Case Study: SATWATCH AI Platform for Enhanced Space Situational Awareness
Client: United States Space Force (USSF)
Industry: Defense / Military / Space Operations
Location: USA (Global/Space Operations)Client: United States Space Force (USSF)

Executive Summary
As the newest branch of the U.S. Armed Forces, established to focus explicitly on the space domain, the United States Space Force (USSF) is charged with organizing, training, and equipping forces to protect U.S. and allied interests in space and to provide space capabilities to the joint force. Operating within an environment characterized by rapidly increasing complexity, congestion, and strategic global importance, the USSF's mission is fundamental not only to national security but also to the stable functioning of modern global society [1]. Critical infrastructure worldwide relies heavily on space-based assets for essential services including precise Positioning, Navigation, and Timing (PNT) via GPS, global communications, financial transactions, weather forecasting, environmental monitoring, and vital intelligence gathering [2]. Maintaining comprehensive **Space Situational Awareness (SSA)**—possessing a deep, accurate, and timely understanding of the location, orbit, characteristics, capabilities, status, and intent of objects in space—is therefore a foundational requirement [3]. Recognizing the inherent limitations of existing systems in this dynamically evolving landscape, the USSF, potentially through organizations like Space Operations Command (SpOC) and its subordinate units such as the 18th Space Defense Squadron (SDS) responsible for space domain awareness, required a transformative leap in its SSA capabilities. The objective was to move beyond simple catalog maintenance towards proactive identification of risks, characterization of activities, prediction of future states, and ultimately, ensuring continued U.S. space superiority and the stability of the space environment [4].

Challenge: Navigating the Complexities of the Modern Space Domain
The space domain is no longer the vast, relatively empty expanse of previous decades but a complex, congested, and increasingly contested operational environment. It is densely populated with thousands of active satellites performing diverse missions, a growing population of defunct spacecraft, and hundreds of thousands of pieces of fragmentation debris ranging in size from large rocket bodies to minuscule but potentially lethal paint flecks [5]. Simultaneously, space is becoming increasingly contested, with potential adversaries developing and demonstrating capabilities ranging from sophisticated electronic warfare (jamming, spoofing) to kinetic anti-satellite (ASAT) weapons, alongside more subtle counter-space activities like close approaches or dazzling of optical sensors [6]. This dynamic reality presented significant, multifaceted challenges for the USSF's existing SSA infrastructure, demanding a more advanced, integrated, and intelligent system capable of proactive monitoring, rapid characterization, accurate prediction, and timely warning. Key hurdles that necessitated a paradigm shift included:

- **Data Diversity, Volume, Velocity, and Veracity:** The USSF relies on a globally distributed, multi-phenomenology network of sensors operated by the Department of Defense (DoD), intelligence agencies, allied nations, and increasingly, commercial SSA providers [3]. This network includes powerful ground-based radars (like the Space

Fence), sensitive optical telescopes (both ground and space-based), passive Radio Frequency (RF) detection systems listening for satellite emissions, and other specialized sensors. Integrating data from these disparate sources presents a monumental data science challenge. Each sensor type has unique characteristics: radars excel at range and range-rate measurements but may struggle with smaller objects or specific orbital regimes; optical sensors provide precise angular measurements but are limited by weather, daylight, and object brightness; passive RF provides valuable functional insights but less precise positional data [7]. Data arrives in unique formats (e.g., state vectors, Two-Line Element sets (TLEs), raw observations like range/azimuth/elevation or right ascension/declination), at different reporting cadences, with varying measurement sensitivities, diverse coordinate systems (ECI, ECF, etc.), and inherent inaccuracies or biases stemming from factors like timing errors, atmospheric distortions affecting optical and radar propagation, and sensor calibration drifts [8]. Correlating sparse or intermittent observations from different sensors to maintain accurate tracks on tens of thousands of objects is computationally intensive. Conflicting reports between sensors or persistent **data gaps** (where objects are not observed frequently enough) can lead to an incomplete, ambiguous, or erroneous operational picture, critically hindering timely and accurate decision-making regarding potential threats or collisions. Furthermore, the sheer **volume** of data generated by these sensors, potentially reaching petabytes requiring sophisticated storage and processing infrastructure, coupled with the high **velocity** of real-time updates, presented significant computational challenges [9]. Ensuring the **veracity** (accuracy and trustworthiness) of data from diverse sources, including commercial or international partners with varying quality control standards, added another layer of complexity.

