Contextual Predictive Model for Early Identification of High-Covariance Conjunctions

T. Olson, C. Reid, J. Stauch, C. Grey, D. Kesler, B. Marchand
Slingshot Aerospace, 841 Apollo Street, Suite 350, El Segundo, CA 90245

ABSTRACT

Satellite operators regularly assess conjunction risks and weigh trade-offs between mission operations and conjunction avoidance actions in the increasingly congested orbital environment. As a conjunction event approaches, operators typically receive a sequence of conjunction data messages (CDMs) from the US Space Force's 18th and 19th Space Defense Squadrons. The CDMs identify objects involved in the conjunction, report a measure of risk and margins for the event, and estimate uncertainty in the form of state covariance matrices. Operators then use the information in the CDMs, along with supplemental data, to determine whether to take action. The methods presented in this paper aim to facilitate operators' decision-making by identifying which conjunction events are likely to benefit from additional data acquisition, and which are likely to self-resolve from routine tracking, well in advance of when they must commit to maneuvers. We introduce a machine learning based model developed by Slingshot Aerospace that takes into account the temporal evolution of CDM sequences, and it predicts, in advance, the conjunctions for which the covariance will not decline enough via typical tracking. The model incorporates contextual information about the objects involved in the conjunction, including physical characteristics and pattern of life details, that enable it to outperform baseline covariance-only approaches. Our model has demonstrated reliable performance, identifying high-covariance CDMs, up to five days prior to the time of closest approach, which equips operators with sufficient advanced notice to request observations and obtain refined state estimates with enough time to incorporate the enhanced information into their collision avoidance maneuver decision.

1. INTRODUCTION

The population of human-made satellites in low Earth orbit (LEO), and beyond, has recently expanded tremendously, and there are further indications of accelerated growth in coming years. With the resulting expansion of resident space objects (RSOs), there has been a corresponding rise in the frequency of near-misses and collision risks that may require active intervention from spacecraft operators. A decision to maneuver for collision avoidance implies costs and risks to operators, including fuel consumption, increased wear on spacecraft subsystems, and reduced operational lifetime, in addition to interruptions in mission execution. Operators frequently weigh the risks of maneuvering against the potential consequences of a collision [1]. It is therefore essential to obtain reliable information about each conjunction to support data-driven decisions.

The conjunction risk assessment process for a typical operator involves the assimilation of conjunction data messages (CDMs) that encode information about the upcoming close approach. Each CDM includes details about the operator's craft, referred to as the "primary" in the conjunction, and the "secondary" object that may be debris or another operational spacecraft. In addition, the CDMs report forecasted miss distances (MD) and state uncertainty matrices in the form of covariance estimates, which are combined through a mathematical model to generate a "probability of collision" (Pc) metric at the time of closest approach (TCA). This latter provides a single number by which the collision risk could be measured. However, nonlinearities in the modeling process lead to an interesting relationship between the covariance, MD, and Pc estimates, whereby the interpretation of the Pc metric can become ambiguous in what is known as the "dilution region" of parameter space [2–6]. Although the dilution region is not the focus of this paper, we provide a discussion about the connection to our predictive model.

There are several ways that the reported CDM information, covariance estimates in particular, can improve over time. One approach is simply to wait until the conjunction is closer in time because the size of the CDM-reported covariance is affected by propagation to TCA. As new observations are routinely collected, the propagation duration from last observation will shorten, and thereby decrease the covariance growth. However, variability in tracking density may lead to insufficient refinement by the Maneuver Commit Point (MCP). Thus an important second option is to decrease the uncertainty by sourcing additional observations that would not have otherwise been collected. These are usually only collected on the secondary object because it tends to be the object with larger covariance, and often the primary object already has precise ephemeris information available, mitigating the need for additional observations.

Sensor availability, viewability, and time constraints imply prohibitive challenges to additional data collection for every secondary object prior to the MCP, especially as the number of objects in orbit grows. Identifying conjunctions as early as possible that would benefit from additional object tracking is therefore paramount, so that only necessary additional observational data is requested and so that it can be collected, fused, and delivered in time to enable operators to make informed decisions. This scaling challenge motivates the work presented in this paper: a machine learning (ML) based model developed by Slingshot Aerospace that predicts temporal evolution of CDM sequences and identifies, in advance, the conjunctions that are likely to need additional tracking before the MCP.

In addition to covariance information, the model incorporates contextual information about the objects involved in the conjunction, including physical characteristics and pattern of life details that enable it to outperform baseline covariance-only approaches. Our model has demonstrated the ability to identify high-covariance CDMs that are likely to benefit from additional tracking as early as five days prior to TCA. This timeframe equips operators with advance notice sufficient to request additional observations and obtain refined state estimates that can be incorporated into their collision avoidance maneuver decision. Our predictive prioritization approach can be integrated into tasking or planning systems to improve conjunction assessment (CA) processes for the growing spaceflight industry.

We first provide background information with a more detailed description of the problem and constraints; next we introduce the data and methods considered in this analysis; and finally we present the results and discussion. The CDM data used in the analysis spanned a time range from June 2023 through June 2024 and was provided thanks to the Eutelsat Group's OneWeb constellation [7] in partnership with the Slingshot Beacon platform [8].

