Orbital State Recommender Validation for Space Battle Management Applications

Oliver Schultz, Joseph Brink, Mark Williams
Lockheed Martin Space

ABSTRACT

Effective Space Battle Management (SBM) in an increasingly contested space environment relies heavily on automated Space Domain Awareness (SDA) software capabilities to enable swift and decisive responses to time-sensitive events. The large quantity of available orbital data necessitates a reliable and efficient means of selecting the most pertinent information, thereby minimizing manual intervention. To address this challenge, we developed the Orbital State Recommendation (OSR) service, which leverages a standardized process integrating qualitative and quantitative analytics to rank orbital data.

Today, a decision-maker might receive multiple conflicting states for thousands of objects with no quick way of determining which states are most reliable. The OSR service's automated ranking removes that uncertainty, allowing operators to quickly determine a high confidence course of action. The primary objective of this research is to validate the effectiveness of the OSR service, ensuring it consistently delivers accurate and actionable recommendations to support informed decision-making. Additionally, we aim to enhance the scalability and effectiveness of the OSR, enabling it to handle the large volumes of data generated by numerous commercial and government data providers.

Our research methodology involves developing a process for evaluating and ranking orbital states, considering factors such as satellite patterns of life, maneuver history, and predicted maneuver activity. The service assesses the longevity of states ensuring that our recommendations are accurate and actionable.

Validating the OSR service is a critical step towards instilling operator confidence in automated SDA software supporting SBM systems. The results demonstrate the OSR's effectiveness in recommending orbital states for various resident space objects and mission scenarios. By providing operators with accurate and actionable information, our OSR service can help reduce the complexity and uncertainty associated with space operations, ultimately supporting more informed and decisive decision-making. Furthermore, our research has the potential to contribute to the development of more effective SDA systems, which are essential for ensuring the safety and security of space operations.

1. INTRODUCTION

Space Battle Management (SBM) enables the United States to effectively operate in space and successfully accomplish military objectives to ensure space superiority. [1] SBM capabilities rely heavily on the accuracy and timeliness of the underlying Space Domain Awareness (SDA) data, which forms the foundation for all decision-making. SBM operators often receive orbital state data (e.g., state vectors, TLEs, etc.) with minimal insights into the provenance or pedigree of the data. Mission critical decisions may be made with incomplete or uncertain situational awareness. The expansion of commercial SDA data providers over the last two decades has improved the United States military's capability to understand the space environment through increased data availability. However, operators still face challenges in leveraging commercial data effectively and determining when to use different sources of data. [2] A poor understanding of commercial SDA capabilities can result in the underutilization of valuable information that could otherwise enhance critical SBM functions.

The accuracy and quality of orbital state estimates are influenced by the capabilities of SDA sensors and environmental factors. Factors such as sensor sensitivity, sensor placement, and processing algorithms can be modified but have limitations. External factors such as space weather, terrestrial weather, dynamic atmospheric conditions, magnetic fields, and solar radiation pressure are harder to accurately model. In an ideal world, all these factors would be modeled without error and all uncertainty would be eliminated. Operators would have perfect knowledge of all actors in the space domain. The SDA community can remove much of the uncertainty to various degrees to enable successful operations through advanced data processing. Accuracy estimates and associated metadata are shared with users in libraries like the Unified Data Library (UDL) which is improving transparency although not universal across sensor providers.

Operators do not have the bandwidth to manually assess the pedigree of all information they receive due to tight decision-making timelines. A high level of trust with automated processes to mitigate varying levels of state accuracy is critical for mission success. Operators often receive insufficient information to perform an in-depth comparison of orbital state data. Efforts such as the SDA TAP Lab have been created to accelerate the development of automated SBM capabilities to overcome some of these challenges. [3]

Lockheed Martin developed a cloud-based software application to deliver automated data pedigree assessments and orbital state recommendations to SBM applications. The Orbital State Recommender (OSR) service shown in Fig. 1 is part of the SDA data insights and pedigree application. Due to limited data often available to operators, a set of heuristic-based algorithms were developed to estimate expected state accuracies and provide a ranked list of recommended orbital states. The service can run analytics on multiple satellite catalogs and generates metrics for data provenance, state longevity, consistency, and latency. Maneuver detection and AI-based pattern of life assessments enable automated state recommendations to SBM applications in near-real-time. This paper provides an overview of how the Orbital State Recommender functions at scale and documents performance metrics.