- **Subtle Anomaly Detection:** Traditional SSA methods, often reliant on established orbital propagation models (like SGP4 for TLEs or high-fidelity numerical integrators) and Kalman filtering techniques, excel at tracking known objects along predictable Keplerian or perturbed paths [8]. However, they inherently struggle with detecting faint signals or subtle deviations buried within noisy data streams. Identifying these subtle anomalies is increasingly critical. Examples include:
    - *Low-Thrust Maneuvers:* Extremely low-thrust, continuous maneuvers using electric propulsion are increasingly common for station-keeping, orbit raising, or potentially covert repositioning. These produce very gradual changes in orbital parameters that are difficult to distinguish from natural perturbations (like atmospheric drag variations or solar radiation pressure effects) using classical orbit determination methods [10].
    - *Rendezvous and Proximity Operations (RPO):* Detecting satellites making subtle adjustments to approach or maintain proximity to another space object requires highly precise tracking and the ability to differentiate intentional maneuvering from natural orbital drift.
    - *Non-Kinematic Changes:* Unexpected changes in communication patterns, RF signal characteristics (frequency, modulation), optical brightness fluctuations (indicating tumbling, venting, or changes in configuration), or faint thermal

signatures (potentially indicating component stress or imminent failure) can be crucial indicators of an object's status or intent but are often missed by systems focused solely on orbital mechanics [7].

  - Minor Deviations: Small station-keeping adjustments slightly outside normal operational parameters or minor deviations from expected behavior patterns might be precursors to larger issues or deliberate actions.

  These subtle events, often possessing a low signal-to-noise ratio (SNR) or mimicking natural orbital perturbations or sensor noise, can be crucial precursors to significant threats (e.g., an impending hostile action), impending system failures (e.g., a satellite breaking apart), or deliberate covert activities. Yet, they are easily lost in the noise or dismissed as sensor error by conventional threshold-based detection systems, potentially leading to strategic surprise or preventable satellite losses.

- **Predictive Capabilities for Anticipatory Awareness:** Truly effective SSA requires evolving beyond forensic analysis (understanding *what happened*) and real-time monitoring (understanding *what is happening now*) towards robust **predictive analysis** (accurately anticipating *what will likely happen next*). The USSF needed the capability to reliably anticipate potential **conjunctions** (close approaches between tracked objects that could lead to catastrophic collisions) well in advance, ideally with sufficient lead time for mitigation planning [11]. This prediction capability needed to be particularly robust for conjunctions resulting from intentional maneuvers (detected or predicted) or unexpected satellite fragmentation events, which can dramatically alter the collision risk landscape in minutes. Furthermore, predicting the likely *intent* or *future behavior* based on observed patterns (e.g., predicting the final target orbit of a maneuvering satellite) is strategically crucial. This predictive power enables proactive collision avoidance planning (often requiring international coordination), allows for more accurate assessment of the intent behind observed actions (distinguishing routine operations from potential threats), facilitates optimized sensor tasking (focusing resources on predicted high-risk events), and ultimately supports maintaining an advantageous operational posture in space by shifting from reactive responses to proactive, anticipatory actions [4].
- **Scalability and Timeliness in a Growing Domain:** The number of tracked objects in Earth orbit is growing exponentially, driven by the deployment of large commercial satellite constellations (like Starlink, OneWeb), the proliferation of smaller satellites (CubeSats), increased global launch rates, and occasional debris-generating events (collisions or ASAT tests) [5], [12]. Any new SSA system needed to be inherently **scalable**, architected using modern software principles and infrastructure to handle this dramatic growth in data volume, object count (potentially hundreds of thousands or millions of tracked objects in the future), and computational complexity without performance degradation or spiraling operational costs. Furthermore, the highly dynamic nature of space operations, especially in congested orbits like LEO, demands **near real-time processing and analysis**. Actionable intelligence regarding potential collisions, imminent threats, or

newly detected objects must be delivered to operators and decision-makers within tactically relevant timelines, often measured in minutes or hours, as delays of even a few hours could render the information obsolete, negate opportunities for effective collision avoidance, or allow a threat to go unchallenged [9]. Meeting these dual requirements of massive scalability and stringent timeliness posed a significant architectural challenge.

