On Detecting Deception in Space Situational Awareness

James Pavur
Ivan Martinovic
[firstname].[lastname]@cs.ox.ac.uk
Oxford University, Department of Computer Science
Oxford, United Kingdom

ABSTRACT

Space Situational Awareness (SSA) data is critical to the safe piloting of satellites through an ever-growing field of orbital debris. However, measurement complexity means that most satellite operators cannot independently acquire SSA data and must rely on a handful of centralized repositories operated by major space powers. As interstate competition in orbit increases, so does the threat of attacks abusing these information-sharing relationships. This paper offers one of the first considerations of defense techniques against SSA deceptions. Building on historical precedent and real-world SSA data, we simulate an attack whereby an SSA operator seeks to disguise spy satellites as pieces of debris. We further develop and evaluate a machine-learning based anomaly detection tool which allows defenders to detect 90-98% of deception attempts with little to no in-house astrometry hardware.

Beyond the direct contribution of this system, the paper takes a unique interdisciplinary approach, drawing connections between cyber-security, astrophysics, and international security studies. It presents the general case that systems security methods can tackle many novel and complex problems in an historically neglected domain and provides methods and techniques for doing so.

CCS CONCEPTS

• **Security and privacy → Systems security**;

KEYWORDS

satellite, space security, ssa, space situational awareness

ACM Reference Format:
James Pavur and Ivan Martinovic. 2021. On Detecting Deception in Space Situational Awareness. In *Proceedings of the 2021 ACM Asia Conference on Computer and Communications Security (ASIA CCS '21), June 7–11, 2021, Virtual Event, Hong Kong*. ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3433210.3453081

1 INTRODUCTION

Space is hard. Satellites operate under constant threat from their environment, besieged by extreme thermal fluctuations and intense radiation. As orbit grows more crowded, a human-made threat has become increasingly salient: space debris. Today, more than 21,000 debris objects measuring >10 cm in diameter whiz overhead in excess of 20,000 km/h [34]. These are joined by an estimated 500,000 1-10 cm diameter particles. Colliding with any one of these objects can debilitate or even destroy a satellite.

To combat this threat, satellite operators rely on a class of data known as Space Situational Awareness (SSA). SSA describes the position, nature, and movement of space objects. Physical astrodynamic models can approximate the movement of debris objects in the short-term, but dynamic factors like solar storms, micro-debris collisions, and drag frustrate this process. As a result, SSA requires continuous updates in the form of observational measurements. Identifying and tracking minuscule objects moving at bullet-like velocities thousands of kilometers away is inordinately difficult, even with sophisticated astrometry equipment.

In this paper, we contend that the complexity of SSA acquisition gives rise to an interesting, but largely unstudied security dynamic. Space surveillance catalogs are beyond the means of most nation-state actors, much less commercial satellite operators. As a result, the vast majority of industry participants rely on shared data from one of two global SSA authorities.

We begin by characterizing a novel threat model, showing that SSA authorities have incentives to mislead data recipients and, indeed, have empirically attempted to do so. Concerns around SSA dependency have recently become a subject of increasing geopolitical tension. However, political assertions regarding the capabilities of attackers and needs of defenders lack technical foundations.

By re-framing SSA deceptions as information integrity attacks, we draw a novel connection between two intuitively distant topics: systems security research on anomaly detection and astrophysical research on resident space object (RSO) characterization. This gives rise to a fundamental question: Using nothing more than an untrusted third-party's description of RSO motion, can an SSA recipient ascertain its nature?

1.1 Contributions

This paper makes several contributions. First, we outline real-world empirical evidence for what is, to our knowledge, a previously unstudied class of data deception attack used by major space powers (Section 3.2). Specifically, we show that states have attempted to deceive others into believing that classified space platforms are pieces of debris. We further show that the mechanism for actualizing these deceptions is strategically limited by a trade-off between privacy and sustainability in the SSA domain (Section 3).

This gives rise to direct technical contributions, starting with a method for replicating these attacks using real-world SSA data

ASIA CCS '21, June 7–11, 2021, Virtual Event, Hong Kong

ACM ISBN 978-1-4503-8287-8/21/06…$15.00
https://doi.org/10.1145/3433210.3453081

(Section 4). These simulations are used to develop what we believe is the first anomaly detection system for space surveillance (Section 6).

