Automated similarity-based inference for space object characterization

Randy Jensen, Arthur Goldblatt, Joseph Birtman,
Devin Cline, William Lin, Gregory Howe
Stottler Henke Associates, Inc.

ABSTRACT

Defending space systems against intentional and unintentional threats depends greatly on space domain awareness (SDA), facilitated by tools to help human operators make faster decisions. The ability to reason under uncertainty is critical not only to operators but also to computer-based tools of the present and future that are essential to the SDA mission. In this paper we describe a similarity-based inferential approach to help with sensemaking from multi-source data, both for operation within automated machine-based analytic workflows and ultimately for helping space operators avoid operational surprise. This approach uses a knowledge graph constructed as a relational network to reason about space objects based on calculated similarities and relationships. This approach has been implemented in a functional tool initially tested in the SDA Tools Applications and Programming (TAP) Lab as part of an integrated battle management framework under development. This paper describes the technical approach with examples based on selected use cases, and walks through the current applied implementation and interoperability within the SDA TAP Lab battle management architecture, as well as plans for future development.

1. INTRODUCTION AND BACKGROUND

With the proliferation of space objects there has also been an increase in sensors, resources, and analytics to collect, process, and reason from observational data to support multiple levels of situational awareness, from perception to comprehension, and then projection. And yet, uncertainty remains a persistent challenge that space operators will continue to face. Thus automated tools that aim to facilitate decision-making workflows must be able to work with uncertainties. Human operators must often rely on informed judgments about likelihoods, based on available data and knowledge of historical behaviors and patterns. This is the motivation for an automated tool providing similarity-based inference, to make informed judgments resembling the practices of human analysts, while taking advantage of data-driven methods to increase scalability, finding patterns and similarities amidst growing data sets.

In general terms, similarity-based inference is not new in SDA applications. There has been a great deal of successful work with algorithms that perform reasoning or data fusion on raw observational data, sometimes referred to as level-one analytics. In situation awareness terms, level-one analytics relate to perception and comprehension. A wide range of increasingly sophisticated tools use statistical and machine learning (ML) methods to interpret observational data (e.g., [1][2][3]). Pattern-of-life (PoL) analytics make inferences essentially based on the level of similarity to past data, to detect and characterize events and anomalies involving resident space objects (RSOs), such as maneuvers, separations, and changes in attitude, stability, radio frequencies, and other signatures. Other level-one algorithms consume observation data to do additional forms of characterization, such as identifying events or behaviors that are out of family, using algorithms that essentially compare an object with a class of related objects, either by rules, statistical and ML methods, or combinations of these.

This paper describes a tool called SHERLOC which implements an approach to similarity-based inference to help with level-two sensemaking, in order to complement the value of level-one tools. Where level-one tools are highly effective at taking observational data to go from perception toward comprehension, the goal for a level-two system is to augment comprehension and even work toward projection, by making inferences across combined information, including the outputs from level-one tools synthesized with other data sources. There have also been other efforts performing level-two analytics. For example, Putman et al describe an algorithm for cross-phenomenology data fusion using Bayesian classifiers for satellite identification [4]. This is essentially a level-two approach, moving up a level from the level-one analytics performed by tools operating on raw data within a single phenomenology. The satellite classifiers use a combination of kinematics, radar cross section, visual magnitude, passive radio frequency, and signal polarity data available from commercial providers via the Unified Data Library (UDL). This multi-source approach also yields insights in a sensitivity analysis for respective phenomenologies. In another particularly relevant effort, Indaco and Guzzetti use dynamic graph structures to represent proliferated LEO constellations as networks of large numbers of individual satellites with edges for relationships or interactions [5]. The graph

structure encodes correlations that help to identify anomalous behaviors within the constellation. The SHERLOC tool uses a similar graph-based method, but applied to a different SDA purpose.

While these efforts are related, in that they apply level-two methods for purposes such as classification, resolution, and anomaly detection, SHERLOC is designed to apply similar methods for sensemaking related to missions and threats. For example, soon after an RSO is placed in orbit from a non-cooperative launch, level-one tools are essential for helping to fill out an initial picture of the object's observable behaviors, but there may still be many unknowns about its mission, capabilities, and possible relationships to other objects. Essentially these additional factors support the reasoning for threat assessments and decisions related to developing courses of action. Human analysts often look for similarities to assess likelihoods for these unknowns. In order to fill knowledge gaps, consideration of similar satellites and their missions can lead to suggestive results by exploring questions like the following:

- Was there more or less secrecy than usual compared to similar past launches?
- Where was the launch site and what was the launch vehicle?
- What kinds of missions are more or less likely given the observed orbital parameters?
- Are there other potential sibling satellites already occupying a similar orbit, such as those from the same constellation or known to operate in coordinated groupings?
- Is this object a potential adversary in a coplanar orbit with a friendly object?
- Were there any noteworthy behaviors during ascent stages or separation from the apogee kick motor?
- How many objects were placed in orbit from the same launch, do they all have similar orbits, and is there any relevant or informative knowledge about the missions of other objects?
- What kinds of missions are more or less likely given observations about size / shape / attitude?