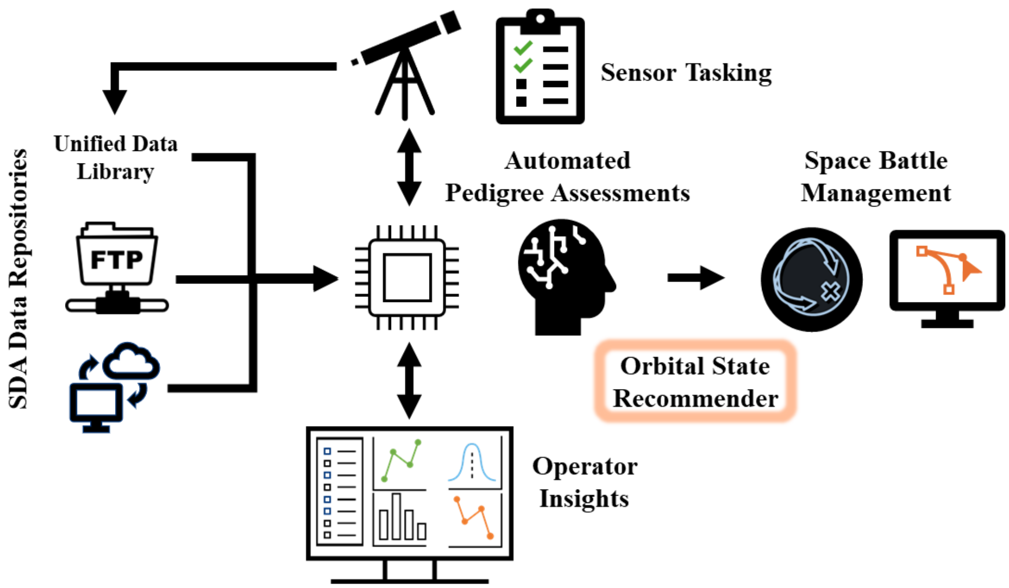

Fig. 1. The Orbital State Recommender shown within the SDA data insights and pedigree application supporting Space Battle Management missions

2. METHODOLOGY

2.1 Dataset

The data used to assess the Orbital State Recommender for this study was accessed via the UDL. [4] The OSR recommends states grouped by RSO ID, source, and type which are defined as follows.

- RSO ID – Unique identifier for each Resident Space Object (e.g., NORAD ID)
- Source – Organization that published the state (e.g., 18th SDS)
- Type – Ephemeris format (e.g., state vector, TLE)

State vectors and TLEs from ten different data providers published between May 19, 2025, and July 28, 2025, were used in this assessment. Over 4.2 million states (~1.1 million commercial and ~3.1 million USGOV) were processed averaging ~60,000 state recommendations per day. Table 1 shows the type of data, number of states processed, and the number of unique RSOs provided by each organization. The identities of the organizations have been replaced with unique identifiers to not associate findings with any specific data provider. Identifiers starting with "C" are commercial entities and identifiers starting with "USG" are United States Government affiliated organizations.

Table 1. Summary of dataset used for Orbital State Recommender validation analysis between May 19, 2025, and July 28, 2025

ID	Ephemeris Type	# States	# Unique RSOs
C-1	State Vector	33,384	16,882
C-2	TLE	22,581	4,160
C-3	State Vector	390,811	2,475
C-4	State Vector	8,723	5,801
C-5	TLE	2,639	19
C-6	TLE	640,328	2,653
USG-1	TLE	2,878,375	27,468
USG-2	State Vector	155,531	5,124
USG-3	TLE	82,995	25,165
USG-4	State Vector	30,097	115
	Total States	4,245,464	

The development of a fully automated, near-real-time processing capability for this data was challenging for multiple reasons. The frequency and quantity of data that is published is inconsistent and often fluctuates. It is common to receive thousands of states in a short period of time and then smaller updates most of the time. Modifications to how states were propagated and cached were required to reduce processing time while managing computational resources. An easier solution would have been to reduce the frequency of data ingests, however minimizing the time between data becoming available on the UDL and generating a state recommendation was a priority. New data is ingested and processed every 10 minutes to support time sensitive needs of SBM applications.