Through the application of classical machine learning techniques, we arrive at a surprising conclusion: even the most fundamental information about an RSO - less than 140 characters of public data regarding its position and motion - contains physical signatures of an object's nature. These signatures can supplement expensive astrometry equipment and uncover concealed satellites (Section 6.3).

Beyond these direct outputs, the paper makes several methodological contributions. Space operations security research has been historically quite limited, in large part due to a lack of methodological foundations and cross-domain perspectives. This paper provides ample background information, assuming an audience of security experts interested in but unfamiliar with the space domain. The intention is to provide foundations for future work at the intersection of space and cyber-space.

2 BACKGROUND AND RELATED WORK

In this section, we discuss the operation of SSA systems in practice and how the unique technical challenges for space surveillance give rise to asymmetric data-sharing dependencies. The section concludes with a discussion of limited related work on SSA integrity.

Today, more than 2,000 satellites orbit the Earth, and this number is expected to increase by an order of magnitude over the next decade [53]. Each space mission runs the risk of adding pieces of space debris to orbit. Debris can be generated as a byproduct of launch operations, weapons tests, or accidental collisions. Space debris objects threaten satellites through high-velocity collisions. Each collision can potentially cause a satellite to break up or disintegrate into thousands of additional debris objects, creating a "cascade effect" with consequences for the broader space environment [46].

The trajectory of debris objects, and even satellites, is not fully predictable. Chaotic physical forces, such as solar weather and drag, will cause deviations from projected paths over time. Space surveillance is partially about detecting these changes as they occur. From the moment a satellite launches to the moment it retires, accurate and recent SSA data is critical to mission safety.

Beyond this immediate operational utility, SSA is also of significant military and strategic value. The world's largest militaries are deeply dependent on space assets for Command, Control, Communications, Computers, Intelligence, Surveillance and Reconnaissance (C4ISR) support. This heavy reliance means space powers are highly sensitive to orbital threats. Accurate SSA allows them to monitor the behaviors of peers in orbit, detect espionage and weapons systems, and enforce diplomatic norms.

2.1 SSA in Practice

SSA measurements can be either reported or observational.

Reported measurements originate from sensors aboard satellites. The type of sensor varies depending on the object's orbit (see Figure 1). Satellites in Low Earth Orbit (LEO), approximately 2,000 km above the Earth's surface, receive signals from Global Position System (GPS) satellites in Medium Earth Orbit (MEO), approximately 20,000 km above the Earth's surface. Thus LEO orbit determination is normally derived from on-board GPS data which is relayed to the

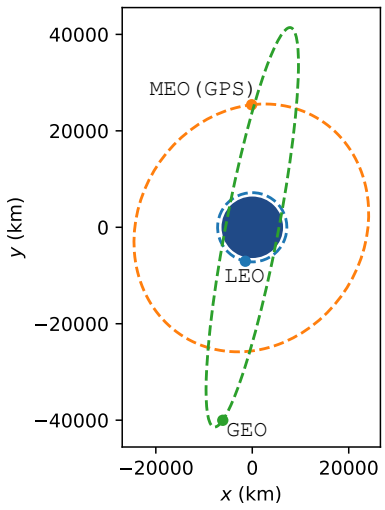

Figure 1: Depiction of Common Earth Orbits. Note that GEO satellites are beyond the range of GPS signals.

Earth as telemetry messages [49]. Higher orbits, such as Geostationary Orbit (GEO), are above GNSS systems and cannot use them for positioning. Instead they employ alternative metrics such as tracking relay satellites, star-trackers, or time-difference of arrival (TDOA) calculations at multiple earth-based antennas [6, 24, 55].

Observational measurements are required to track space debris and defunct satellites which cannot transmit telemetry directly. These measurements require sophisticated astrometry platforms. For objects in LEO, radar sensors are typically used, while, at greater distances, electro-optical telescopes are necessary [25]. A single ground station cannot reliably track objects. Instead, many observations correlated from from sites distributed across the Earth.

The cost of such a system is immense. While precise numbers are scarce, the latest round of capacity upgrades for the US Space Surveillance Network is believed to have exceeded $6 billion in procurement costs [15]. This puts in-house SSA beyond the means of not just commercial actors, but the majority of states. Even if states have sufficient resources and will (such as in the case of China or the EU), they may lack the territorial reach and diplomatic leverage to deploy radar and telescope systems across the planet.