While several of these questions are informed by data from level-one tools, they can collectively help to make level-two inferences about an object based on similarities in these various dimensions. Some very basic questions, such as an estimate that the satellite's mission is either military or commercial, may be readily answered with high confidence. More nuanced details are likely to involve more complex reasoning and estimates with varying confidence levels. The SHERLOC inferential tool uses a relational network with nodes for space objects, launch events, launch vehicles, countries, orbital parameters, and other related features. Extended multi-source data include nodes for characterizations of dynamic events and verified behaviors. Relationships between nodes are represented as edges. The network uses a similarity function that works like a "friend recommender" in social network systems, based on the strength of edges between nodes, together with feature analysis within nodes. This allows for automated inferences in ways that are comparable to how human operators reason about unknowns regarding space objects. For one or more objects under investigation, by first identifying similar objects and any available knowledge related to the questions above, this gives insights about likely mission, capabilities, and potential threat.

Although this kind of sensemaking is less common in existing tools, at least one additional example is highly relevant. Witman et al describe an approach for assessing similarities between satellite behaviors by using action-free inverse reinforcement learning (AFIRL) in anomaly detection [6]. This is tested in a tool for managing time-series comparisons across heterogeneous space-related data sets. An example application is the detection of anomalous satellites within a constellation, by evaluating similarities to identify cases such as unusual mass, potentially indicating that a satellite is carrying a "hidden" payload unlike its siblings. This is very similar to the SHERLOC tool's goal of applying automated level-two analytics to help uncover unknowns about space objects that may be related to mission and threat concerns. Ultimately with the thread of uncertainty underlying many SDA challenges, the best way forward is likely to involve ensembles of automated tools that focus on different elements of the problem, and there's value in exploring commonalities between different approaches. For the development and testing of the SHERLOC tool described in this paper, the Apollo Accelerator program at the SDA TAP Lab has provided an extremely valuable collaborative environment for developing an ensemble of tools in this manner, working toward integration into a prototype battle management system [7][8]. One of the benefits is the interdependency allowing for a decomposition of problems so that common data sets and messaging protocols provide a baseline available across tools for different specialized algorithms. In addition to discussing the SHERLOC tool design, this paper also highlights examples where this collaborative environment has helped to streamline concept development.

2. METHODOLOGY

The SHERLOC tool is implemented with parallel high-level functional goals, to support integration with other components in a battle management system with automated machine-to-machine data flow, and to provide a user interface for operators and analysts to explore potential threats and also see the justifications for automated assessments about threats. It is constructed with back-end and front-end components, containerized and deployed on a hosted platform. The core capability is a relational network approach to representing space objects together with characterization data for object attributes and temporal data for behaviors. This is coupled with an inference engine for reasoning about space objects based on similarities to other objects and known threats, and for assessment and prioritization of potential hostility. Machine-to-machine automation is based on standardized platforms and methods for data throughput, such as the UDL and the database and messaging utilities under development in the SDA TAP Lab's battle management system. Integration with other tools includes the ingestion of catalog data as well as event data and indicators produced by upstream analytic algorithms that characterize meaningful behaviors such as a maneuver or an unusual photometric change. Similarity-based inferences are passed as outputs to downstream tools, to support a variety of use cases by providing insights about the following:

- Objects identified as similar to known threats.
- Likely values for unknowns, such as the likely identity or constellation for unnamed objects in a catalog, or likely capabilities and behaviors.
- Multi-object relationships such as functional siblings or newly launched payloads working in conjunction with existing objects on orbit.

Along with the implementation of the relational network and automated data flows, a human-machine interface supports operator investigation to visualize the network and see patterns or clusters of similarities. This provides explainability as well, by making the justification for inferences apparent in terms of which objects are similar or related, and why. Additional user interface utilities support filter, search, and customization of the feature-wise weights used in calculating similarity. The following subsections go into more detail about elements of the tool's preliminary implementation.

2.1 Source data

The SHERLOC system ingests both structured and unstructured data. The current implementation uses catalog data from Space-Track for its master list of objects, and then associates additional data with these objects from a variety of sources. One external component that the system benefits from is the Target Model Database (TMDB) currently being developed in the SDA TAP Lab, led by a team from Millennial Software. The TMDB serves several purposes, to collect "clean" catalog-related data and temporal data in a repository as a source of truth that can be used across integrated components and algorithms. When tools throughout the architecture produce messages specific to RSOs for temporal events such as maneuvers, separations, stability changes, radio frequency link changes, and various other indicators, these are preserved in the TMDB for later reference. As a central repository, the TMDB can also maintain a record of data pedigree for traceability in any downstream analytics that are based on the data.