Solution: The SATWATCH AI Platform - Fusing Data and Intelligence

To overcome these formidable challenges, the **SATWATCH AI platform** was meticulously developed and engineered by 577 Industries, leveraging their unique convergent expertise in Artificial Intelligence, advanced data science, high-performance computing, and astrodynamics. SATWATCH represents a paradigm shift in SSA processing and analysis, functioning not as a new sensor system, but as an advanced **analytical engine** or "brain" that intelligently ingests, processes, fuses, and analyzes multi-source data. It employs a suite of cutting-edge AI and Machine Learning (ML) techniques, specifically adapted for the space domain, to provide deeper insights, more reliable predictions, enhanced anomaly detection, and ultimately, actionable intelligence far exceeding the capabilities of traditional systems.

- **Advanced Data Fusion Engine:** At its core, SATWATCH incorporates a sophisticated data fusion engine designed explicitly to handle the complexities, uncertainties, and multi-modal nature of SSA data [13]. This engine moves beyond simple data aggregation or basic track correlation. It intelligently processes, cleanses, correlates, and synthesizes information from diverse sources using advanced algorithms:
    - *Sensor Bias Correction:* Employs techniques to estimate and correct systematic errors or biases inherent in different sensor measurements (e.g., timing offsets, range biases).
    - *Rigorous Data Cleansing:* Utilizes statistical methods and consistency checks to identify and mitigate outliers, noise, and potentially erroneous data points within sensor feeds.
    - *Precise Time Synchronization:* Implements algorithms to accurately align observations from disparate sources onto a common time frame, critical for accurate state estimation and fusion.
    - Robust Uncertainty Quantification: Associates rigorous uncertainty metrics (e.g., covariance matrices) with all observations and state estimates, providing a probabilistic understanding of the data's reliability [8].

      By intelligently correlating observations from multiple sensors (radar, optical, RF) and diverse data types (kinematic measurements, brightness, RCS, signal characteristics, potentially telemetry where available) and dynamically weighting them based on factors like assessed sensor reliability, observation geometry (e.g., giving higher weight to observations taken perpendicular to the direction of motion), data timeliness, and quantified uncertainty, the fusion engine constructs a unified, high-fidelity, and probabilistic understanding of the state (position, velocity, characteristics) of each object in the space environment. This process effectively resolves ambiguities inherent in single-sensor views, fills observational gaps by leveraging complementary sensor coverage, and provides crucial confidence levels associated with its state estimates and predictions [14].

- **Synergistic AI/ML Techniques for SSA:** SATWATCH's analytical power stems from its integrated suite of AI/ML algorithms, not applied generically, but specifically chosen, adapted, and synergistically combined to address the unique physics and data characteristics inherent to the space domain:
  - *Advanced Trajectory Prediction:* Going significantly beyond standard orbital mechanics propagators (like SGP4, often used for TLEs but known to have limited accuracy, or basic numerical integration with simplified force models), SATWATCH utilizes cutting-edge techniques:
    - **Physics-Informed Neural Networks (PINNs):** These models embed the fundamental physical laws governing orbital motion (e.g., N-body gravitational interactions, atmospheric drag models, solar radiation pressure equations) directly into the neural network's learning process, typically by adding physics-based residual terms to the loss function [15], [16]. This ensures predictions remain consistent with physical reality, improves generalization to unseen scenarios, and potentially requires less training data compared to purely data-driven models for complex dynamics.
    - **Advanced Sequential Models:** Networks like **Long Short-Term Memory (LSTM)** or **Transformers** are employed, as they excel at learning complex temporal dependencies directly from observational data [9]. These models can capture subtle, non-linear dynamics influenced by factors like unpredictable atmospheric drag variations (driven by space weather), complex solar radiation pressure effects (dependent on satellite shape and orientation), third-body gravitational perturbations (from the Moon and Sun), and even potentially inferring the presence and effect of continuous low-thrust propulsion signatures [10]. The combination results in significantly more accurate long-term trajectory predictions compared to traditional methods, especially crucial for objects exhibiting non-cooperative, anomalous, or highly perturbed behavior.
  - *Multi-Modal Time-Series Anomaly Detection:* SATWATCH analyzes the diverse time-series data streams associated with each space object (orbital parameters derived from tracking, telemetry readings if available, optical brightness fluctuations, RF emission characteristics) using a combination of unsupervised and model-based techniques to detect deviations from expected norms:
    - **Unsupervised Learning:** Models like **Autoencoders** learn a compressed representation of the 'normal' operational baseline for each object (or class of objects) and flag significant reconstruction errors as potential anomalies [17]. **Isolation Forests** efficiently identify outliers in the high-dimensional feature space derived from the multi-modal data [7]. **Gaussian Mixture Models (GMMs)** can detect subtle shifts in behavior

patterns or transitions between different operating modes. These unsupervised methods are crucial for detecting *unknown* or *novel* anomalies for which specific signatures have not been pre-defined.