2.1.1 SSA Repositories. The result of these cost barriers is that the majority of satellite operators rely on third-party SSA data.

By far, the dominant source of SSA is the United States Space Surveillance Network (SSN). The SSN comprises more than 20 locations, leveraging the US military's extensive network of forward deployed military installations to achieve geographic distribution. It is believed to be the only system capable of tracking smaller objects measuring 5-10 cm in LEO and 1 m in GSO [28, 45].

The closest competitor is the Russian Space Surveillance System (RSSS). Its nature and capabilities are opaque, but it is believed to consist of at least 8 sensing sites (primarily within former USSR territory) and to have a catalog about 1/3rd the size of the US SSN [45]. A nominally civilian network - the International Scientific Optical

Network (ISON) - is managed by the Russian Academy of Sciences and includes some smaller academic research installations [21].

A handful of smaller networks exist, including the European Space Agency's (ESA) Space Surveillance and Tracking program, and systems operated by China, Canada, India, Japan, Korea, France, and Ukraine [45, 54]. Recently, commercial services promising independent civilian SSA have emerged, although, at present, none offer data comparable to the SSN [10, 30, 36].

2.1.2 SSA Sharing. Most operators receive SSA data from the US SSN. The US military publicly posts SSA through Space-Track.org in the Two-Line Element Set (TLE) format.

This data standard was developed in the 1970's to facilitate the sharing of an object's *ephemeris* (its projected orbital path and position) using two 80-column punch cards [20]. TLEs were designed in conjunction with orbital propagation models to forecast the motion of an object through orbit. In particular, TLEs are designed for use with the Simplified General Peturbations Model (SGP4) [20].

The use of archaic data standards is ostensibly to maintain backwards compatibility. However, it offers some practical strategic benefit. As the precision of TLEs is limited, information can be disseminated more freely that if the data were more granular. For example, TLEs are sufficiently precise for satellite communication purposes and even for detecting many collision threats, but they are not precise enough for ASAT weapons targeting.

Where more precise data is required, the SSN takes a case-by-case sharing approach. This requires a formal agreement which imposes additional requirements on recipients, such as that they provide the SSN with reported telemetry measurements from their own satellites [48]. This bolsters the US military's SSA advantage and allows them to control the flow of potentially dangerous information to untrusted states while also sharing "no questions asked" public ephemerides to prevent environmental catastrophes.

The precise accuracy of TLEs is difficult to generalize and varies depending on object location and orbit. As a rough approximation, at *epoch*, the TLE's time of issue, it is accurate to within 1 km in any dimension, this degrades at a rate of 1-3 km per day for the first week [3]. Beyond this point, SGP4-propagated TLEs become increasingly meaningless, eventually describing impossible orbits.

The TLE format itself, along with notation used to reference orbital elements, is summarized in Figure 2. An overview of those elements most directly related to an orbit's physical properties can be found in Figure 3. It is important to remember these values are intended as inputs to SGP4 and are not direct physical attributes.

As TLEs go out of date, Space-Track provides repository updates. While the SSN claims to screen all objects daily, daily updates are not typically provided. This is because the previous TLE are often deemed sufficiently accurate. Space-Track also notes in its user agreement that TLE updates may be withheld for "national security reasons" or as a result of sensor anomalies. Several classified objects are also never listed in public SSN catalogs.

2.2 Related Work

There has been only limited technical systems security research on satellite missions, much less the specific context of SSA. Excluding research on terrestrial GPS equipment, a cursory search found only three satellite-security papers published in any of the "big four" systems security venues (IEEE S&P, Usenix Security, NDSS, and ACM CCS) in the past 30 years. All three stuck to the relatively comfortable territory of communications privacy [2, 19, 38].

This lacuna extends far beyond the systems security niche [11]. Meta-analysis of this gap has highlighted various causes, ranging from limited appetite for collaboration between the aerospace and security communities to the restricted and expensive nature of satellite hardware [11]. A key motivation for this paper is to demonstrate that fundamental methods used in systems security can generate relevant knowledge for aerospace operators.

The most directly related prior work is our own publication discussing cyber-mediated counterspace capabilities and the use of SSA tampering attacks to create misleading space debris collision forecasts [40]. Due to heavy classification barriers and restricted information availability, that research necessarily rested on a large number of assumptions regarding the debris forecasting and conjunction avoidance process used by national militaries. This paper arose, in part, from our recognition of the need for a deeper technical look at a simpler, but empirically validated, threat model to supplement our prior work.