With the variety of space object catalogs, there is a need for canonical names to help reduce mismatches for automated system analytics. Among its other functions, the TMDB helps to resolve synonymous or conflicting names, while also avoiding duplication of effort by different providers collecting data from multiple sources. For example, various data resources refer to launch vehicles using different variations of names and designators. Some of these variations simply involve different wildcard symbols or capitalizations (e.g., "Soyuz-2.1b/Fregat" vs. "Soyuz-2-1B Fregat"), while others involve alternatives for commonly used synonymous terms (e.g., "Chang Zheng 6A" vs. "CZ 6A" vs. "Long March 6A"). The practice of maintaining canonical names in a central repository is particularly useful for the SHERLOC tool as it runs similarity analytics on different objects and nodes in the network, relying on positive matches where appropriate.

The current system also ingests additional data from multiple providers in the UDL, including the Notices to Space Operators (NOTSOs) among others. A data product of the U.S. Space Command's Joint Commercial Operations (JCO) cell, the NOTSOs are messages containing curated information and analysis products characterizing space activity or events [9]. The NOTSOs are a useful information source for patterns of behaviors such as verified

maneuvers and rendezvous and proximity operations (RPOs), complementing temporal event data that will be available in the SDA TAP Lab's battle management architecture and the TMDB. In addition to the content within NOTSOs, considerations at one level higher in abstraction are useful for similarity-based analytics, where for examples the frequency and quantity of NOTSOs can be a broad indication of noteworthy activity levels for an RSO.

2.2. Relational network

The relational network is implemented as a persistent graph database, incorporating principles from Microsoft's GraphRAG approach. This enables complex relationship modeling, scalability, and an enhanced ability to extract insights from unstructured data. The knowledge graph construction forms the backbone of SHERLOC's analytical capabilities. The graph database's capabilities support efficient and complex queries and traversals. Community detection algorithms can be used to identify densely connected node groups, revealing hidden patterns and relationships in the space object network. The initial implementation includes nodes for RSOs, countries, launch events, launch vehicles, orbital regimes, object types, radar cross-section categories, mission type, shape categorization, and NOTSOs by event type. Object nodes contain further details with non-enumerated values such as orbital parameters (apogee, perigee, inclination, RAAN, and longitude in the case of objects in geostationary (GEO) orbit). Future development plans include expanding the information and attributes in the knowledge graph, with multi-modal data to help uncover complex relationships across many factors.

As other efforts have found, there is value in using hierarchical classes for features. Ruttenberg et al describe a probabilistic programming approach to classifying RSOs, using a hierarchical taxonomy for space objects, characterized by features such as orbital regime, attitude, mass, brightness, mission, etc., with sub-classes as needed [10]. The application in this instance was for automated reasoning to triage sensor resources. A hierarchical organization for discerning features has similar utility in the SHERLOC tool, where some node types are associated with hierarchical classes in order to allow for more refined similarity analyses, beyond simple comparison lookups. For example with satellite vehicles, even with the use of canonical names to counteract inconsistencies from different data sources, there are still more detailed factors that are useful to consider in automated comparisons. In the publicly available website Jonathan's Space Pages maintained by Jonathan C. McDowell, there are proposed groupings of launch vehicles into families based on the first stage and various parameters [11]. There are cases where differently named launch vehicles in fact use the same first stage rocket and have similar expectations for max apogee for objects placed in orbit. Thus, it makes sense for a machine-based representation of launch vehicles to capture hierarchical classes that can represent these kinds of groupings. The TMDB maintains a hierarchical representation for launch vehicles, and this is used in the SHERLOC relational network.

The current implementation of the relational network primarily uses edges to represent associations between nodes and launch events, and between RSOs and their unique attributes. So there are edges between a launch event node and the nodes for the country of origin, launch site, launch vehicle, and objects in orbit. The nodes for RSOs have additional edges to nodes for other information uniquely associated with the objects themselves. Where no information is available, for example regarding an object's mission type, no node or edge is added connecting to the object. Future enhancements are anticipated to capture other kinds of relationships such as formation flying, sibling and cooperative functionality, other multi-node events beyond launches, operational relationships such as communication links to ground control stations, and temporal relationships such as phases of expected lifetime.

2.3. Similarity analysis engine

The similarity analysis engine uses a multi-faceted approach to determine space object similarity, considering network connections in the graph, together with other information and calculations using additional object attributes (both enumerated and non-enumerated values). A network similarity component evaluates the effective resistance between nodes in the graph. Other specific factors include a calculation of orbital similarity based on non-enumerated parametric values (inclination, apogee, perigee, etc.), and other similarities associated with any known specific identifying information such as country, object type, size, shape, event histories reflected in NOTSOs, mission, and use category (government, military, civil, commercial). An additional element in progress but not complete as of this writing involves the incorporation of additional dynamic and temporal data produced by upstream tools in the SDA TAP Lab's battle management system, such as maneuvers, separation events, and changes in stability, RF links or attitude. In particular, the system architecture involves a subsystem dedicated to the generation of indicators related to camouflage, concealment, deception, and maneuver (CCDM), which will be

posted to a system-wide message bus and also captured in the TMDB. Some of this kind of information about dynamic events is also available via NOTSOs, but the CCDM indicators have a valuable potential additional role in SHERLOC's similarity-based inference. Taking all available factors in the relational network, the similarity engine combines factors in an algorithm using a weighted average to produce an overall similarity score.