2. BACKGROUND

2.1 Conjunction Data Messages

Conjunction data messages (CDMs) are generated by a conjunction assessment (CA) process that predicts close approaches of various orbital objects and estimates associated conjunction metrics that include the miss vector and probability of collision. The states of RSOs, sourced from catalogs and/or operator ephemerides, are propagated forward in time with high fidelity dynamics models to estimate the future positions and velocities, along with associated covariance matrices representing uncertainties in the estimates. A pairwise screening process is then run to detect possible close approaches (conjunctions) between objects. For each object, a volumetric screening is applied to detect if the forecasted trajectories of other objects intersect the protected volume around the object in question. When an intersection that satisfies conjunction criteria is identified, a CDM may be issued to alert the spacecraft's operator.[1]

The 19th Space Defense Squadron (SDS) works alongside the 18th SDS to perform conjunction assessment and issue CDMs as appropriate [9], using their own high accuracy catalog (HAC) along with operator-provided ephemerides to perform the screening process [10]. A single conjunction screening cycle often yields multiple versions of CDMs for the same conjunction event, depending on which source of state estimates was used in the calculation. New CDMs are reissued regularly with revised forecasts as additional observations are routinely collected. The new observations are fused with prior catalog information to update state estimates and conjunction metrics. Consequently, operators usually receive a sequence of CDMs preceding each TCA.[2] The included covariance (uncertainty) estimates in those messages tend to decrease as a function of the amount of time remaining to TCA due to an increase in information and decrease in propagation time (uncertainty grows as a function of propagation time) [9].

Slingshot's Beacon platform [8] aggregates CDMs issued by the 18th and 19th SDS. The Beacon data ingestion process associates CDMs with either an existing event or a new one if no prior CDMs have been seen with similar conjunction details; this tagging allows for the entire sequence of CDMs to be uniquely associated with a single conjunction event. The sequence then can be addressed as a unit, rather than treating each message individually.

[1] There are benefits to tracking debris-on-debris and non-cooperative operator conjunctions for purposes of SDA and analytics, but those are out of scope for the principal focus of this paper, which is intended to aid active operators with their decision-making process in the lead up to a future conjunction event.

[2] Note that the screening process may simultaneously produce multiple CDMs for a given event during a single update cycle, which should be distinguished from the multi-CDM sequences that lead up to a conjunction event. The latter CDMs encode updated predictions with new information, whereas the former represent different source data from the time period.

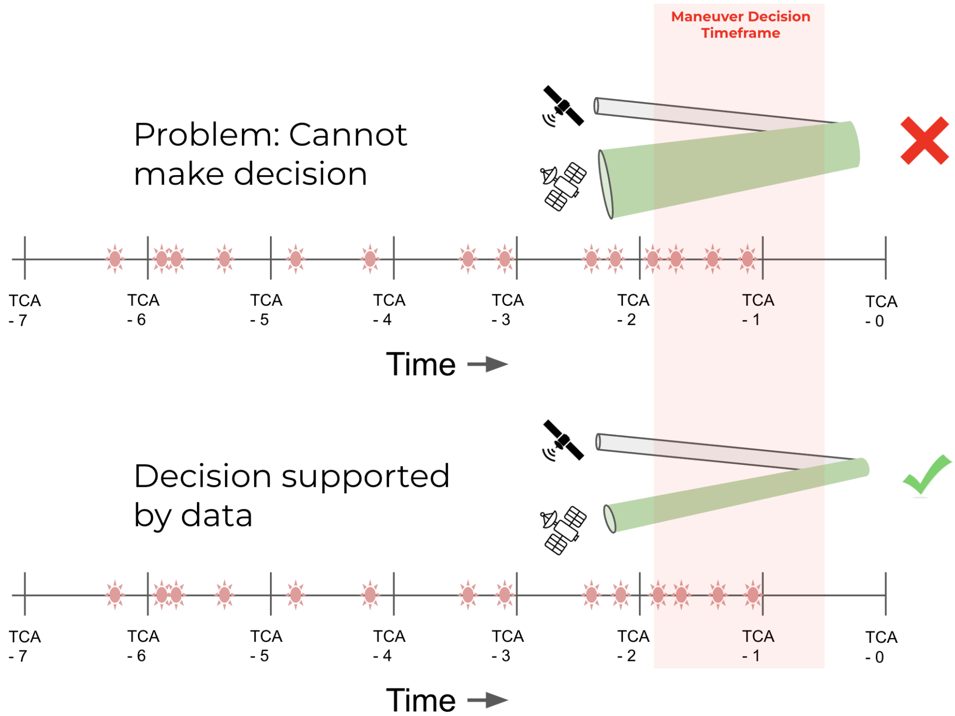

Fig. 1: Operators evaluate the information available leading up to the MCP and determine how to respond to the upcoming conjunction event. If the uncertainty is small enough at the MCP, then the decision whether or not to maneuver for collision avoidance is justified by the available information; however if the uncertainty is too large, the operator is unable to make an informed maneuver decision. Each mark on the timeline represents a newly issued CDM. They are generally issued on an eight hour cadence, with some variability. The grey cylinder illustrates smoothed evolution of the TCA-propagated state uncertainty of the primary object as reported in the CDM sequence, and the green cylinder is analogous for the secondary. They visually shrink as TCA approaches, in accordance with the shorter propagation time to TCA from subsequent CDMs.

