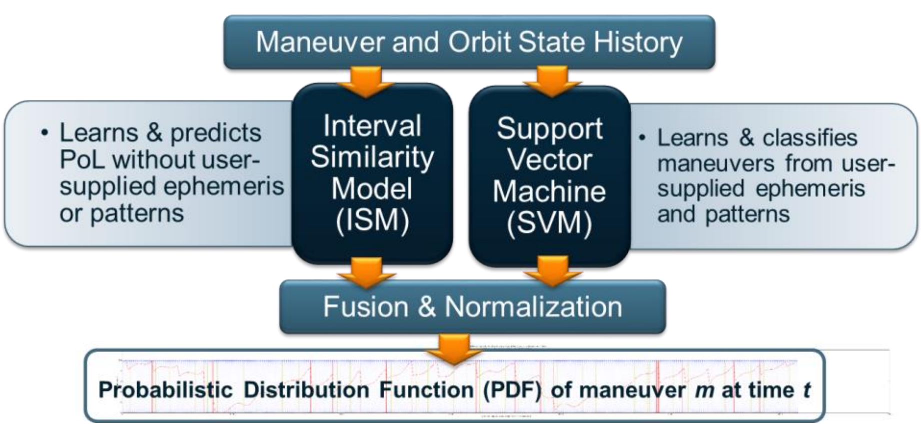

Figure 1: Multiple model approach capitalizes on the mathematical strengths of each algorithm while incorporating multiple features and inputs

Interval Similarity Model (ISM)
ISM effectively calculates the probability that a satellite is executing a pattern of maneuvers that are similar to historical PoL. Inspired by similarity-based clustering [1], ISM's output is a probability density function (PDF) detailing the probability that a maneuver will occur with respect to time. It avoids strict clustering in favor of a probabilistic approach to allow for learning of patterns that can generalize with new observations. ISM populates an interval similarity matrix that connects consecutive intervals strongly or weakly based on the similarity between the two intervals. Ultimately this method produces a matrix estimating the probability that each interval is likely to repeat in the future, and this allows for future prediction of maneuvers. The columns of the similarity matrix are generated one at a time, one per each maneuver. Whenever a new maneuver occurs, it creates a new interval between itself and each maneuver that has occurred previously. We are primarily interested in representing how likely it is that that interval will repeat in the future. An interval is likely to repeat if it is part of a pattern of repeating intervals, and intervals in a repeating pattern are likely to be similar.

Formally, suppose we have labeled maneuvers, *i* through *k*, and intervals *ij*, *ik*, and *jk*. Our approach uses Bayesian probability calculation to estimate likelihood that maneuvers are part of a repeating pattern. Equation 1 computes the similarity of interval length where σ is the estimated standard deviation of the interval *ij* duration, and where interval *ij* and interval *jk* are the durations of those intervals. This is the "similarity" part of the interval similarity model, and is the probability that maneuver *k* at time *t* would be observed when the model was given that interval *ij* and interval *jk* are part of the same pattern. Next, we take this similarity indicator and fit use it to estimate how likely it is that that interval will repeat in the future. This estimate is a Bayesian calculation which includes the similarity between intervals *ij* and *jk*, the initial probability that *ij* would repeat, and adaptive priors.

$$P(k\ occurs\ at\ time\ t \mid jk\ is\ a\ repeat\ of\ ij) = \frac{e^{\frac{-(interval1 - interval2)^2}{2\sigma^2}}}{\sigma\sqrt{2\pi}}$$

(Equation 1)

Once the probability of repeat is estimated for each interval, the next step is to use that information to predict when the next interval will occur. A probability distribution for when the next maneuver will occur is generated for every interval. It is weighted by the interval's calculated likelihood of repeat. All of these predicted and weighted probability distributions are combined into one distribution and combined with the outputs of the SVM in the next subsection.