The weights of factors used in similarity scoring can be adjusted to emphasize different aspects of similarity based on the specific analysis needs. For example, considering the factor for an RSO's associated object type (payload, rocket body, or debris), there are different conditions where it may be desirable to assign this factor a higher or lower weight for similarity scoring. In a general case such as using similarities to predict unknowns about a space object, the existing type designations may be useful, perhaps for the goal of comparing payloads to payloads. However, in an alternative case looking to uncover hidden relationships between objects, it may make sense to reduce the weight for the existing object type values, perhaps to find unusual cases where an object tagged in the catalog as one type (e.g., rocket body) has in fact exhibited similarities to another type (e.g., payload), based on other evidence factors and potentially behaviors.

SHERLOC allows for different configurations of factor weights to be used in the similarity engine, which can also be specified or modified in the operator interface, as shown in Fig. 1. In this configuration, individual factors are assigned to tiers, allowing for adjustments to individual weights across factors within a tier, and then respectively across overall tiers.

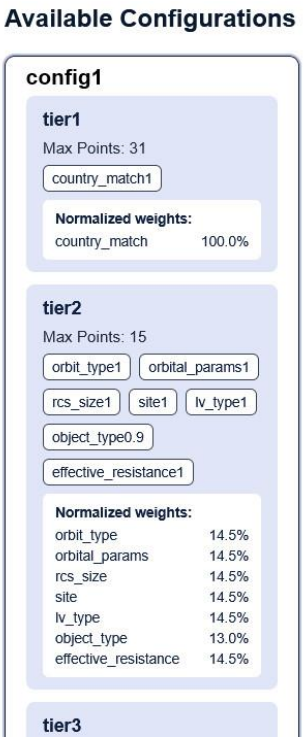

Fig. 1. Partial view of the SHERLOC operator interface for configuring factor weights used in similarity analysis

2.4. Machine-to-machine data flow

As an automated tool, SHERLOC is specifically designed to work with machine-to-machine data flows to maximize scalability and reduce decision times by identifying patterns and similarities amidst growing data sets. SHERLOC is integrated with the UDL as a standard platform for space data, both as a consumer of data and a service provider writing to the Notifications endpoint with information about objects similar to known threats. Within the SDA TAP Lab, the battle management system architecture has been a highly effective context for development and testing. Initially this began as a collaborative environment with cross-pollination among an ensemble of tools and algorithms addressing different SDA challenges. But as it has evolved into a unified system architecture concept, this adds the value of factoring realistic integration requirements into the early stages of individual tool development.

The emerging battle management system architecture includes an Apache Kafka message bus for distributed data streaming. The message bus uses topics with defined schemas, where messages conforming to these topics can be written and read by different system components. SHERLOC does both, consuming messages on topics populated by upstream tools, and currently writing to one topic with a CCDM indicator to identify objects that are similar to known threats, for downstream tools. The cadence for SHERLOC's message publishing essentially starts with an initial state for any object that hasn't been updated on the message bus, followed by future updates as conditions change (new objects, changes in the status of comparison objects as threats or objects of interest, etc.). In addition to the interface with the message bus, SHERLOC also ingests data from the TMDB discussed above. This initial integration will continue to evolve, with additional input and output messages as both the SHERLOC tool and the overall battle management system architecture undergo further development.

2.5. Operator interface

In addition to the messaging and data flows for machine-to-machine operations, the user interface also supports operator visualization and browsing of the network. Elements of the main visualization are shown in Fig. 3 and described below.