- Model-Based Detection: Anomalies are also detected by comparing observed behavior against the high-fidelity predictions generated by the advanced trajectory and behavior models (like PINNs or LSTMs). Significant deviations between observation and prediction trigger alerts.

  This multi-pronged, complementary approach allows for the detection of a wider spectrum of anomalies, ranging from sudden, impulsive events (like a large maneuver or breakup) to gradual performance degradations (like a slow drift off-station or increasing satellite tumble) or subtle changes in operational status (like unexpected transmitter activation/deactivation).

- *Behavioral Pattern Recognition:* Trained on vast, curated datasets encompassing decades of historical SSA observations combined with high-fidelity physics-based simulations representing various operational scenarios, supervised classification algorithms (such as **Support Vector Machines (SVM)**, **Random Forests**, and **Gradient Boosting Machines** like XGBoost) learn to recognize the characteristic multi-dimensional signatures associated with specific events and behaviors [18]. This includes accurately identifying different types of **orbital maneuvers** (e.g., distinguishing between routine station-keeping burns, inclination changes, phasing maneuvers for constellation management, drag make-up burns, and potentially more concerning rendezvous or intercept trajectories) and, crucially, distinguishing these intentional actions from natural orbital perturbations, sensor noise artifacts, or signatures potentially associated with system failures or uncontrolled tumbling based on subtle patterns across kinematic, photometric, and RF data streams.
- *Enhanced Conjunction Prediction and Assessment (CA):* By tightly integrating the high-accuracy trajectory predictions generated by the physics-informed models with the outputs from the maneuver detection and behavioral pattern recognition modules, SATWATCH forecasts potential conjunctions (close approaches) with significantly greater confidence and substantially **longer predictive lead times** (often extending reliable predictions further into the future) than traditional CA screening methods that rely on less accurate orbit propagation [11]. Critically, SATWATCH can specifically predict and flag high-risk close approaches that are identified as the likely consequence of detected ongoing maneuvers or anticipated future maneuvers (based on recognized patterns), providing crucial context for risk assessment. This intelligent filtering dramatically **reduces the false alarm rate** associated with conjunction warnings (which plagues traditional systems, leading to alert fatigue and unnecessary avoidance maneuvers), allowing operators to focus on the genuinely highest-risk events [4].

Implementation: Integrating SATWATCH into the USSF Ecosystem
Deploying the sophisticated SATWATCH platform within the USSF's secure, complex, and mission-critical operational environment was a significant systems engineering and integration undertaking. It involved careful planning, rigorous testing methodologies, phased execution, and continuous collaboration across several key dimensions:

- **Challenging Multi-Source Data Integration:** Establishing robust, secure, and high-bandwidth interfaces to ingest near real-time data streams from a highly heterogeneous mix of sources—including legacy DoD radar and optical systems, modern sensors like Space Fence, allied partner contributions, commercial SSA providers, and potentially national intelligence assets—required substantial systems engineering effort. Developing sophisticated **data normalization and transformation pipelines** was essential to handle the diverse array of formats (e.g., TLEs, OMMs (Orbit Mean-Elements Messages), VCMs (Vector Covariance Messages), proprietary telemetry formats), varying coordinate systems (e.g., ECI, ECF, TEME), differing units of measurement, and often stringent, differing security protocols and data handling caveats (including the complex task of securely integrating classified and unclassified data streams using accredited cross-domain solutions). Ensuring data integrity, establishing clear **data lineage** (tracking data origin and transformations), and maintaining **data provenance** (documenting data history and quality) throughout the entire processing pipeline was paramount for building trust and ensuring auditability [9], [12].
- **Rigorous AI Model Training and Validation (T&V):** The AI/ML models underpinning SATWATCH required extensive, iterative training and validation cycles to ensure their accuracy, robustness, and reliability for mission-critical applications. This involved leveraging decades of archived historical SSA observation data alongside meticulously crafted, **high-fidelity simulated datasets**. These simulations, potentially generated using tools like AGI's Systems Tool Kit (STK) or custom physics-based simulators, were crucial for representing a wide spectrum of nominal satellite operations, known historical maneuvers (for validation), potential component failure modes, rare orbital phenomena, and plausible hypothetical adversarial tactics for which real-world examples might be scarce or non-existent. The models underwent continuous, iterative refinement based on performance evaluation against withheld test data and, crucially, against **live, real-world events** as they occurred to ensure their accuracy, robustness against noise and data gaps, generalization capability to unseen objects or behaviors, and resilience to unforeseen operational scenarios. Strategies for handling **concept drift** (systematically retraining or adapting models as satellite behaviors, operational procedures, or the space environment itself change over time) and incorporating **human-in-the-loop feedback** from experienced USSF analysts (especially during initial validation phases to confirm or correct anomaly classifications) were essential components of the T&V process [18], [21]. Validation metrics focused not just on statistical accuracy but also on operational relevance (e.g., timeliness of warnings, reduction in false alarms).