2.2 Evasive Action Decisions

As TCA approaches, operators decide whether information in the CDMs warrants a collision avoidance maneuver. Operators also consider relevant supplemental data and trade-offs between impacts to mission operations and risk profiles. As illustrated in Fig. 1, many operators require a sufficiently small uncertainty estimate to make an informed maneuver decision.[3] The operator must commit to a course of action prior to the maneuver commit point (MCP), the latest time in advance of the conjunction event (i.e., closest to TCA) within which they can deliver any required commands to the spacecraft. MCPs vary by operator and communications opportunities leading up to TCA and typically range from hours to days before TCA, depending on available communication windows and spacecraft capabilities.

CDMs issued multiple days before TCA typically are characterized by larger uncertainty values than CDMs issued closer to TCA due to the longer propagation time to TCA. Although ongoing observation activity tends to improve the uncertainty estimates over time according to a quadratic scaling law, the amount by which it decreases may not be enough to enable a data-driven decision by the MCP. Fig. 2 illustrates several possibilities: whether the initial uncertainty is large or small, as TCA approaches and states for the objects of interest are updated, the estimated uncertainty might decrease far enough on its own to make a decision, but it might not. If the uncertainty remains large at the MCP, then the key metrics like MD and Pc that are used to evaluate the risk of a collision may not be reliable indicators of the true danger. Without remedy, the operator may be left without enough detail to make an informed decision, as discussed extensively in the "dilution region" literature [2–6].

[3]The uncertainty estimate is computed as the square root of the largest element in the CDM-reported radial-transverse-normal (RTN) covariance matrix.

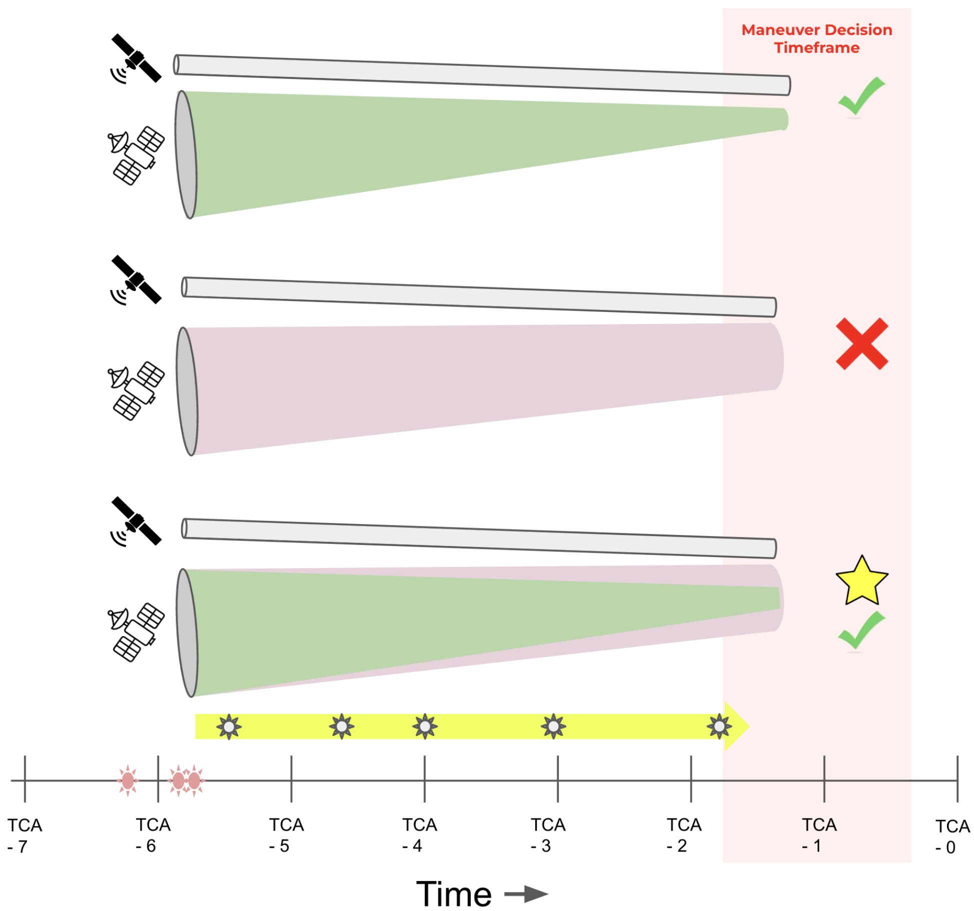

Fig. 2: Whether the uncertainty is relatively small or large in early CDMs around TCA-5 days, routine tracking operations tend to reduce the uncertainty in later CDMs. It may shrink by enough to support an informed maneuver decision at the MCP (top, pink), or it may remain too large at the MCP (middle, pink). This paper aims to use the limited data available at TCA-5 (illustrated via CDM marks only prior to TCA-5 on the timeline) along with contextual information about the objects to predict which situation will be reached at the MCP, near TCA-1. If the uncertainty would be too large (bottom, pink), additional observation data could be sourced and fused during the intervening time (yellow sunshine cartoons) to reduce uncertainty at the MCP and better support an informed decision (bottom, green).