Beyond that paper, we found one additional related publication which considers SSA from a privacy perspective [8]. In it, the authors consider a satellite operator who wishes to share reported SSA measurements with a central authority (e.g. the SSN), but to limit the accuracy of this data for privacy reasons. They propose a mechanism for the injection of artificial measurement noise for a controlled trade-off between accuracy and privacy. This is not directly relevant to our threat models, which focus on end-product TLE ephemerides rather than astrometry measurements, but it does support our argument that safety and security can trade-off in SSA information sharing contexts (see Section 3).

Peripheral work can be found in the blockchain community, where businesses have emerged leveraging political distrust in US-military data as an impetus for a decentralized SSN alternative [12]. This aligns with academic work which contends that SSA costs and responsibility paradigms may be well-suited to blockchain [50].

Ultimately, SSA security remains a largely unstudied topic. In particular, to our knowledge, no research exists which considers how SSA recipients can protect themselves against deceptions made by centralized SSA authorities.

3 THREAT MODELING

SSA repository owners have many environmental incentives to disseminate data freely. This is due to the aforementioned "cascade effect" risk, whereby debris from a satellite collision can threaten the SSA operator's own satellites.

However, strong incentives also exist to withhold SSA data. Accurate SSA regarding military and intelligence satellites can enable adversaries to target these systems via ASAT weapons or "shadow" them with sophisticated eavesdropping platforms [16]. This is more than an abstract theoretical risk. At least four major space powers (United States, Russia, China, and India) have demonstrated ASAT missile capabilities and three (United States, Russia, and China), have demonstrated dual-use proximity operations technology. As kinetic counterspace capabilities increase, SSA operators may feel

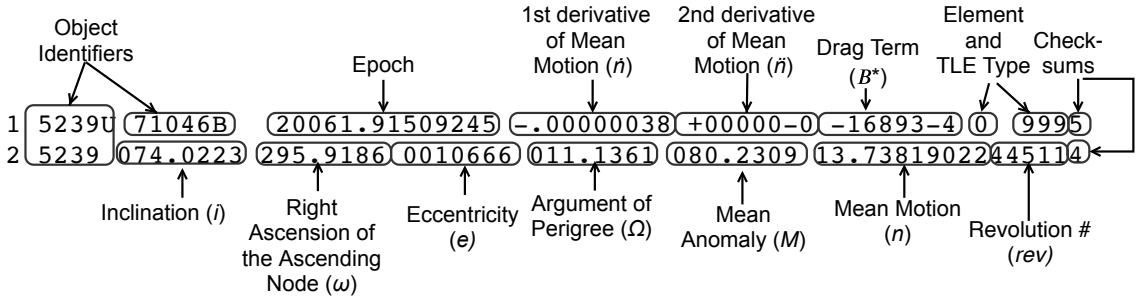

Figure 2: The TLE format. The notation in parenthesis is used throughout this paper and summarized further in Figure 3.

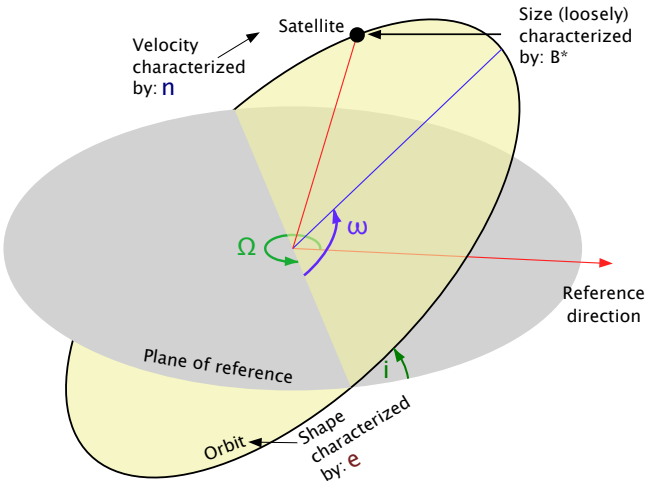

Figure 3: A summary of orbital elements. Adapted from [29].

pressure to conceal the precise location of key satellites, or even their existence altogether.