Support Vector Machine (SVM)
The SVM a supervised machine learning algorithm that is typically used as a classifier to determine if observations belong to a specific category such as a maneuver or non-maneuver class. Unlike the ISM, the SVM does not learn temporal intervals of PoL. Instead, it classifies whether an object will imminently maneuver. The SVM uses a set of labelled ephemeris data to discover a set of separating rules between classes. These rules can then be applied on a set of new unlabeled data points such that we can classify them as points belonging to a class. More specifically, in the SVM formulation any set of data points, is treated as a set of points (or feature vectors) in a multi-dimensional vector space. To discover the separating rules (also known as 'training the algorithm') means to find a separating hypersurface in this multi-dimensional vector space such that all (in practice most) of the points belonging to each class lie on different sides of the hypersurface. For this paper, we used ephemeris data of labeled historical maneuvers and orbit state as feature vectors to train the model. The SVM was trained to recognize pre-maneuver orbit states – that is the orbit state immediately before a maneuver. We hypothesized that since station-keeping maneuvers are designed to return the satellite to a desired orbit state there would be a correlation between orbit state and maneuvers times.

Once trained, the SVM outputs a likelihood that a given orbit state was "pre-maneuver". To use the SVM for prediction we first used an orbit state propagator to predict the future orbit state of the satellite. After each maneuver we used the propagator to calculate the orbit state for every four minutes of the next fifteen days. We then fed the fifteen days of predicted orbit state to the SVM to produce a probability of maneuver vs. time result.

The final step was to fuse the SVM results with the ISM results. This was done by taking the probability of maneuver vs. time produced by both the SVM and the ISM and multiplying them together, then renormalizing the result.

4. DATASET OVERVIEW

Our approach was tested using the maneuver times for the Galaxy 15 (NORAD ID: 28884) geo-synchronous satellite, during a four-year period (2011-2015). The dataset was synthetically generated by the Air Force Research Labs (AFRL) Space Vehicles Directorate and demonstrated realistic levels of collection cadence (up to six days without observations of object) and noise (up to 90 microradians). This data was astrometric and had orbit state. It had four sources, each collecting at a different cadence and from a different earth-based location. In this paper, we worked on a subset of 2013 to 2015 which contained 144 maneuvers. Galaxy 15 was selected for experimentation purposes for three reasons. First, ephemeris data was freely available for use in validation and training of the SVM. Second, it demonstrated Patterns of Life (PoL) when performing station-keeping maneuvers. Lastly, there were maneuvers which did not follow established PoL which could be used as a test case for anomalous maneuver detection. Figure 2 shows PoL in the form of the duration between maneuvers. On average, the satellite maneuvered every 3.5 days, however there were instances where there was as little as 1 day between maneuvers and as much as 19.5 days between maneuvers. In the results section, we demonstrate the ability to alert the user to these anomalies in advance.

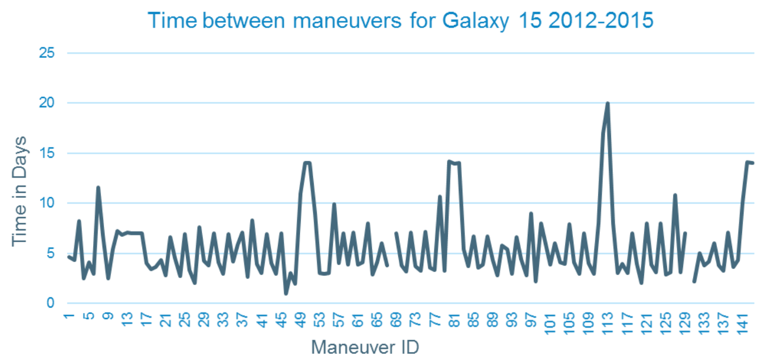

Figure 2: The dataset contained 144 maneuvers performed by Galaxy 15 from January 2013 to May 2015. On average the satellite maneuvered every 3.5 days but there were instances where there was as little as 1 day between maneuvers and as much as 19.5 days between maneuvers.