2.2 State Longevity

Fig. 2 shows a high-level flow diagram of how data is ingested from a data library and processed to generate a state recommendation. Starting on the lefthand side of Fig. 2, a data processor is used to query data from multiple databases (e.g., UDL, FTP server, etc.). The data is cleaned and cached in an internal database. Data cleansing is an important step due to the varying levels of metadata that data providers include when publishing states. Data can also be provided in different formats depending on the originating source, therefore, normalizing the state data is critical to make comparisons in later processing steps. In the case that data does not include maneuver information, a maneuver detection and characterization process is used to detect any changes from previous states. If a maneuver is detected, a Feedforward Neural Network is used to characterize the maneuver as "station-keeping" or "non-station-keeping". [5] Metrics such as state longevity and data dissemination latency are calculated and incorporated into the orbital state recommendation algorithm. The algorithm uses quantitative and qualitative metrics to generate an ordered list of orbital states to users. Historical metrics used by the state recommender are updated and cached in a database. Finally, old data that is no longer needed is purged from the database to manage data storage needs. Historical metrics are updated on a rolling basis and old data ages out as new information becomes available.

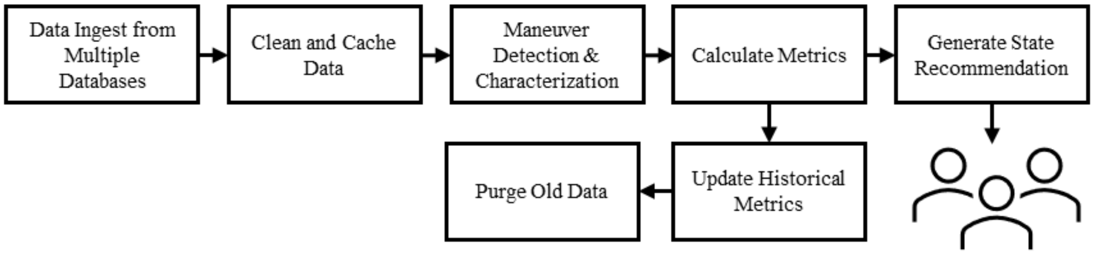

Fig. 2. Orbital State Recommender data flow diagram

A *state longevity* metric is one of the core quantitative heuristics used by the Orbital State Recommender to estimate state uncertainty over time. This provides an estimate for how far into the future an orbital state can be propagated while maintaining accuracy and is particularly beneficial for states that do not include covariance, such as TLEs. When a new state is ingested, the state longevity metric is updated. Previous states with the same RSO ID, source, and type are propagated to the latest state epoch and the relative ranges are calculated as shown in Fig. 3. Similar approaches have been previously published; however, these applications only used the NORAD catalog and were calculated offline. [6] The authors are unaware of an implementation that compares government and commercial datasets in near-real-time and at the scale described here. TLEs are propagated with the Astrodynamics Standards SGP4 propagator, and state vectors are propagated using an implementation of Draper Semi-analytical Satellite Theory (DSST). The DSST propagator accounts for 4x4 tesseral and zonal gravity models, solar radiation pressure, third body (sun and moon) and atmospheric drag perturbations.