This creates a dilemma. On the one hand, environmental concerns support the sharing of accurate and complete SSA; a classified satellite is no less at risk of collisions than an unclassified one. On the other hand, this same data can be abused by adversaries for counter-space or counter-intelligence operations. Moreover, if an SSA operator is caught withholding or modifying shared data, this choice can cause unwanted attention or reputational harm.

At present, the ephemerides of tens of thousands of objects are physically unknowable to anyone other than the US military. Foreign governments, including major space powers like China and Russia, face an unpalatable reality of total reliance on the goodwill of a foreign military for the safety of their own space missions. Historically, SSA sharing has been a well-regarded function of space diplomacy and international SSA partnerships have proven resilient to broader geopolitical tensions. In recent years, however, both China and Russia have grown increasingly skeptical of the US military's benevolence in the SSA domain [17, 33]. These states, and others, have begun investing in expanded domestic capabilities. However, the inordinate expense and complexity of such efforts means that, for the foreseeable future, SSA-sharing will remain common practice by necessity if not by choice.

3.1 Classification Deception Attacks

In this paper, we focus on one specific abuse of SSA trust. In it, the "attacker" is an SSA repository owner who wishes to deceive a third party SSA recipient (the "defender") into believing that a spy satellite is a piece of space debris. In doing so, the attacker seeks to bypass the aforementioned safety/privacy tradeoff: sharing the location of the object accurately to avoid collisions, but also preventing adversaries from learning the true nature of a strategically important space asset.

This has several benefits for the attacker when compared with merely omitting the secret satellite from the SSA repository altogether. By sharing accurate data indicating that there is *some* object in that orbit, an attacker can issue conjunction alerts to protect that object without arousing suspicion. Moreover, while other SSA operators may be able to determine the presence of an object in that orbit, determining the object's purpose is a much more complex task requiring significant manual effort and sophisticated equipment. Against a backdrop of tens of thousands of other debris objects, defenders must prioritize and are more likely to allocate limited resources to studying objects explicitly designated as satellites.

For an attacker, success would allow them to prevent defenders from altering their behaviors during overpasses - allowing for the collection of imagery intelligence (IMINT) data. Alternatively, it allows the attacker to protect key communications missions from ASAT weapons, or to block competitor access to a high-value orbit without negotiating appropriate rights.

For a defender, the primary objective is to identify which objects in an SSA repository are maliciously misdesignated. Under our threat model, we assume the defender is a nation-state actor with limited SSA capabilities relative to the attackers. For example, states like China and Japan have only small SSA networks. By flagging the most suspicious objects, they could optimize limited resources.

As a point of clarification, we refer to situations where SSA deviates from physical reality as "attacks" and those causing these deviations as "attackers." This is not intended as a value judgment or implication of malfeasance. There are legitimate intelligence and security motivations for obfuscating SSA data. While this paper focuses on defensive techniques, its findings are also relevant to agents designing satellite missions which evade detection.

3.2 Empirical Support

A key barrier to technical security research on space operations is the dearth of unclassified empirical data. Our decision to focus on

object-type deceptions, rather than more convoluted threat models targeting object location and motion characteristics, arises in part from the fact that highly similar attacks have happened in practice.

The core conceit - using space debris as a disguise for critical space systems - is nearly as old as spaceflight itself. Throughout the Cold War, both the US and USSR dedicated significant resources to the interception and monitoring of that included disguised hardware used for transferring mission data to Earth [9].

Much more recently, Russia was credibly accused of deceiving the international community by claiming a piece of orbital debris has been generated from a Rokot-Briz launch on May 9th, 2014 when, in fact, the object in question was a nano-satellite [51]. At least in unclassified contexts, this deception worked, with the US military tracking object *2014-28E* as a piece of space debris for more than six months. In November 2014, the object began to engage in orbit-altering maneuvers, revealing its true nature.

In 2015, Oleg Maidanovich, head of Russian space command, reported that they had identified a cluster of espionage satellites masquerading as space debris [52]. Maidanovich declined to provide further detail, but Russian press coverage heavily implied US involvement. Regardless, this demonstrates that Russian military officials perceive this threat model as a plausible attack vector.

Unofficial statements from individuals within the defense community further support an abstract interest in developing nanosatellites which are small enough to be only trackable by the SSN [37]. While comments from the US National Reconnaissance Office (NRO) on their use of nanosatellites are characteristically vague, intelligence experts have argued that the ability to evade detection by foreign SSA networks is one driving motivation [37].