- **Seamless System Integration and Workflow Adaptation:** SATWATCH was architected with modularity and standardized interfaces (such as REST APIs, potentially message queue interfaces like Kafka for asynchronous alerts) to facilitate seamless integration into the USSF's existing Command and Control (C2) frameworks, primarily the evolving **Space C2** system and potentially interfacing with the **Unified Data Library (UDL)** [4], [22]. The goal was to *augment*, not replace, existing systems and minimize disruption to established operator workflows. Standardized data exchange formats (like JSON, XML, potentially leveraging space-specific standards like CCSDS, and potentially STIX/TAXII for sharing threat-related intelligence) ensured interoperability with other analytical tools and C2 elements. The insights generated by SATWATCH—such as prioritized anomaly alerts with associated confidence scores, detailed conjunction event data including probability of collision (Pc) and contributing factors, and characterized maneuver assessments—were carefully designed to be clearly and intuitively visualized on operator consoles, often integrated into existing geospatial map displays (showing orbital tracks and predicted events) or analytical dashboards, thereby maximizing user adoption and operational utility. Significant effort was invested in designing **intelligent alert prioritization** algorithms, considering factors like event severity (e.g., predicted Pc), the criticality of the involved objects (e.g., national security asset vs. debris), the confidence level of the prediction, and potential mission impact, to actively mitigate operator alert fatigue [1].
- **Secure and Scalable Cloud Deployment:** The SATWATCH platform was deployed onto a certified, secure cloud infrastructure (such as DoD-approved instances within AWS GovCloud or Azure Government), strictly adhering to stringent Department of Defense cybersecurity standards and processes, likely following the **Risk Management Framework (RMF)** for authorization [23]. This cloud-native approach provided several essential operational benefits:
    - *Computational Elasticity:* The ability to dynamically scale computing and storage resources up or down based on fluctuating data volumes (e.g., during large constellation deployments) and processing demands (e.g., running intensive simulations for specific events), ensuring performance while optimizing cost.
    - *High Availability & Resilience:* Leveraging the inherent redundancy and fault tolerance capabilities of cloud infrastructure to ensure high system availability and mission assurance, critical for 24/7 SSA operations.
    - *Scalability:* The ability to rapidly scale the entire system horizontally (adding more processing nodes) as the number of tracked objects and ingested data sources inevitably grows over time, ensuring the system remains performant well into the future [12].
    - *Centralized Management & Updates:* Facilitating continuous monitoring, maintenance, security patching, and periodic retraining/updating of the AI models within a secure, controlled environment, potentially enabling CI/CD practices adapted for secure defense systems [24].

Results: Transforming Space Domain Awareness

The operational deployment and integration of the SATWATCH platform yielded substantial, measurable, and operationally significant improvements in the USSF's SSA capabilities, fundamentally transforming key aspects of space domain monitoring, assessment, and prediction:

- **Unprecedented Maneuver Detection Accuracy & Timeliness:** SATWATCH consistently achieved **95% accuracy** in detecting and correctly classifying specific categories of satellite maneuvers, encompassing both traditional impulsive (high-thrust, short duration) burns and, critically, the challenging **low-thrust** (continuous, gradual acceleration) maneuvers often used for electric propulsion systems. Detection often occurred within hours of maneuver initiation, a significant improvement over previous baseline methods which might take days or rely on manual analysis. This represented a marked improvement over previous methods, significantly reducing both false positives (spurious maneuver alarms that waste valuable analyst time and resources) and critical false negatives (missed maneuvers that could indicate impending threats, repositioning for advantage, or the initiation of proximity operations), allowing operators and analysts to focus their efforts far more effectively on events of genuine operational interest or concern [10].
- **Actionable Predictive Lead Time for Collision Avoidance:** The platform reliably provided operators with a crucial **72-hour predictive lead time** for accurately forecasting potential high-risk conjunctions, particularly those identified as resulting from detected ongoing or predicted future satellite maneuvers. This extended warning window allows sufficient time not only for detailed orbital safety analysis (including high-precision orbit determination updates and Pc calculations) and necessary international coordination (if involving foreign operators) but also for the efficient planning and execution of safe, optimized, fuel-efficient collision avoidance maneuvers by satellite operators, thereby preserving valuable satellite operational lifetimes and mission capabilities [11]. It also facilitates better allocation of scarce high-accuracy tracking resources (like dedicated radar or optical sensors) for verifying predicted high-risk events.
- **Enhanced Operational SSA Picture (Clarity & Depth):** SATWATCH delivered a richer, more comprehensive, dynamic, and significantly more timely understanding of the space environment—a true multi-dimensional operational picture, moving beyond simple object tracking. Operators gained deeper insights into not just *where* objects were located, but *how* they were behaving (maneuvering, station-keeping, tumbling), potentially *why* (inferred intent based on patterns, or inferred status based on signatures), and *what* they might do next based on predictive analytics. This enhanced understanding allowed for a more efficient and effective allocation of national SSA resources (sensors, analysts), enabling focused attention and sensor tasking on the objects, events, or orbital regions posing the greatest interest, uncertainty, or potential risk [3], [4].
- **Tangible Reduction in Collision Risk:** The combination of enhanced predictive accuracy for trajectories, longer warning lead times for potential conjunctions, and significantly reduced false alarm rates provided by SATWATCH's intelligent conjunction assessment

capabilities directly contributed to a measurable reduction in the overall probability of accidental collisions in space. This is particularly critical in increasingly crowded orbital regimes like Low Earth Orbit (LEO) and Geosynchronous Orbit (GEO), thereby safeguarding valuable national security assets, critical commercial communication and navigation satellites, and human spaceflight missions [5], [11], [12].

- **Improved and Accelerated Threat Assessment:** By rapidly identifying, characterizing (e.g., estimating maneuver delta-V, direction, duration, and likely purpose), and predicting the future path associated with anomalous or unexpected satellite behavior, SATWATCH enabled quicker, more confident, and fundamentally data-driven assessment of potential non-cooperative, irresponsible, or overtly hostile actions in the space domain [6]. This acceleration provides strategic decision-makers with the critical, timely intelligence needed to understand adversary capabilities, potentially infer intent from observed patterns of activity (e.g., coordinated maneuvers, RPO activities), and formulate appropriate, proportional, and timely responses, enhancing deterrence and strategic stability.

Conclusion: Securing the High Ground with AI-Powered SSA

The SATWATCH AI platform, expertly developed and implemented through the collaborative efforts led by 577 Industries, marks a pivotal and transformative advancement in Space Situational Awareness capabilities for the U.S. Space Force. It establishes itself as an indispensable analytical tool underpinning modern space domain operations. By masterfully integrating diverse, complex, multi-source data streams and applying a sophisticated, synergistic suite of AI/ML techniques specifically tailored for the unique physics and challenges of the space environment, SATWATCH delivers unparalleled capabilities for the timely detection, accurate characterization, reliable prediction, and contextual understanding of anomalous satellite activities and potential conjunction events [14], [15], [18].

Its successful implementation and seamless integration into USSF operational workflows have demonstrably enhanced the clarity, timeliness, predictive depth, and overall reliability of the Nation's space picture. This significantly improves spaceflight safety by proactively mitigating collision risks in increasingly congested orbits and, critically, arms commanders and strategic decision-makers with the actionable intelligence needed for effective threat assessment and response [4], [6]. SATWATCH thereby materially strengthens the United States' ability to protect its vital interests, deter aggression, preserve freedom of action, and ensure long-term stability in the increasingly vital and contested space domain. It positions the USSF for continued leadership in a rapidly evolving technological and geopolitical landscape. Future iterations of SATWATCH aim to incorporate even more diverse data sources (e.g., integrating space-based SSA sensor data, commercial RF monitoring constellations), enhance AI model adaptability to novel adversary tactics through continuous learning paradigms, further improve the explainability and trustworthiness of AI-driven insights for decision support, and potentially explore AI-driven optimization of sensor tasking for more efficient SSA resource allocation.