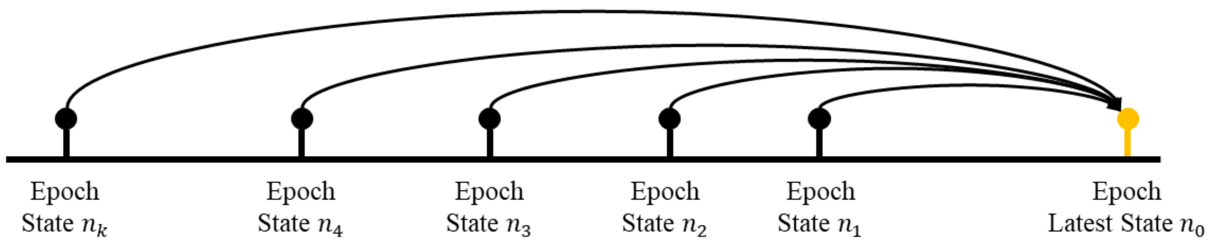

Fig. 3. States are propagated to latest state epoch and the relative position errors are calculated

Fig. 4 shows how state longevity is compared across different data providers. Each marker represents the difference in reported position between two consecutive state epochs for each source and type. A best fit line is estimated for each provider and used to predict uncertainty for future states. This method provides a heuristic measure of position uncertainty growth over time. In this example, states from Provider 3 (green line) diverge much more quickly from future state updates than states from Providers 1 or 2 (black and blue lines). In general, a state from Provider 1 or 2 would be expected to have less state uncertainty growth over time. However, if the most recent state from Provider 2 is 40 hours old, a state from Provider 3 that is only 6 hours old would be expected to have less uncertainty.

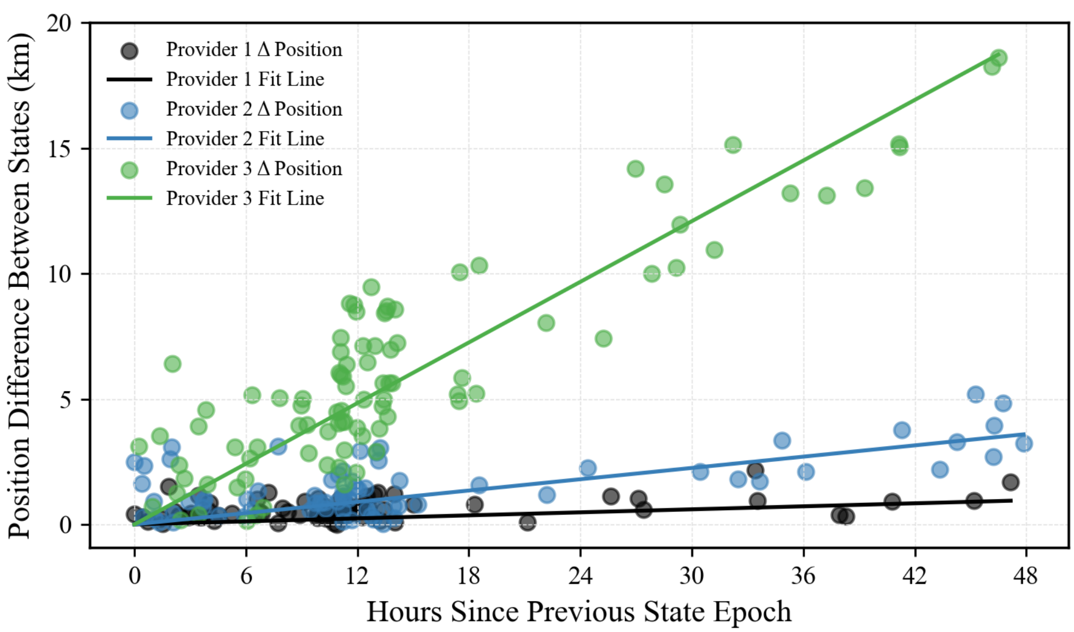

Fig. 4. State position differences vs. time between states for three different data providers

2.3 State Longevity Prediction Results

The state longevity was calculated for all records described in Section 2.1. The predicted position difference, based on the state longevity metric, observed position error, and the *Δ Position Error* (observed error - predicted error) was calculated and documented for each new state. Fig. 5 shows the cumulative distribution function (CDF) of the absolute value of *Δ Position Error* for each data provider across all RSOs. The plot can be interpreted as *X percent of the predicted state errors were within Y km of the observed state error*. For example, the state longevity metric predicted state differences for C-3, C-5, C-6, and USG-1 greater than 80% of the time within 2 km of observed error. Table 2 shows the middle 50th and 90th percentile *Δ Position Error* accuracies for each data provider.