2.3 Dilution Region

Conjunction risk assessment processes frequently depend on miss distance, probability of collision, uncertainty (covariance), and hard-body radius. These criteria are widely used across the industry, yet non-linear relationships among the elements lead to non-intuitive interpretations. In particular, a low Pc value may either be an indicator of low risk or an indicator of low quality information, depending on the size of the projected covariance ellipse in the CA calculation. In the latter case, the possibility of collision may still be substantial, but large uncertainty in one or more of the objects' state estimates effectively smears or "dilutes" the probability density over a wide area. The range of parameters for which this occurs has previously been labeled the "dilution region" [2–6]. In such situations, a low Pc may not be a reliable estimate of the intrinsic collision risk; this bears notable connections to the motivations of our proposed predictive model for high-covariance CDMs. It could be illuminating to study potential complementary effects and synergies between our model and dilution region phenomena in a future study.

2.4 Action Criteria

The criteria by which operators determine whether an avoidance maneuver will be executed depend on mission requirements and spacecraft capabilities. They may be defined by minimum values for Pc, miss distance, or various projections of the miss vector [11]. Operators also often consider the uncertainty represented by the magnitude of the covariance elements or number of observations used in the state estimate calculation – a metric closely related to uncertainty. At the time of writing, there were no globally agreed-upon policies in the spaceflight industry for determining whether an active avoidance maneuver is required or recommended, although the NASA handbook suggests Pc $\geq 10^{-4}$ is a common threshold, while Pc $\geq 10^{-5}$ or less offers extra conservatism [11]. Ultimately, each operator sets their own criteria and thresholds for making a maneuver decision.

We do not consider any specific maneuver-commit criteria because the details of whether or not to maneuver are out of scope for this discussion. Rather, we focus on predicting which CDMs will have a large uncertainty at the MCP. Specifically, the model presented herein predicts the likelihood that **the square root of the largest component of the RTN covariance matrix exceeds 1 km at an MCP of approximately TCA-1 day (24 hours prior to the CDM-reported TCA)**. We label those events as "high-covariance conjunctions". Importantly, we are not attempting to predict the actual covariance directly, rather the model predicts the likelihood that (the square root of) it exceeds the threshold value.

2.5 CDM Enhancement

When considering the conjunction events for which uncertainties are going to exceed the target threshold, additional effort may be warranted to source independent data for fusion with the available measurements to further refine the decision parameters. External data could be obtained from third party providers, which typically involves tasking optical or radar sensors to collect additional observations of the object(s) of interest. Those observations can be combined with existing data, either using the original data from the CDM-provider or derived from the CDMs if applicable, to generate enhanced conjunction assessments.

Depending on the objects' orbital parameters, sensor scheduling and viewability constraints, and details of the data fusion processes, the cycle of gathering data and generating enhanced CDMs can require up to several days. It is therefore paramount that this process begins well in advance of TCA and continues through to the MCP, so that the enhanced CDMs are made available with sufficient time remaining to make an informed decision.

It is also important to identify objects that are most likely to require tracking in order to prioritize resources. Indeed, if observation resources were unlimited, simply tracking every object in the catalog more frequently and precisely would resolve most issues. Of course that is not feasible with current technology, and given the expected growth in space traffic it seems unlikely to become more feasible in the near term. Even reducing the scope of the problem to providing additional tracking on all secondary objects would require a substantial expansion in tracking capacity. Yet, most conjunction events reach sufficiently small uncertainty by the MCP without additional tracking, so most of those additional requests would result in unnecessary, wasted effort.

The need for early notice to collect observations, along with prioritization of those collection requests, motivates development of a predictive binary classification model. The model-generated probabilities could yield a numerical ranking of upcoming conjunction events in order of expected likelihood of exceeding the high-uncertainty (insufficient information) threshold. Tracking the associated objects could be prioritized to obtain the requisite observations while making best use of the limited observation windows and resources. Fig. 3 illustrates the potential impact of the model presented in this paper. See the Evaluation section for additional explanation.

Fig. 3: *Top* – Categorization of conjunction events into those that support informed maneuver decisions with sufficiently small uncertainty (blue) and those that do not support the decision due to unacceptably large uncertainty at the MCP (orange). *Bottom* – The right-shift of well-supported decisions enabled by this approach (yellow overlapping the orange zone) with minimal wasted effort (yellow overlapping the blue zone). The yellow box illustrates the events impacted by the output of the model; it mostly overlaps the orange zone, which indicates the value lift in that those events have been positively affected with improved information, while the small overlap of the blue zone indicates that those events would have had sufficient information even without intervention.

2.6 Special Characteristics

An important question is whether there are special object characteristics or specific subsets of the spacecraft catalog that consistently fail to meet the desired uncertainty bounds. For instance, it has often been suggested anecdotally that a certain set of debris objects may be poorly tracked in general, such as small pieces with low radar cross section or optical reflectivity. If all secondaries with high uncertainty could be categorized in this way, then a concerted effort to enhance the tracking of those objects would effectively resolve the vast majority of high-covariance events and thus eliminate the need for any sort of predictive model. However, although prior research has identified certain objects as being statistically associated with "higher risk" [12], there is no indication that specific characteristics or orbits are globally and persistently associated with high-covariance conjunctions. Our analysis presented herein differs from prior work by contextualizing the question to predicting only those high-covariance events within the domain of ongoing CDM-sequences. Within that limited scope, object characteristics like size or hardware properties may provide valuable indicators of the likely evolution of associated conjunction metrics. Several forms of supplemental information, like estimated radar cross section, are already included in standard CDMs, and additional metadata about the spacecraft can be sourced from external sources such as Slingshot Aerospace's Seradata space object catalog [13].