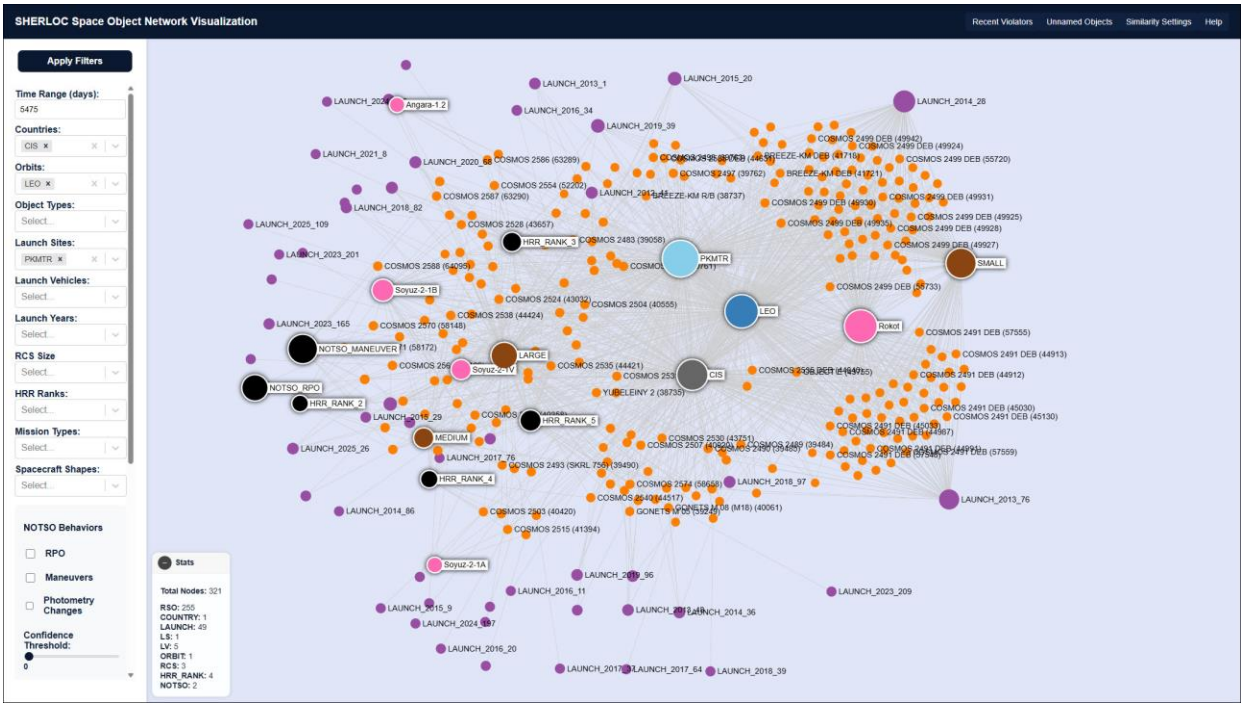

Fig. 2. SHERLOC operator interface

In the view shown above, the network visualization occupies the main area of the screen, with a filtering panel on the left. Users may apply multiple filters simultaneously, with the graph updating in real-time as filters are adjusted. The current view above is filtered to show Russian (CIS) objects in low earth orbit (LEO) originating from the Plesetsk launch site (PKMTR), that have not decayed. There are 255 object nodes in this filtered view, and 321 total nodes. The network visualization supports zooming, panning, and node selection for detailed exploration. When a node is selected, it highlights connected nodes and edges, enhancing the visibility of relationships. The visualization implements color-coding based on node types as follows:

Resident space object (RSO)	orange
Orbital regime	blue
Country	gray
Launch event	purple
Launch site	light blue
Launch vehicle	pink
Radar cross section	brown
NOTSO event behavior	black
High-rate-revisit (HRR) rank	black

The relative positions and distribution of nodes for display purposes are handled automatically and do not convey specific meaning, other than the fact that nodes with many edges tend to be placed in central locations. Node sizes are also related to the number of edges; those with many connections are larger.

When a node is selected, a details panel is displayed on the right side, to present specific information about the node. The nodes for RSOs have the most information to be presented in this panel, as shown in Fig. 3 below.