5. PREDICTION RESULTS

We evaluated our results across five dimensions: *precision*, *recall*, *confidence*, *timeliness* and *likelihood performance*. Our original hypothesis was that the ISM+SVM combined model would outperform the ISM model across board. What we learned is that the combined model performs better under some circumstances and that the ISM model performs better under different circumstances. The precision and recall scores reported below did not have significant differences between the two models, but there were advantages to using the combined approach. In general, the combined model was more confident in its predictions with a higher average probability for correct predictions than the ISM. The ISM, on the other hand, predicts maneuvers further in advance than the combined model. It appears that each model worked to its strengths with the ISM outperforming the combined approach on longer, repeating patterns whereas the combined approach was better able to handle dynamic patterns. This is demonstrated in Figure 3 in which the vertical red line represents the true maneuver time, the blue line represents the predictions of the ISM+SVM combined models and the black line represents the predictions of the ISM model. This color scheme will be consistently used throughout future figures. Frequently, the ISM and combined approaches were in sync (Figure 3, left), but when they were not, it indicated a lower confidence prediction and sometimes a departure from an established PoL (Figure 3, middle and right).

Precision and Recall

Precision and Recall are standard metrics within the classification research area. Precision measures *correctness* while recall measures *completeness*. Precision evaluates how many predicted maneuvers were correct out of the total number of predictions made. Recall evaluates how many predicted maneuver times were correct out of the total number of true maneuver times. A third metric, F-Score is frequently used as a single score to balance out precision and recall. For these dimensions, we scored based on the difference in hours between the predicted maneuver time and the true maneuver time. We scored for each hour up until 24 hours and then again for 36, 48 and capping it at 72 hours.

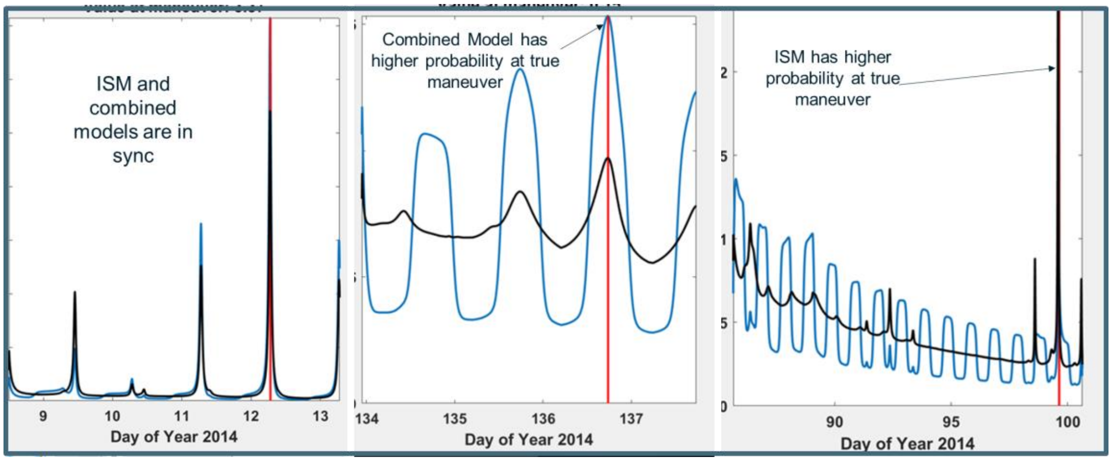

Figure 3: The ISM (black line) and ISM+SVM (blue line) combined models were often in sync (left) when predicting maneuvers (red line) and when they were not in sync, it indicated a lower confidence prediction.