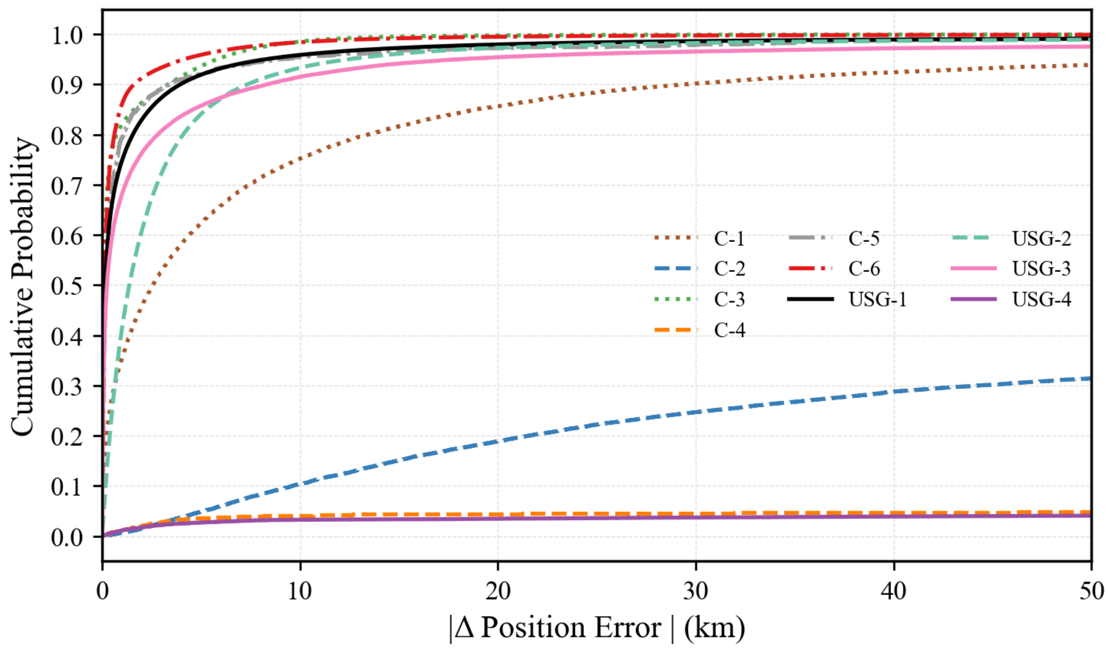

Fig. 5. CDF plot of the absolute *Δ Position Error* between observed and predicted longevity errors for 10 different sources of RSO state information

The state longevity algorithm can predict errors within a few hundred meters 50% of the time and within a few kilometers 90% of the time for most of the data providers. The algorithm has poor performance predicting state longevity for some of the data providers (e.g., C-2, C-4, USG-4). The errors do not appear to be correlated with the type of state or due to propagation errors. The authors believe that some of the poor prediction outcomes can be explained by the intended purpose or task for which the states were created.

Table 2. Middle 50th and 90th percentile *Δ Position Error* accuracies from May 19, 2025, to July 28, 2025

ID	Ephemeris Type	50th Percentile (km)	90th Percentile (km)
C-1	State Vector	[-2.8, 2.3]	[-31.5, 27.5]
C-2	TLE	[-294.5, 406.8]	[-2516.6, 3570.3]
C-3	State Vector	[-0.2, 0]	[-3.3, 3.3]
C-4	State Vector	[-1339.3, 5548.9]	[-6760.7, 10982.4]
C-5	TLE	[0, 0.5]	[-0.65, 8.1]
C-6	TLE	[-0.1, 0.1]	[-1.0, 2.8]
USG-1	TLE	[-0.1, 0.2]	[-2.96, 5.6]
USG-2	State Vector	[-1.7, 1.0]	[-6.8, 7.9]
USG-3	TLE	[-0.34, 0.15]	[-8.8, 7.8]
USG-4	State Vector	[-645, -263]	[-1251.9, 913.5]