3. METHODOLOGY

Our aim was to develop a model that predicts which secondary objects were likely to be involved in conjunction events that violate the 1 km uncertainty threshold at TCA-1, as described above (Section 2.4). We first developed a candidate set of predictive models to identify which objects would be involved in conjunction events violating the uncertainty criterion. Then we evaluated the performance of each of the candidate models using a masked test set to evaluate the prediction improvements brought by the addition of contextual information.

3.1 Data Sources

For this investigation, we processed CDMs that reported Eutelsat Group's OneWeb constellation as the primary objects, excluding any in which the secondary was also from that constellation. The data was drawn from Slingshot's Beacon platform. The models were trained on data from June 2023 through November 2023, and evaluation data spanned January 2024 through June 2024. December 2023 was excluded as a gap between training and evaluation to eliminate any cross-information contamination between the two time periods. Altogether the combined training and evaluation datasets included CDMs associated with approximately 955,000 events from 17,000 distinct secondary objects.

We retained only those CDMs which were generated by screening an operator ephemeris file for the primary object against a secondary object based on the 18th SDS's high accuracy catalog (HAC). We also excluded CDMs that failed to pass minimal data quality checks, including reporting default values for covariance or NaN values in the probability of collision field. Finally, since we were interested in predicting high-covariance conjunctions well in advance of the MCP, we only considered CDM sequences wherein the first message originated at least 5 days prior to TCA. After those filters, the final number of conjunction events was reduced to approximately 590,000 events from 9,000 distinct secondaries.

Contextual data was collected from CDM records and metadata in Beacon and merged with details from Slingshot's Seradata space object catalog [13]. From the former, we extracted information like object type (payload, rocket body, debris, etc) and estimated radar cross section, while Seradata provided additional details like thrust capabilities and activity status.

3.2 Lost Conjunctions

A key challenge arises from training and evaluating models on real-world data. Both the primary and secondary objects continuously interact with their environment and may change state for a variety of reasons. In particular, maneuver plans and associated ephemerides can be updated and revised over time, especially if the operator decides to take action in response to an upcoming conjunction. Those decisions can be made as late as the MCP, but it is also possible to commit preemptively to a maneuver before that deadline. In such dynamically evolving situations, some conjunction events will effectively resolve prior to the MCP, which terminates the CDM sequence before the final report of the covariance estimates. We addressed that challenge conservatively in our analysis by treating those "lost" conjunctions as having low covariance at the MCP. Thus the performance numbers represent a lower bound on the predictive quality, as some of those lost conjunctions may have ended up as high-covariance CDMs had they not been preemptively resolved. We also show the effect of removing all lost CDMs a posteriori; although it is not a true representation of the model quality, similar filtering conceivably could be approximated with additional knowledge or advance prediction of lost conjunctions.

3.3 Model Details

We developed a multi-feature ML model to predict which secondary objects would be involved in high-covariance conjunctions as defined in Section 2.4. This model was then evaluated against three other models using the same criteria. The set of included features for each model was developed through a combination of importance studies and consultation with SDA and CA subject matter experts (SMEs).

The central multi-feature ML model synthesizes early-CDM (well before TCA) covariance estimates and number of observations along with contextual features like object size and orbital parameters. We employed a gradient-boosted decision tree framework using a binary log-loss function, specifically the LightGBM framework [14]. We chose this architecture because it includes native support for both numerical and categorical variables, efficient resource utilization at the relevant data scales, sufficiently high performance on the experimental datasets, and relatively straightforward interpretability and explainability. We also considered other model formats including logistic regression, but none outperformed or provided sufficient benefits compared to the selected framework. Hyperparameters of the LightGBM models were selected by comparing F1-scores across training runs; only the top-performing model was used to report the final evaluation benchmarks below.

For comparison, we defined three other models with different levels of sophistication. The reference model with the least level of complexity "predicts" by randomly selecting from the secondaries. Next we compared with a SME-motivated model in which the early-CDMs are sorted by combined covariance size, and the secondaries with largest covariance are flagged; this model explores the common assumption that there are certain poorly tracked objects that are easily identifiable by their large covariance matrix elements. We call both of those reference approaches *rule-based* methods because they use fixed logic based on a set of predefined heuristic rules.

We also compared with a *trained* method that uses two CDM-based input features, again motivated by SME discussions. This reference model consumes both the early-CDM covariance along with the number of observations used to generate the secondary state. To maximize potential performance of the model and enable more direct comparison with the main method, we trained both the Covariance+Obs baseline and the full multi-feature model with the same LightGBM architecture, though the reference model only considered those two features.

The four models are enumerated below for later reference:

1. [*Rule-based*] **Random:** Randomly selected (conjuncting) objects
2. [*Rule-based*] **High Cov:** Only objects with highest covariance at TCA-5
3. [*Trained*] **Cov & Obs:** Covariance + number of observations at TCA-5
4. [*Trained*] **Multi-feature contextual:** TCA-5 covariance + number of observations + cross-sectional area + orbit + object characteristics

3.4 Evaluation

The models were evaluated based on two metrics: Value Lift and Wasted Requests, using the N(=20, 40, 80, or 160) highest scored objects by each model. We define those metrics as follows:

- **Value lift** (also known as Recall): Of all batch objects with large uncertainty at the MCP, what fraction the model correctly flagged for additional observations.
- **Wasted requests** (also known as the complement of Precision): The model requested additional observations for an object by scoring it in the top N, but ultimately that object did not have large uncertainty at the MCP.