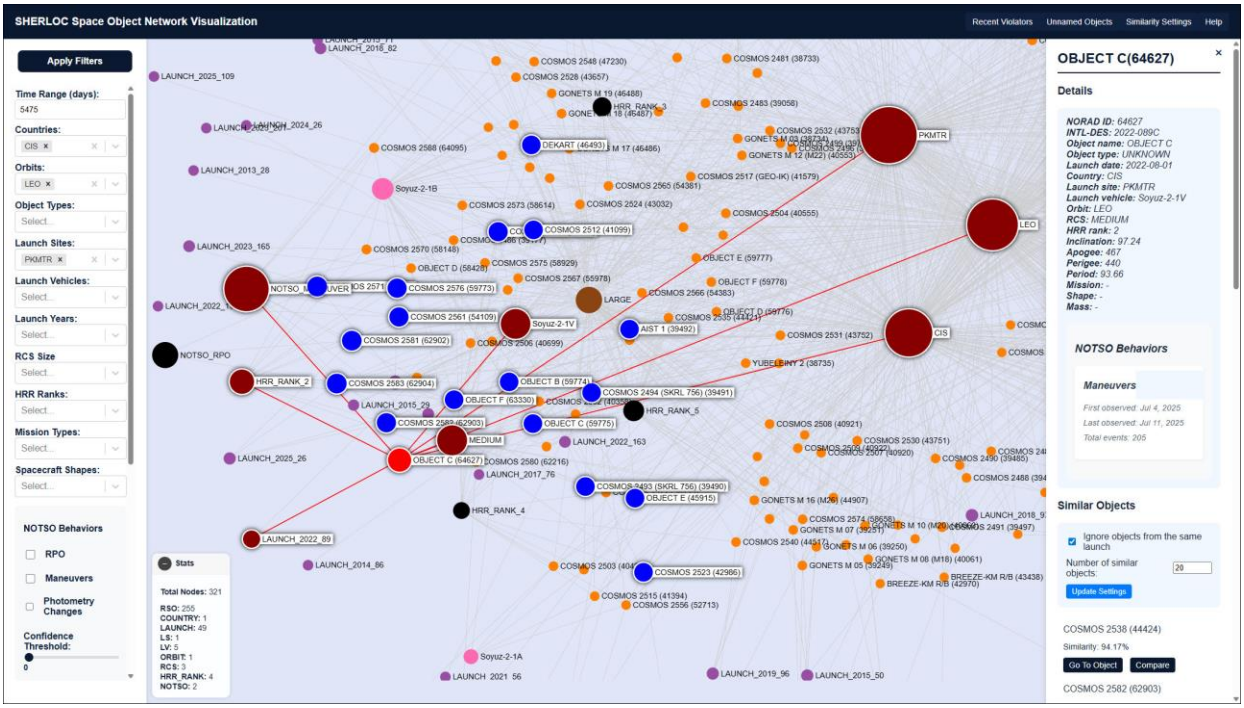

Fig. 3. SHERLOC operator interface, with object 64627 ("OBJECT C" in the catalog) selected

When a node is selected, its display color is changed to red, and the edges to connected nodes are also shown in red, while the connected nodes are also highlighted and shown in burgundy.

One of the primary use cases is to learn about object similarities, to draw inferences. When an RSO is selected, the nodes for the top most similar objects are highlighted in royal blue, as shown above. These are also displayed in list form in the details panel, in descending order starting with the objects returning the highest similarity scores. The case shown above is discussed further in the following section as an example use case.

3. EXAMPLE INFERENCE

This section walks through an example of SHERLOC's similarity-based inference, using the operator interface for illustration purposes. This example involves the same object shown selected in Fig. 3 above, assigned the ID and name 64627 / OBJECT C in the Space-Track catalog from the 18 Space Defense Squadron.

To set the stage for this example, 64627 / OBJECT C was released on June 26 2025 from Cosmos 2558 (53323). The latter, Cosmos 2558, has been an object of interest since its launch in August 2022, when it was placed into a coplanar orbit with a US satellite. Analysts believe Cosmos 2558 is in the Russian Nivelir family of inspector satellites. Past Nivelir satellites Cosmos 2519 and Cosmos 2542 have exhibited behaviors of concern, including the release of sub-satellites in events considered tests of anti-satellite (ASAT) weapon capabilities [12]. Analysts believe Cosmos 2558 is likely in the same family as these prior Nivelir satellites, based on a variety of similarities including its initial launch coplanar to the US satellite. However, Cosmos 2558 maintained orbit for nearly 3 years without releasing any sub-satellites. Thus when 64627 / OBJECT C separated from Cosmos 2558 in June 2025, this was already a situation deserving close attention. Initially there was uncertainty about whether 64627 / OBJECT C was a piece of debris or a payload, but when a series of maneuvers were detected soon after, this provided greater certainty that 64627 / OBJECT C is a sub-satellite.

Fig. 4 below explores this example closer in the SHERLOC operator interface, with the view in the network visualization expanded, and the details panel scrolled down to see the list of objects found to be similar to 64627 / OBJECT C.