Figure 4 shows the comparative precision for both the ISM only (blue) and ISM+SVM approaches (orange). Along the x-axis is the number of hours separating the predicted maneuver time from the true maneuver time and the y-axis represents probability up to 1. For precision, there was very little difference between the two models. The combined model approach started off strong with a 22% precision over the ISM only approach of 19% for maneuvers predicted within an hour of a true maneuver. In other words, of the total number of correct predictions that the combined model made, 22% of the predicted maneuver times were within 1 hour of a true maneuver time. The ISM rapidly caught up though and by the 24-hour mark, overtook the combined approach slightly. For both approaches, over 40% of the predicted maneuver times within the same day (within hours) of the true maneuver. Over 70% of the predicted maneuver times were within three days of the true maneuver.

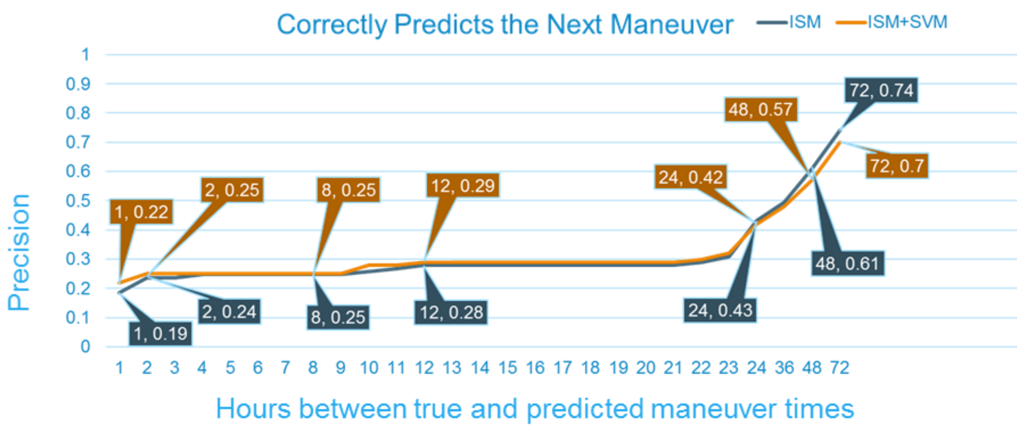

Figure 4: Precision Results for the ISM (blue) and ISM + SVM combined models (orange) were very close. The combined model correctly predicted 22% of maneuver times within one hour of a true maneuver. Both models achieved over 70% prediction of maneuver times within 3 days of a true maneuver.

On the flipside of precision is recall. While the precision scores were fair, the recall for both ISM and ISM+SVM combined approaches were low. Figure 5 shows the recall at each hour starting from a low of 13% of true total maneuvers being predicted within an hour of a true maneuver time by the ISM model only. In other words, out of the total number of maneuvers that actually occurred, only 13% of them were predicted to occur at a time within 1 hour of the true maneuver time. It peaks at 50% of true total maneuvers being predicted within 72 hours of a true maneuver time also by the ISM model. It follows the same trend as observed in precision, where the combined model initially starts with stronger scores, but is overtaken by the ISM only approach by around the 24-hour mark. These low scores are due to the prevalence of false negatives (missed maneuvers) and the probabilistic nature of this approach. Although only the most probable result is returned, the system produces a vector of maneuvers and associated probabilities that can be accessed by adjusting the probability thresholds. Figure 6 shows the recall if the three most probable maneuvers are returned. By returning the top three results, the recall increases to 24% within 1 hour for the combined model and 81% for within 72 hours of a true maneuver.