Fig. 3 in the Background section illustrates the situation and the relative portions of value lift and wasted requests.

For evaluation purposes we grouped conjunction events into batches distinguished by the UTC day of TCA and applied a conservative max aggregation over all secondaries in each batch. Thus if a secondary showed up in multiple conjunction events with TCA on the same UTC day, each of those events was processed separately and the maximum predicted likelihood was assigned as that object's final score. In a production environment, this approach would generate a daily list of high-interest secondaries for which additional tasking could be requested.

We assumed for this demonstration that the operator (or observation provider) has a finite budget, limiting the number of secondaries for which excess observations can be requested. This is consistent with the earlier remark that collecting additional observations on the entire growing space catalog is, or soon will be, infeasible. It is crucial to optimize the deployment of that budget and maximize the value lift while minimizing the number of "wasted" requests on objects whose MCP uncertainties would end up small enough on their own. For concreteness we considered four cases with request limits of 20 objects/day, 40/day, 80/day, and 160/day, where the number refers to the count of distinct objects (length of the high-interest list) and not the number of observations collected.

4. RESULTS

4.1 Overall data metrics

The breakdown of conjunction events across object classes shown in Fig. 4 emphasizes our earlier observation that requesting additional tracking on all secondaries associated with daily CDMs would be infeasible for most operators. The ability to predictively filter the candidate object list to a finite number of high priority targets directly addresses this challenge. The figure further illustrates the overwhelming contribution of cataloged debris objects to the conjunction landscape in the relevant orbital regimes for this analysis. There were some seasonal fluctuations across the time period, as well as an overall shift in the total number of daily events; all of which may be heavily influenced by the highly active solar cycle and associated effects on drag.

4.2 Inferred Feature Importance

The contextual model was trained on the first six months of the time period. Although numerous features were available, only a small subset were found to improve predictive capabilities. Fig. 5 shows the ten features with highest importance weights identified by the model training algorithm. The scale should be interpreted as approximate relative indications of the importances; individual predictions may not be influenced equally, depending on the CDM and contextual features. The list of top features includes the TCA-5 day uncertainty and number of observations, in line with expectations from subject matter experts. In addition to those standard features, the other most prominent contributors that stand out are three orbital elements of the secondary object: semi-major axis, eccentricity, and inclination; these elements appear to be quite important predictors, at least for the objects most commonly seen during the training and evaluation phases of this project. We note that the high importance weight assigned to those attributes does not imply that there are certain objects or orbital regimes that are poorly tracked in general. Rather it indicates that there are some aspects of those elements that, in combination with the rest of the features, are informative of the likelihood that large uncertainty will persist through the MCP. This information is useful for assessing the relative prioritization of additional tasking requests on certain objects. Additional predictive power was gained from details like the size, type, and thrust level of the secondary, and incremental value was found in examining the country that owns the spacecraft. We reiterate that these features are not necessarily related to tracking quality in a direct fashion, but instead may indicate small but significant biases in the underlying statistics that enable improved predictive strength. The relative importance of any particular feature may vary significantly between different operators and orbital regimes.

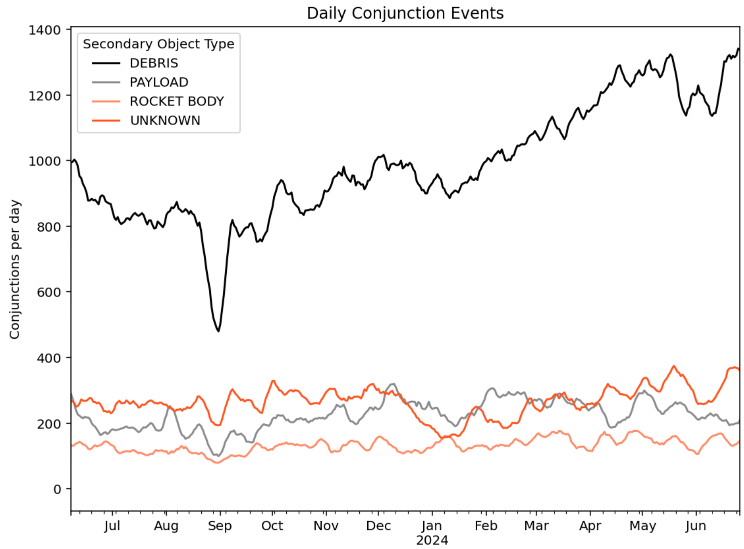

Fig. 4: Seven-day rolling mean number of conjunction events per day in the analysis dataset after applying the filters described in the Data Sources section. The data are segmented by the secondary object's CDM-reported type, illustrating the large relative proportion of debris objects.