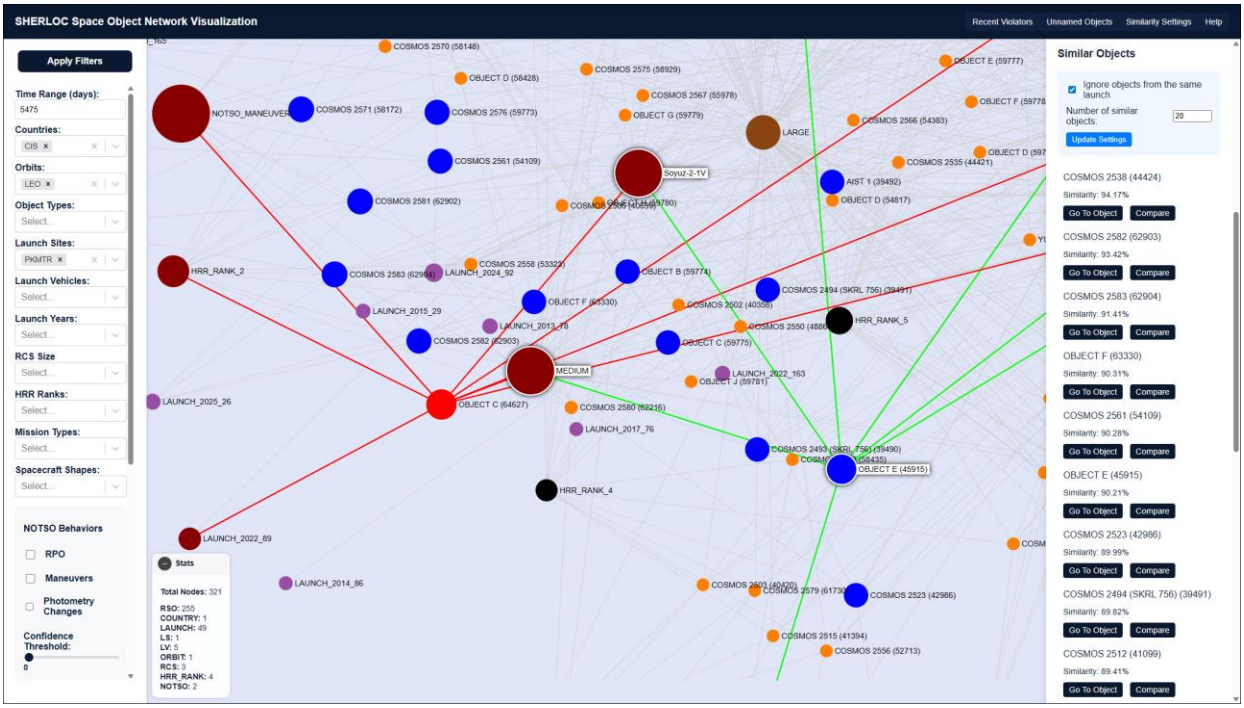

Fig. 4. SHERLOC operator interface, viewing similar objects to 64627 / OBJECT C, with current focus on 45915 / OBJECT E

In the list of similar objects for 64627 / OBJECT C (right side panel), the sixth item is 45915 / OBJECT E. The node is also highlighted in the view of the network visualization above. Although this object is assigned a generic name in the catalog like many others, it is actually quite significant. 45915 / OBJECT E is the sub-satellite projectile released from Cosmos 2543 (itself a sub-satellite of Cosmos 2542) in July 2020. In fact it was ejected at high velocity, and this was the event considered an ASAT test. This similarity is automatically identified by the SHERLOC tool, paralleling the assessments of human analysts. In a variety of operating situations, this kind automated inference has the potential to add value, for example in this case at the earliest stage after 64627 was released from its parent. Incidentally, there are also NOTSOs relevant to this example that provide a valuable data source, as these were issued for a series of verified maneuvers conducted by 64627 / OBJECT E between July 4 and July 11, supporting the reasoning that it is a payload not debris. The behavioral data from NOTSOs are also currently represented in the SHERLOC relational network.

Also noteworthy are Cosmos 2582, Cosmos 2583, and another unnamed object 63330 / OBJECT F (listed consecutively second through fourth in the Similar Objects list, in Fig. 4 above). Analysts consider these to be objects of interest for similar reasons based on similar patterns. Cosmos 2581, 2582, and 2583 were launched into orbit in February 2025. Although they were not launched into an orbit coplanar with a US satellite, they do have similarities to objects of interest, and possibly the Nivelir satellites. The trio were involved in a series of RPOs with each other, and Cosmos 2583 released 63330 / OBJECT F in March 2025. Notably, SHERLOC identifies this object 63330 as similar to both 45915 / OBJECT E and also the more recent 64627 / OBJECT C sub-satellites.

The same information visible in the operator interface is also available in the machine-to-machine data flow to downstream tools for a battle management system.

4. DISCUSSION AND CONCLUSION

The foregoing examples are intended to illustrate how automated similarity-based inference can be used to help uncover unknowns about potential missions or capabilities, similarities to objects of interest and potential threats, and also the relationships between multiple objects.

Although this paper presents concepts, methods, and use cases, there is also merit in considering risks. As with any automated tools, or more broadly the increasing use of artificial intelligence to augment human decision-making, care must be taken in terms of the trust placed in autonomy. Probabilistic approaches such as similarity-based inference are intended to help make best use of data in the analysis of situations involving uncertainty, which is the hallmark of SDA decision-making. With the SHERLOC tool or any solution offering an assessment of likelihoods, it should not be either presented or received as a source of certainty. Although human analysts use analogies and similar cases to reason under uncertainty, they still encounter completely new situations that fail to match expectations derived from prior cases. By the same token, automated systems may help to facilitate similarity-based or experience-based reasoning with the power to process larger data sets, but this does not make them invincible in the case of events that lack precedents. As long as these limitations are understood by both the tool provider and the end user to operate under the proper context, automated analytics can be extremely useful. Also, an ensemble of tools can help combat brittleness for different operating conditions, as different solutions may have different strengths and weaknesses to help work around new or unexpected situations.

The preliminary SHERLOC tool is operational and initially integrated with the developmental battle management system at the SDA TAP Lab, and yet significant work remains in terms of algorithm refinements, data ingestion, and integration with workflows. The following are some areas of planned future work:

- Dynamic threat lists or objects of interest. The existing implementation uses a simple representation for threats or objects of interest, with reference lists for identifying objects similar to known threats. Development plans include refining this approach to consume more dynamic data, including the CCDM indicators for objects of interest, produced by upstream tools in the SDA TAP Lab's system architecture.

- Similarities based on parametric input. The methods used for similarity-based inference within the SHERLOC tool to date have an underlying assumption that comparisons are being made between two (or more) existing catalog objects. The same inferential capabilities can be applied to prospective or hypothetical objects by simply using anticipated parametric inputs, in place of the requirement for an existing catalog object. For example, before a launch, publicly available information may identify the launch site, launch vehicle, likely ascent trajectory, and potentially even the payload, long before there is ever a catalog object or orbital state. This is another situation where human analysts often reason from past cases, so by the same token this is an opportunity for automated similarity-based inference to help.