References

[1] United States Space Force, Space Capstone Publication: Spacepower, Doctrine Note, Jun. 2020.

[2] P. L. Hays, J. M. Smith, A. R. Wilson, and J. D. Douglas, Eds., Space Handbook: A War Fighter's Guide to Space, Vol. 1, Air University Press, Maxwell AFB, AL, 2009.

[3] B. Weeden and P. Samson, Eds., Global Space Situational Awareness Sensors, Secure World Foundation, Broomfield, CO, Mar. 2010. [Online]. Available: https://swfound.org/

[4] U.S. Government Accountability Office (GAO), Space Situational Awareness: DoD Needs to Improve How It Manages and Acquires Capabilities, GAO-20-316, Mar. 2020.

[5] National Research Council, Limiting Future Collision Risk to Spacecraft: An Assessment of NASA's Meteoroid and Orbital Debris Programs. Washington, DC: The National Academies Press, 2011.

[6] T. Harrison, K. Johnson, T. G. Roberts, and M. P. Gleason, Space Threat Assessment 2024, Center for Strategic and International Studies (CSIS), Washington, D.C., Apr. 2024.

[7] D. A. Vallado, Fundamentals of Astrodynamics and Applications, 4th ed. Hawthorne, CA: Microcosm Press, 2013.

[8] O. Montenbruck and E. Gill, Satellite Orbits: Models, Methods, and Applications. Berlin, Heidelberg: Springer-Verlag, 2000.

[9] A. K. Kembhavi and B. W. Weeden, "Big Data Challenges in Space Situational Awareness," in Proc. Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conf., Maui, HI, Sep. 2014.

[10] T. S. Kelso, "Analysis of the Low-Thrust Maneuver of USA 193 (NROL-21)," Celestrak, Feb. 2008. [Online]. Available: https://celestrak.org/ (Illustrative example of low-thrust analysis)

[11] F. Laporte and A. Lamy, "Space Debris Collision Probability Analysis," Acta Astronautica, vol. 66, no. 9–10, pp. 1427–1437, May–Jun. 2010.

[12] D. B. Spencer, Ed., Space Traffic Management. Washington, DC: AIAA Progress in Astronautics and Aeronautics Series, Vol. 263, 2019.

[13] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, Estimation with Applications to Tracking and Navigation: Theory, Algorithms, and Software. New York: John Wiley & Sons, 2001.

[14] J. L. Speyer and D. H. Jacobson, Primer on Optimal Control Theory. Philadelphia, PA: SIAM, 2010. (Note: Foundational reference relevant to state estimation/fusion).

[15] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," J. Comput. Phys., vol. 378, pp. 686–707, Feb. 2019.

[16] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, "Physics-informed machine learning," Nat. Rev. Phys., vol. 3, pp. 422–440, Jun. 2021.

[17] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. Cambridge, MA, USA: MIT Press, 2016.

[18] C. M. Bishop, Pattern Recognition and Machine Learning. New York, NY, USA: Springer, 2006.

[19] T. S. Kelso, "Validation of SGP4 and IS-GPS-200," Celestrak, Jul. 2007. [Online]. Available: https://celestrak.org/ (Illustrative example of orbit model validation)

[20] D. S. Alberts and R. E. Hayes, Power to the Edge: Command... Control... in the Information Age. Washington, DC: CCRP Publication Series, 2003.

[21] A. Holzinger, G. Langs, H. Denk, K. Zatloukal, and H. Müller, "Causability and Explainability of Artificial Intelligence in Medicine," WIREs Data Mining Knowl. Discov., vol. 9, no. 4, art. e1312, Jul./Aug. 2019.

[22] DoD, Risk Management Framework (RMF) for DoD Information Technology (IT), DoDI 8510.01, Mar. 2014.

[23] AWS, "AWS GovCloud (US) Regions," Amazon Web Services. [Online]. Available: https://aws.amazon.com/govcloud-us/ (Illustrative example of secure cloud infrastructure)

[24] M. Fowler, "Continuous Integration," martinfowler.com, May 2006. [Online]. Available: https://martinfowler.com/articles/continuousIntegration.html (Conceptual reference for CI/CD)