Finally, in 2020, the US Department of Defense and Department of Commerce began negotiations on a series of regulations restricting the imaging of space objects. Of particular relevance are the requirements that space-to-space photographs of satellites and debris have a maximum resolution of 50 cm and a blanket ban on photographing any objects not listed in the Space-Track.org repository [18]. Many nano-satellite platforms are smaller than 50 cm in any dimension. For example, CubeSat form-factors start at 10 cm^3.

In short, recent developments point to a growing interest from both of the world's leading SSA authorities in concealing operational satellites as pieces of space debris. For states which rely on either, detecting SSA deceptions will become increasingly relevant.

4 EXPERIMENTAL DESIGN

Given the lack of public specifics on historical SSA attacks, defense evaluation hinges on the development of realistic simulations. We seek to minimize the assumptions required by focusing on a foundational threat model: lying about a space object's purpose.

A traditional conceptualization of this challenge would revolve around improving physical sensors and measurement techniques. However, we take the unconventional approach of treating it as an attack on digital information integrity. This lets us leverage existing work on anomaly detection to both build and evaluate our system. Treating this as an integrity attack also helps incorporate real-world data. Rather emulating the launch of spy satellites through imprecise physical simulations, we can instead "tamper" with the contents of real-world public ephemerides.

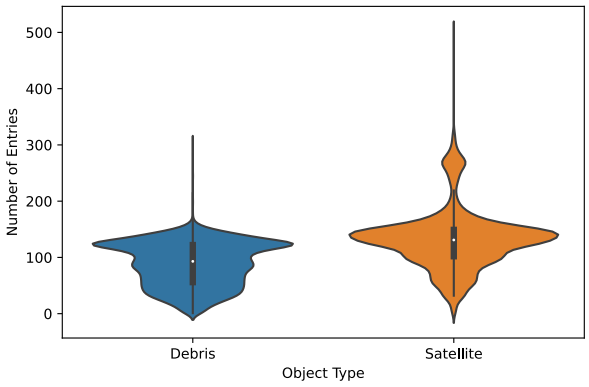

Figure 4: The distribution of SSA entries by object type. The width of each plot is proportional to the number of objects which have been updated at a given annual rate.

That said, there is at least one critical assumption required. We must assume that the real-world ephemeris data used is reasonably trustworthy. The physical reality of SSA collection means that, to have any real-world data, we must start with the assumption that the US military is not already trying to deceive us at scale. The reasonableness of this depends on personal political perspectives.

However, we need not have perfect faith in US SSN data. Even if this "ground truth" information includes a small number of maliciously concealed satellites, they are highly unlikely to represent a meaningful proportion of the overall dataset due to the reputational and environmental dynamics discussed in Section 3. Given that the ability for defenders to determine *anything* about the nature of an object from its public ephemeris alone would represent a substantial improvement, a theoretically perfect system is far from necessary. Nevertheless, it is worth remembering that, when we discuss our model's accuracy at recognizing satellites/debris objects, this is, in fact, shorthand for its accuracy at predicting whether or not the US military would publicly label an object as such.

4.1 Data Description

Our source of SSA data is the United State's SSN catalog. We consider one year (March 1, 2019 to March 1, 2020) of public SSA. This amounts to approximately 2 million TLEs describing around 19 thousand objects. Roughly 40% are satellites and 60% pieces of orbital debris. However, satellite ephemerides tend to be updated at slightly greater frequencies, meaning that the dataset is roughly balanced with respect to individual TLE classes (see Figure 4).

A descriptive summary of individual orbital elements appears in Table 1. We have excluded the second time derivative of mean motion ($\ddot{n}$) from analysis as the vast majority (>99%) of records report this value as zero. Neither $\ddot{n}$ nor $\dot{n}$ are used by SGP4 but are legacy holdovers from older propagators [26].