Some of the providers appear to be creating states to support sensor scheduling for specific periods of time. It is likely that the states are generated to be accurate for a limited time span and are not meant to be propagated for longer durations. Today, there are limited automated mechanisms for SBM operators to quickly assess the nuances of the orbital state data available to them. The Orbital State Recommender is designed to help provide operators quantifiable metrics to make assessments even when data providers do not explicitly provide a disclaimer for the intended use of the orbital state information.

2.4 State Recommendations Accounting for Maneuver Predictions

The *Δ Position Error* results are skewed by maneuvers occurring between state updates. In many instances, especially for core SBM use cases, maneuvers occur with minimal warning and are challenging to predict. However, routine station-keeping maneuvers can be accounted for and predicted in advance. In previous work, a process to detect and characterize station keeping maneuvers was developed using a Feedforward Neural Network. [5] As an example, this process was used to characterize the station keeping maneuvers of IRNSS-1G, a GEO navigation satellite. IRNSS-1G performs regular station keeping maneuvers shown in the top plot of Fig. 6. East-west station-keeping maneuvers occurred around June 3rd and July 3rd, 2025, to maintain the satellite's orbital longitude slot.

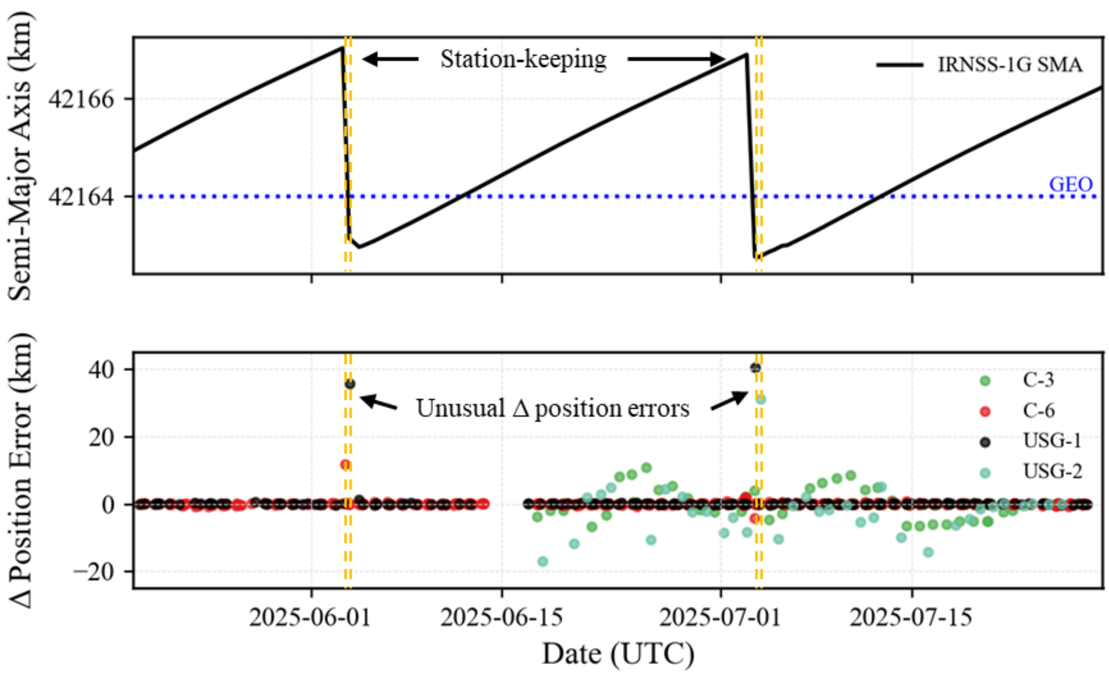

Fig. 6. IRNSS-1G semi-major axis time history showing station-keeping maneuvers (top) and the *Δ Position Error* highlighting increased error when maneuvers occurred (bottom)