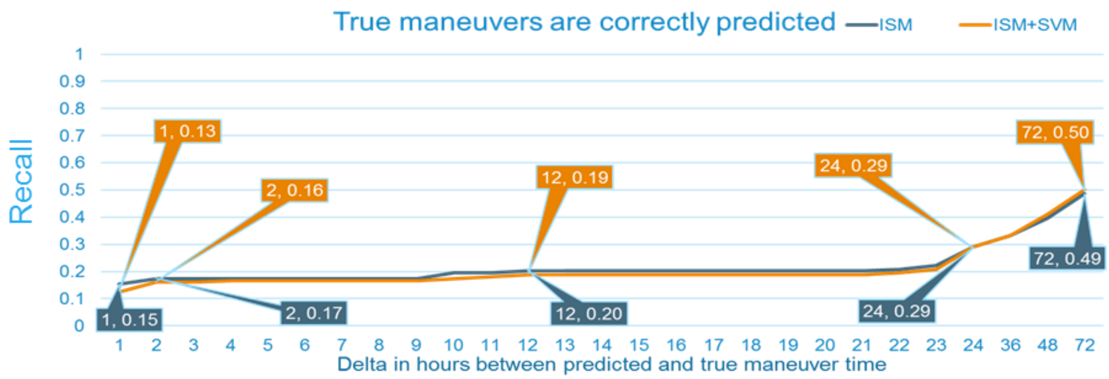

Figure 5: Recall results for the ISM (blue) and ISM+SVM combined (orange) models were very close. The ISM slightly outperformed the combined model by correctly predicting 15% of true maneuver times to occur at a time within 1 hour of the true maneuver time. It peaks at 50% of true total maneuvers being predicted within 72 hours of a true maneuver time also by the ISM model.

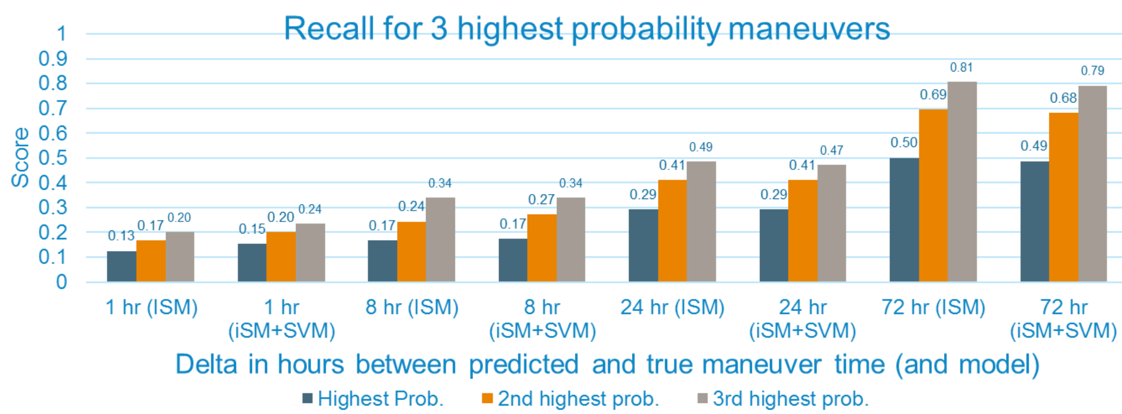

Figure 6: Recall results for the three most probable predictions are greatly improved over the recall results for just the most probable prediction (Figure 4). By returning the top three results, the recall increases to 24% within 1 hour for the combined model and 81% for within 72 hours of a true maneuver for the ISM model.

The drawback with returning additional results is that precision can take a hit because more predictions are being made, but not all of them are correct. We can use the probabilities themselves as a deciding factor as to whether to return just the most probable results or the top three most probable results for a prediction. When the system is confident on its prediction, the associated probability is much higher than the second most probable result. On average, when the most probable result is correct, it's associated probability is four times higher than the most probable result for an incorrect prediction for both the single and multiple model approaches. When there is a small delta between the top two or three probabilities, it indicates that the models are not confident in their prediction and the rankings between them are not reliable. Figure 7 illustrates this effect. On the right, is a confident, correct prediction has a much higher probability than the second most probable prediction. On the left, is a less confident prediction where the delta between the first, second and third most probable results is small. By returning multiple results only for low confidence predictions, we can boost recall without flooding the user with a large uptick in results and hitting to precision. This process is tunable to the user to allow them to achieve the right balance of precision and recall.