Note that several values are distributed across the full range of the TLE format. This holds for many angular elements (Ω, ω, M, and, to a lesser extent, i). This makes sense as orbits are dictated by mission requirements and variation in angular elements encapsulates their spread around the Earth. For other features, the

Table 1: Distribution of Orbital Element Values

Element	Mean	Std. Dev.	Min	Median	Max
$\dot{n}$	0.00004141	0.00104152	-0.11202822	0.00000097	0.99999999
B^*	0.0003716	0.0231546	-0.9934500	0.0000486	20.7840000
i	72.8193	31.5599	0.0000	82.7970	144.6450
Ω	177.8059	108.2238	0.0000	175.6950	359.9998
e	0.0563491	0.1644054	0.0000002	0.0039360	0.9184180
ω	175.9580	104.4194	0.0006	172.2482	359.9998
M	183.0258	107.9505	0.0004	186.4745	360.0000
n	11.773592	4.846328	0.037454	14.026936	16.474780
rev	37,138	28,991	0	31,266	99,999

Note: The number of significant figures is representative of the precision of the TLE data format. The only exception is B^*, which has variable precision.

physical requirements of maintaining a stable orbit manifest clearer boundaries on the parameter space. For example, n, roughly a proxy for satellite velocity and altitude, has an upper-bound of around 17 revolutions per day. Although the TLE format could allow for the representations of satellites with greater velocities (up to 200,000 km/h in LEO) such claims are physically implausible. For this paper, we assume that the data is presently trustworthy. That is to say, the US military has not, over the one-year period concerned, engaged in SSA deceptions. The reasonableness of this assumption depends largely on personal political perspectives. However, even if the US military does presently engage in SSA classification attacks, they are unlikely to do so for a statistically meaningful proportion of the SSN repository. Thus, while treating SSN data as "ground truth" might lead to minor accuracy discrepancies, it is unlikely to cause radical performance deviations.

A second source of SSA, such as Russia or China's, would have been desirable. Unfortunately, such data is not generally available. Commercial SSA is similarly difficult to access, with models revolving around direct contracts with national militaries or agreements with defense-industrial companies [23]. To our knowledge, the US military is the only entity which publicly shares a robust SSA catalog, although some organizations will repackage and redistribute this data with additional processing.

Fortunately, this SSA is representative of real-world practice. It is likely that nearly all space-faring nations at least consider the SSN as a factor in their SSA processes. This is particularly true for smaller objects, such as 10 cm CubeSats and small debris particles, which are believed to be only trackable using SSN's technology. The main variation between defenders is thus less in terms of whether they must trust US data and more on the extent of that dependence.

4.2 Attack Implementation

In the SSN, debris objects are designated by a naming convention whereby the suffix "DEB" is appended to the name of their originating mission. In the simplest sense, we can replicate an object deception attack by appending the "DEB" label to an satellite's descriptor. It is important, however, to do so for all TLEs belonging to an object from the moment of launch onwards, or defense models may receive prior information which biases performance. Additionally, TLE parameters such as the object's ID and name are trivially alterable by an attacker. We thus assume that the attacker optimally alters all such metadata elements.

As mentioned in Section 3, the attacker has strong incentives *not* to alter the orbital elements themselves. This is because accurate positional information is necessary to issue conjunction alerts which protect the concealed satellite from orbital collisions.

5 SIGNATURE EXTRACTION APPROACHES

Differentiating between debris and satellites on the basis of their motion alone is a complex problem largely neglected in prior work. In the status quo, RSO classification is handled by SSA operators who leverage billions of dollars in labor and equipment to perform this task. Even with high-end astrometry equipment, both satellites and debris objects appear as little more than tiny dots, making determining their nature a challenge. This may explain why the US SSN failed to identify *2014-28E* following its launch. When observation is required, the process can entail many measurements in order to build composite signatures from physical attributes, such as light reflections [14].

Fortunately, such effort is only rarely required. SSA operators typically can rely on reported data from satellite owners and launch operators. Following a launch, they collaborate with owners to hunt down and make radio contact with platforms. Any objects not identified in this process can be considered debris.

5.1 On Physical Signatures

An initial, but misleading, intuition might be to use maneuvers, such as the one which ultimately revealed *2014-28E* (see Section 3.2), as a signature to identify satellites. However, many small satellites never maneuver. Nano-satellites reach orbit by "hitchhiking" through vehicle-sharing agreements. At some point in the launch sequence, they are pushed away from the rocket - normally using a simple spring-powered mechanical plate [39]. From this point, they spend the entirety of their lifespan adrift. Even when larger satellites *do* have maneuvering systems, they use them only sparingly. Moreover, space debris objects themselves can appear to "maneuver" as the result of space weather effects.