- Multi-object relationships. By encoding a variety of multi-object relationships in the relational network, this can help predict new cases based on similarity. Multi-object relationships include parent / child relationships resulting from spawning events, link and communications relay relationships, sibling relationships based on mission, constellation, and/or specific behaviors like formation flying, and others.

- Dynamic event data. A certain amount of dynamic event data is already ingested into the relational network in the form of NOTSOs and the patterns of behaviors they represent for objects. Development plans include supplementing the NOTSOs with more granular and timely dynamic event data from upstream tools in the SDA TAP Lab's battle management system architecture.

- Time series data across events and behaviors. Building on dynamic event data, it is also useful to consider behavior patterns across discrete events, or even across the lifetimes of different objects. If a certain class of objects frequently makes an attitude change before a maneuver, these kinds of temporally associated events can be captured as behavior signatures to further refine similarity-based analytics. Time-series analysis methods and sequence alignment techniques can be used to detect recurring behaviors and deviations from expected patterns.

- Explainability. One of the benefits of a similarity-based inferential approach is that the retrieval of specific similar cases or objects provides a basis for explainability. Automated reasoning drives the characterization that an object is similar to known threats, or likely to have certain capabilities or mission characteristics or relationships to other objects. This reasoning can then be explained in terms of the retrieved similarities, with confidences based on the sources of data supporting similarities, and articulated in a human-readable summary of the basis for the algorithm's results.

5. ACKNOWLEDGEMENTS

The research reported in this document was performed in connection with contracts for the U.S. Air Force. The views, opinions, and conclusions expressed in this paper are solely those of the authors and do not reflect the views or policies of the organizations or institutions with which they are affiliated with or their sponsors or customers.

The authors are grateful for the support of the Air Force Research Laboratory Information Directorate in conducting this research. The authors are also grateful to the SDA TAP Lab and its leadership and staff, for providing an extremely constructive environment for rapid concept and tool development, and to the other teams participating and collaborating in this environment with enthusiasm and a remarkable range of technology and space domain expertise. Additionally the authors are grateful to subject matter experts from Patchplus Consulting who contributed valuable knowledge from operational space domain experience.

6. REFERENCES

[1] M Jah, and R Madler. Satellite Characterization: Angles and Light Curve Data Fusion for Spacecraft State and Parameter Estimation, *Proceedings of the 8th Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference*, Maui, HI, 2007.

[2] T Kelecy, D Hall, K Hamada, and D Stocker. Satellite maneuver detection using Two-line Element (TLE) data, *Proceedings of the 8th Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference*, Maui, HI, 2007.

[3] R Linares and R Furfaro. Space objects maneuvering detection and prediction via inverse reinforcement learning, *Proceedings of the 18th Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference*, Maui, HI, 2017.

[4] T Putman, C McComas, and J Faustman. Multi-Phenomenology Fusion for Satellite Identification, *Proceedings of the 25th Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference*, Maui, HI, 2024.

[5] M Indaco and D Guzzetti. Transformer-based anomaly detection in p-leo constellations: A dynamic graph approach, *Acta Astronautica*, vol. 218, pp. 177–194, 2024.

[6] D Witman, T Olson, B Williams, D Kesler, and B Marchand. Action-Free Inverse Reinforcement Learning for Evaluating Satellite Similarity and Anomaly Detection, *Proceedings of the 25th Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference*, Maui, HI, 2024.

[7] G Furlich, A Crews, J McGuigan, T McLauglin, C Burns, P Balster, G Jones, G Hofer, IS Bartlett, L Hettiarachchi, H Spolar, and SP Allen. Automated, Collaborative Applications to Close Kill-chain Gaps, *Proceedings of the 25th Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference*, Maui, HI, 2024.

[8] SP Allen. SDA TAP Lab Using Commercial Technology to Avoid Operational Surprise, *Proceedings of the 25th Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference*, Maui, HI, 2024.

[9] B Golf and A Konnath. Joint Commercial Operations (JCO) Introduction and Way Forward, *Proceedings of the 25th Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference*, Maui, HI, 2024.

[10] BE Ruttenberg, MP Wilkins, and A Pfeffer. Reasoning on Resident Space Object Hierarchies Using Probabilistic Programming, *Proceedings of the 18th International Conference on Information Fusion*, Washington, DC, 2015.

[11] JC McDowell. GCAT: General Catalog of Artificial Space Objects – Launch Vehicles, URL https://planet4589.org/space/gcat/web/lvs/index.html [Online; accessed July 2025]

[12] B Weeden, and V Samson (eds). Global counterspace capabilities: An open source assessment, *Secure World Foundation*, URL https://swfound.org/media/207541/swf_global_counterspace_capabilities_2023_es_en.pdf [Online; accessed October 2023]