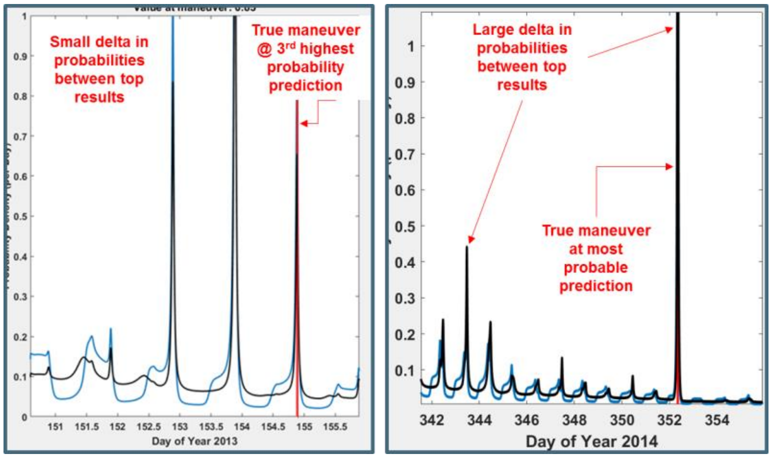

Figure 7: The system produces a vector of most probable maneuver times. For confident predictions, there is a large delta between the probabilities of the first and second most likely maneuvers (shown as blue (ISM+SVM) and black (ISM) peaks on right). For less confident predictions, there is a small delta between the probabilities of the most likely maneuvers (left) and sometimes the true maneuver (red vertical line) is not the one predicted with the highest probability.

Timeliness

A critical component of this approach is to predict maneuvers as far in advance as reasonable. Therefore, we evaluated our solution on how many days in advance can a maneuver time be correctly predicted. On average, true maneuvers were correctly predicted 8.5 days in advance, with correct predictions in as little time as 1 day in advance and up to over 2 weeks in advance.

Likelihood Performance

Finally, we compared the ISM and ISM+SVM Combined model results for likelihood metrics. As is seen in Table 1, the combined model performed slightly better than the ISM on instantaneous log-loss and mean PDF and performed similarly in the other metrics. Log-loss is a measure the likelihood of the observed data given the prediction. In our case it measures the likelihood that the maneuvers would occur when then did given the predictions made by the algorithm. A baseline used for comparison (which modeled maneuver events as a Poisson process – meaning it assumed no pattern) achieved a log-loss score of 2.73. Since it is a "loss" metric, smaller values are better than larger ones. Comparing the score of the ISM+SVM model (2.11) to the baseline score shows it had approximately 86% increased performance (using the exponentiation of the difference in scores.

Table 1: Metrics comparing the performance of the ISM vs. the ISM+SVM combined models.

Galaxy	Instantaneous Log-Loss	Prob Transformation	Mean PDF	Mean Daily Probability
Baseline (Poisson)	2.73	2.71	0.073	0.073
ISM	2.13	2.68	0.45	0.088
ISM +SVM	2.11	2.68	0.5	0.088

6. PREDICTION OF ANOMALIES

In the results section above, we presented how well we predicted future maneuvers. However, the purpose of this system is not just to predict future maneuvers correctly, but also to rapidly detect when there is a deviation from expected PoL. When we use the term anomalous, we are simply referring to the observation that the maneuver is unusual or has different characteristics from the majority of other maneuvers occurring in the dataset. As mentioned above and shown in Figure 2, maneuvers occurred every 3.5 days on average in this dataset, but there were periods when multiple maneuvers occurred in 24 hours and when only 3 maneuvers were performed over the course of 3.5 weeks. Figure 8 shows three instances of anomalous maneuvers in the dataset and our results over them. On the left, there were two maneuvers that occurred within 24 hours. We correctly predict the first one with a very high spike in probability. We also predict the second maneuver with a low probability 7 days in advance. For the middle and right charts, the models have a very strong bias to predict maneuvers during a 17- and 20-day gap in maneuvers, respectively. In both cases, many maneuvers are predicted in rapid succession which indicates to the user that there is a deviation from expected PoL within that timeframe. In all three of these anomalous cases, the system is alerting the user to a change from expected behaviors.