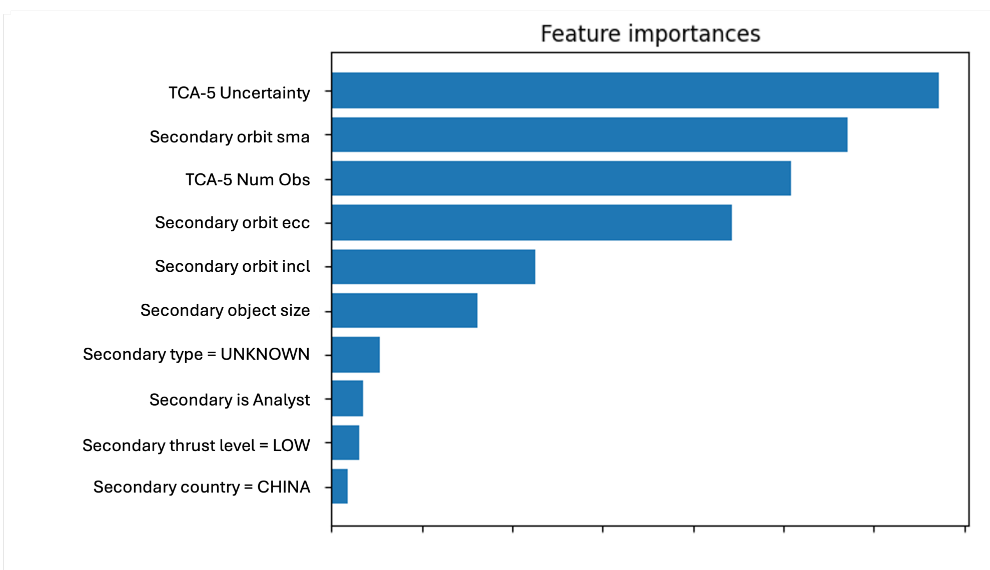

Fig. 5: Ten highest-ranked trained contextual model features. The extent along the x-axis represents the relative importance of each feature as measured by a split criterion [14].

Table 1: Performance evaluation metrics using the top-20, 40, 80, and 160 daily objects flagged by each model, averaged over the evaluation period.

20 requests / day	**1: Random**	**2: High Cov**	**3: Cov & Obs**	**4. Multi-feature**
% value lift (higher is better)	3.7%	9.1%	10.1%	**11.3%**
Avg wasted requests (lower is better)	14.5	6.4	5.0	**3.2**
40 requests / day				
% value lift	7.5%	18.5%	19.8%	**21.7%**
Avg wasted requests	28.7	12.4	10.5	**7.7**
80 requests / day				
% value lift	14.9%	37.2%	38.3%	**41.2%**
Avg wasted requests	57.8	24.4	22.8	**18.5**
160 requests / day				
% value lift	29.8%	61.0%	70.2%	**72.7%**
Avg wasted requests	115.5	57.7	55.1	**51.5**

4.3 Model Performance Comparisons

The top-N secondary-selection benchmarks for performance comparisons were generated using the evaluation dataset spanning the first six months of 2024. As explained above, "value lift" measures the recall fraction by each model, and "wasted observations" indicates the daily average number of false positives generated from the model over the evaluation period. Results are shown in Table 1; each section of the table represents a different hard limit on the number of daily flagged objects as indicated in the top left cell. The best value in each row is highlighted in boldface.

There are several key points illustrated in the tables:

- The multi-feature model (model #4) outperforms all comparison models in both of the critical business metrics, capturing a larger portion of the important objects while wasting fewer tasking requests on irrelevant objects.
- The "conventional wisdom" that objects with high covariance (uncertainty) in early CDMs tend to have large uncertainty near the MCP is reflected in the performance of model #2.
  - It is fairly effective, especially when requests are very limited (20-request portion of Table 1).
  - However, it is outperformed by other methods with increasingly wide margins as the number of available requests rises.
- The number of observations is a useful secondary feature to include alongside early covariance values.
- Although the early covariance and number of observations are good predictors of large uncertainty in later CDMs, the additional information included in the multi-feature model adds measurable improvement in all four cases.

4.4 Performance Without Lost Conjunctions

We also report results obtained by training and evaluating the models on an augmented dataset from which all lost conjunction events – those for which the CDM sequence ended prior to the representative MCP at TCA-1 – have been discarded. While this is not feasible in a production system because it requires *a priori* knowledge of which events will be lost, it provides an alternate perspective on the model characteristics and is suggestive of the best possible performance.

The feature importances inferred by the trained model very closely mirrored those shown in Fig. 5, indicating that the underlying data distributions were not significantly different after removing lost events. However, the performance metrics showed improvements across the board, which further emphasized the point that many of the lost conjunctions would likely have retained large covariance through the MCP, and those likely outcomes were reflected in the model

Table 2: Performance evaluation metrics using the top-20, 40, 80, and 160 daily objects flagged by each model after removing lost conjunction events, averaged over the evaluation period.

20 requests / day	**1: Random**	**2: High Cov**	**3: Cov & Obs**	**4. Multi-feature**
% value lift	4.2%	13.6%	13.6%	**13.7%**
Avg wasted requests	14.0	0.6	0.5	**0.4**
40 requests / day				
% value lift	8.2%	26.4%	26.6%	**26.9%**
Avg wasted requests	28.2	2.2	1.8	**1.5**
80 requests / day				
% value lift	16.3%	48.0%	49.5%	**50.7%**
Avg wasted requests	56.6	11.2	9.0	**7.3**
160 requests / day				
% value lift	33.4%	70.7%	81.5%	**83.6%**
Avg wasted requests	112.2	44.2	43.1	**40.1**

predictions. As before, the multi-feature contextual model consistently outperformed the other models in every situation. All three non-random models are roughly equally effective at capturing the most egregious cases as evidenced by the near-identical metrics in the 20-requests portion of Table 2. However, the separation between evaluation scores becomes more pronounced on larger tracking lists, which we interpret as an indication that the contextual model is able to identify important relationships between the contextual and numerical features. The additional information enables it to handle the less egregious cases more effectively.