The bottom plot in Fig. 6 shows the *Δ Position Error* of states from 4 different data providers. Most of the predictions are accurate within several hundred meters; however, a few noteworthy predictions have errors greater than 30 km. *Δ Position Error* greater than 30 km for USG-1 (black markers) is expected to occur less than 1% of the time for all state updates, as shown in Fig. 5. The orange, vertical dashed lines in Fig. 6 show the state updates that were published right after a station-keeping maneuver was assessed to have occurred which align with the large USG-1 *Δ Position Errors*. The authors assess that the significant prediction errors in this case are a result of the station-keeping maneuvers.

The Orbital State Recommender algorithm uses a maneuver prediction algorithm to provide a likelihood of when the next station-keeping maneuver will occur. This information is provided to users as an estimate of when an orbital state is expected to expire. For example, a subset of 97 IRNSS-1G historical maneuvers characterized as "station-keeping" were used to generate a statistical model to predict when the next station-keeping maneuver is likely to occur. Fig. 7 shows the empirical likelihood that a station-keeping maneuver will occur (blue) given the time since the last station-keeping maneuver and a model fit to the empirical data with a gamma distribution function (orange). The model can be used to quickly predict the next station-keeping maneuver without needing to use the entire time history of maneuver events.

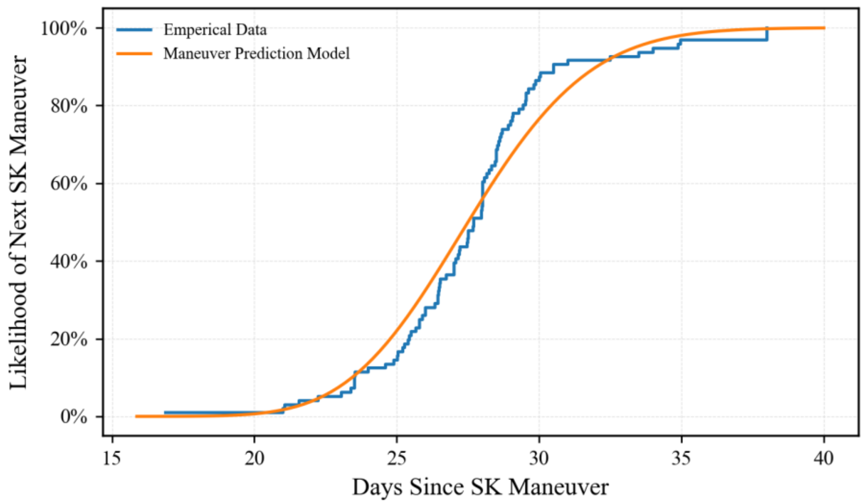

Fig. 7. The likelihood of the next station-keeping maneuver for IRNSS-1G given the time since the last station-keeping maneuver

The gamma distribution model has a mean value of 27.6 days between maneuvers with a standard deviation of 3.4 days. 90% of station-keeping maneuvers are expected occur between 22.4 and 33.4 days since the last station-keeping maneuver. As time progresses without a maneuver, it becomes increasingly more likely that a maneuver will occur.

The likelihood that a maneuver will occur in a given time frame and cause the current state to expire can now be calculated. In this example, if this process were run on June 27th, it would have been 24 days since the last maneuver. There is a 20% likelihood that the current state will not be accurate due to a maneuver within 24 hrs. (day 25) and a greater than 80% likelihood that it will not be accurate after another week (~31 days). IRNSS-1G conducted a station-keeping maneuver on July 3rd, approximately 30 days later which is within 1 standard deviation of the mean time between maneuvers. If a maneuver occurs outside of the predicted timespan, then an operator should be alerted that activity outside of normal patterns of life for the satellite may be occurring.