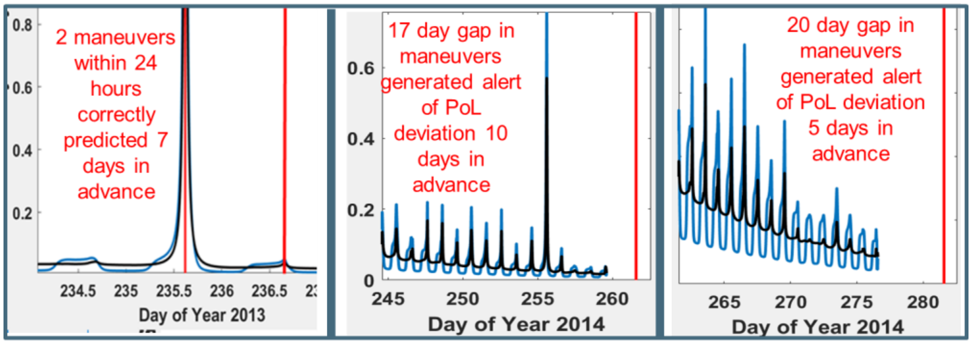

Figure 8: Models were able to flag anomalous maneuvers (red lines) that deviated from Galaxy PoL both in terms of unanticipated maneuvers that occurred (left) and anticipated maneuvers that did not occur (right).

7. CONCLUSIONS

Our approach has demonstrated the ability to predict 70% of maneuver times within 3 days of a true maneuver time and 22% of maneuver times within 24 hours of a maneuver. We have also been able to detect deviations from expected maneuver patterns up to a week in advance. Although the incorporation of the SVM did not significantly boost overall scores, it did provide benefits in increased confidence and increased precision of predicted maneuver time.

This technology can be used to increase timeliness of unanticipated dynamic events to provide operators with maximum time to generate and execute a COA, if warranted. The application of this work extends beyond maneuver prediction. It can be incorporated into data association tasks for Uncorrelated Track (UCT) correlation. It can be used to dynamically task a constellation of sensors to decrease observation gaps. And it can be used for left-of-event prediction of large scale, long term patterns of life. Future work is planned for validation on larger datasets, additional objects and model extensions.

8. ACKNOWLEDGEMENTS

We want to thank our sponsors at the Air Force Research Labs Space Vehicles Directorate for funding this phase of our work through contract number FA9453-12-C-0302.

9. REFERENCES

1. Balcan, Maria-Florina, Avrim Blum and Santosh Vempala (2008). *Clustering via Similarity Functions: Theoretical Foundations and Algorithms*. In Proceedings of the 40th ACM Symposium on Theory of Computing (STOC), Cambridge, MA.
2. Gao, J.; Ling, Haibin; Blasch, Erik; Pham, Khanh; Wang, Zhonghai; Chen, Genshe (2013). *Pattern of Life from WAMI Objects Tracking based on Context-Aware Tracking and Information Network Models*. In Proceedings of SPIE 8745, Signal Processing, Sensor Fusion, and Target Recognition, Baltimore, MD.
3. Long, Letitia (2013). *Activity Based Intelligence: Understanding the Unknown*. In “The Intelligencer Journal of U.S. Intelligence Studies.” Volume 20, Number 2 (pp 7-15).
4. Masunaga, Samantha (2016). *Small satellites are back, with down-to-earth expectations.* “LA Times, http://www.latimes.com/business/la-fi-adv-small-satellites-20160519-snap-story.html” accessed 1 September 2015.
5. Scholkopf , B and Smola, A. (2001). *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press Cambridge, MA.
6. Shabarekh, C. Kent-Bryant, J, Garbus, M., Keselman, G., Baldwin, J. and B. Engberg (2016). *Efficient Object Maneuver Characterization to Support Space Situational Awareness*. In proceedings of the Technical Track at the 32nd Space Symposium, Colorado Springs, CO.