5. CONCLUSIONS AND FUTURE WORK

Our results demonstrate the predictive value that a machine learning approach incorporating contextual features can contribute to conjunction assessment and risk analysis (CARA) processes. In this application, a model developed by Slingshot Aerospace that included contextual features, alongside more traditional signals, demonstrated a measurable lift in the ability to accurately forecast which upcoming secondary objects are most likely to retain unacceptably high covariance near the MCP. Such information can be useful for operators by identifying conjunctions that may benefit from added tracking. Since these predictions were based on data from TCA-5 day CDMs, well in advance of the MCP, it also enables sufficient advance warning to permit the potentially time-intensive process of gathering and assimilating extra observations to create more precise estimates of future spacecraft states.

Contextually augmented predictive models as proposed here could be incorporated into existing CARA systems. Operators (or CDM providers) could then work with observation data providers to source additional measurements as needed to support the action decisioning process. It would be interesting to evaluate the potential ancillary benefits of identifying and filtering early CDMs via this type of model, too. Advance prediction of which events *will not* require additional tracking could reduce or eliminate the need for frequent intermediate, yet unnecessary, updates for those events, while also providing advance warning about those for which more information will likely be needed. Sorting and filtering in this way would optimize the processing needs of both the CDM-provider and operator. Outside of the conjunction domain, similar models could also be beneficial for general-purpose catalog maintenance applications by enhancing existing methods with additional contextual data to improve tasking and scheduling prioritization.

Throughout the course of this work, numerous follow-up questions were identified, including

- How well does a single predictive model generalize to multiple operators or different orbital regimes?
- How well do these methods extend to predicting other conjunction-related metrics and criteria, especially those associated with high-risk and diluted CDMs?
- What other contextual features are important for predicting those metrics?

- We picked particular points in time, relative to the MCP, at which to evaluate the criteria (TCA-5 and TCA-1, respectively), but the relevant values for risk assessment (especially miss distance, probability of collision, etc) tend to fluctuate significantly as TCA approaches due to uncertainty in the state estimates. How can we better manage the forecasting and decisioning process in the context of this variability?

In addition to those topics, it would be very interesting to more precisely quantify the relative impacts of fusing observations with existing CDMs.

The analyses discussed herein represent an early demonstration of the value of modern machine learning methods in synergy with traditional techniques for analysis and forecasting in the space domain. As the number of spacecraft in orbit rises, along with the scale and volume of associated data, the opportunities and benefits afforded by the inclusion of ML models will likewise continue to increase. The next few decades are sure to be an exciting time for SSA practitioners with many new and interesting challenges ahead!

6. ACKNOWLEDGEMENTS

The authors are grateful to the leadership at Slingshot Aerospace for supporting this research, and to the many teammates and subject matter experts who provided valuable feedback and advice throughout this project. We very much appreciate the generous contribution of OneWeb CDM data from the Eutelsat Group, without which this analysis would not have been possible.

REFERENCES

[1] L. K. Newman, A. K. Mashiku, M. D. Hejduk, M. R. Johnson, and J. D. Rosa, "Nasa conjunction assessment risk analysis updated requirements architecture," in *AAS/AIAA Astrodynamics Specialist Conference*, no. AAS 19-668, 2019.

[2] S. Alfano, "Relating position uncertainty to maximum conjunction probability," *The Journal of the Astronautical Sciences*, vol. 53, pp. 193–205, 2005.

[3] M. S. Balch, "A corrector for probability dilution in satellite conjunction analysis," in *18th AIAA Non-Deterministic Approaches Conference*, p. 1445, 2016.

[4] M. D. Hejduk, D. Snow, and L. Newman, "Satellite conjunction assessment risk analysis for "dilution region" events: issues and operational approaches," in *Space Traffic Management Conference*, 2019.

[5] M. S. Balch, R. Martin, and S. Ferson, "Satellite conjunction analysis and the false confidence theorem," *Proceedings of the Royal Society A*, vol. 475, no. 2227, p. 20180565, 2019.

[6] C. Miller, "Efforts in solving the dilution problem for orbital collisions," Master's thesis, Purdue University, 2022.

[7] Eutelsat Group, "Eutelsat OneWeb." `https://oneweb.net/`, 2024.

[8] Slingshot Aerospace, "Slingshot beacon." `https://www.slingshot.space/solutions/applications#beacon`, 2024.

[9] 18th & 19th Space Defense Squadron, "Spaceflight safety handbook for satellite operators v1.7." `https://www.space-track.org/documents/SFS_Handbook_For_Operators_V1.7.pdf`, 2023.

[10] "Conjunction data messages : Recommended standard," Tech. Rep. CCSDS 508.0-B-1, The Consultative Committee for Space Data Systems (CCSDS), 2013.

[11] F. J. Krage, "Nasa spacecraft conjunction assessment and collision avoidance best practices handbook," tech. rep., NASA, 2023.

[12] M. Stevenson, D. McKnight, H. Lewis, C. Kundstadter, and R. Bhatia, "Identifying the statistically-most-concerning conjunctions in leo," in *Proceedings of the AMOS Technical Conference*, 2021.

[13] Slingshot Aerospace, "Seradata." `https://spacetrak.seradata.com/`, 2024.

[14] G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T.-Y. Liu, "Lightgbm: A highly efficient gradient boosting decision tree," *Advances in neural information processing systems*, vol. 30, 2017.