2.5 Value Proposition for Space Battle Management Applications

Operators have a great need for automated detection, prediction, and confidence assessments of space-related events. They should not need to be concerned with manually assessing every state database update because automated capabilities can ingest more data in less time. Decision makers need a consistent methodology to select states which is less dependent on specific user knowledge and skillsets. The Orbital State Recommender helps improve confidence in SBM solutions and reduces manual operator workload by providing automated state recommendations at high volumes. The service has been successfully demonstrated to run in near real-time on multiple catalogs, processing over 60,000 states per day.

In contrast, without this service, operators are likely to choose the latest state from a predefined provider. However, when making decisions several days in advance, an older state with better longevity is more likely to provide accurate results. Timely state recommendations, paired with maneuver detection is critical, enabling operators to evaluate the confidence in their assessments.

The ability to quickly assess the suitability of an RSO state for mission planning is important for SBM applications. This allows operators to rapidly determine whether a state is reliable for planning purposes. Additionally, the orbital state recommendation service can enhance critical missions such as sensor orchestration, collision avoidance, and threat assessment.

3. FUTURE WORK

The data used in this paper was logged using an automated process during day-to-day processing activities. However, the statistics and graphs were created manually. In the future, this information will be automatically updated, and performance benchmarks will be routinely published for users to view. This will include historical state longevity prediction accuracies for each data provider.

Currently, the authors have electro-optical observations of GEO satellites from a limited number of providers. With access to additional sources of observation data, there is a desire to compare recommended states to observations occurring after the recommendation time. The measurement residuals compared to the states can be used to show how accurately states align with future observations. This process has been demonstrated with success; however, the source of the observations is the same as some of the available states, which currently does not provide an independent assessment of state accuracy.

There is desire to enhance maneuver prediction models with a statistical mixture model. A gamma distribution generally fits the time between maneuver data well; however, a more robust model would likely account for maneuvers occurring at different cadences and irregular frequencies. We expect a mixture model to better represent the time between station-keeping maneuvers and improve prediction accuracy.

4. CONCLUSION

Overall, the Orbital State Recommender service has shown promise in supporting Space Battle Management operations by providing an automated capability to evaluate and rank orbital state data. The ability to accurately predict state longevity and account for maneuver uncertainty has been demonstrated through testing with operational data. Prediction errors are less than a few kilometers for most sources of data assessed. The service has proven it can process large amounts of state data from multiple data providers in near real-time, making it well-suited to support timely decision-making in operations. Further development and refinement of the service will continue to enhance its performance and utility in SBM applications.

5. ACKNOWLEDGEMENTS

The material for this paper was generated solely with Lockheed Martin Space internal research and development funding. The authors would like to acknowledge and thank Shankar Khanna, Dave Walton, and Anna Garcia for their support and technical review. This research was enhanced through our participation in the SDA TAP Lab, which provided access to a collaborative community and rich datasets. The opportunity to leverage these resources has been invaluable for this project. The paper was reviewed and approved by Lockheed Martin Public Information Release Authorization (PIRA) # SSS2025081874.

6. REFERENCES

[1] "Space Delta 2," USSF, April 2024. [Online]. Available: https://www.spaceforce.mil/About-Us/Fact-Sheets/Fact-Sheet-Display/Article/3738430/space-delta-2/. [Accessed July 2025].

[2] United States Government Accountability Office, "Space Situational Awareness: DOD Should Evaluate How It Can Use Commercial Data," GAO, Washington, DC, 2023.

[3] "SDA TAP Lab," [Online]. Available: https://sdataplab.org/. [Accessed July 2025].

[4] "UDL," [Online]. Available: https://unifieddatalibrary.com/. [Accessed July 2025].

[5] O. Schultz, J. Brink, M. Williams and C. Curry, "Data Insights, Pedigree, and Automation for Space Domain Awareness," in *Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS)*, Kihei, 2024.

[6] G. Peterson, R. Gist and D. Oltrogge, "Covariance Generation for Space Objects using Public Data," in *AAS/AIAA Space Flight Mechanics Meeting*, San Diego, 2001.