Still, there is intuitive reason to expect differentiation between debris and satellite orbits. While space missions are diverse and their orbits vary, all satellite orbits are fundamentally *selected* by sentient individuals while debris orbits are *incidental* to this selection process. The selection criteria for space missions is also generally guided by a few common factors. For example, communications constellations will attempt to maximize the time at which certain points on the Earth's surface have line-of-sight to one or more satellites. Imagery and remote sensing systems will prefer orbits which offer sensor coverage of specific regions.

Even more fundamentally, cost and convenience influences the selection of orbits. Some missions seek to maximize longevity and minimize drag. Others may be required by regulators to de-orbit after some time and may prefer orbits which naturally cause this. Finally, small satellite payloads are often forced into orbits based on the requirements of primary payload they "ride-share" with.

Debris faces none of these constraints. While debris is made up of materials that, at one point, belong to a "designed" mission, its position is the result of chaotic physical dynamics.

It is tempting to imagine a hypothetical checklist of physical characteristics which separate the signatures of "designed" and

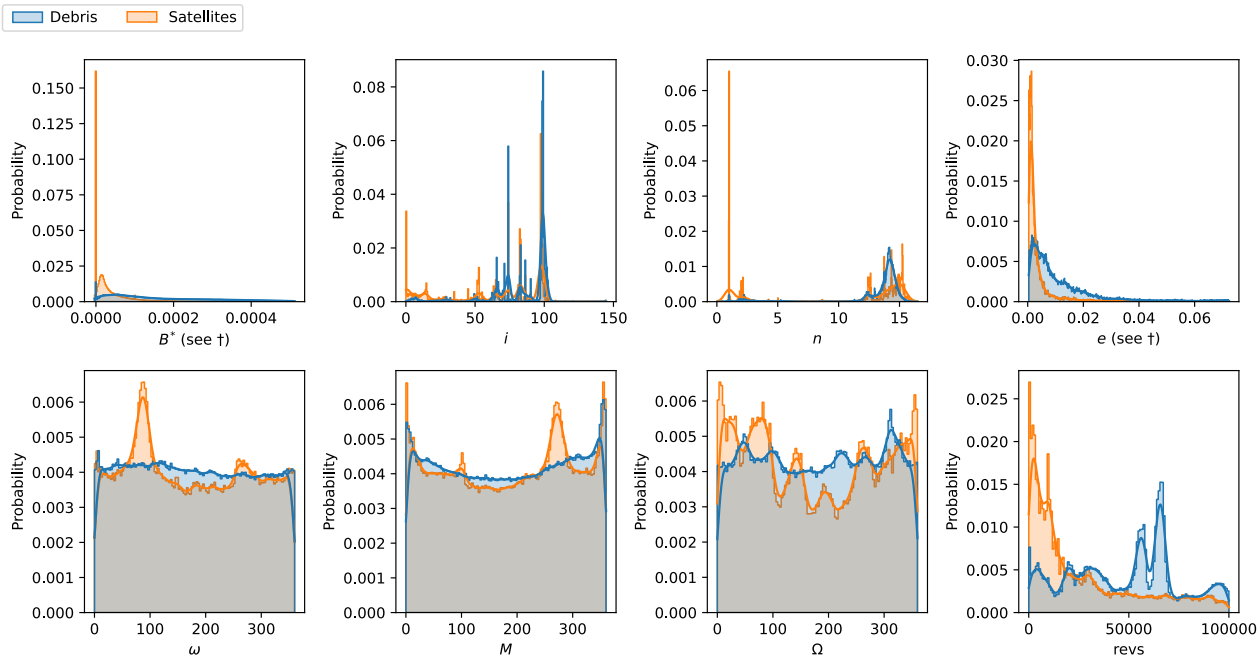

Figure 5: Distribution of key orbital elements by object type. † The B^* and e plots are "zoomed in" around the 10th-90th percentile values to improve legibility due to heavy clustering in both elements.

"incidental" orbits. However, the diversity of both spacecraft missions and debris objects increases the complexity of such a task substantially. State-of-the-art methods for space object classification and shape determination thus continue to rely on photometric and other observed characteristics [22, 35].

In our threat model, the fundamental challenge for defenders is their lack of access to this direct physical knowledge. Instead, they must leverage a restrictive set of public TLE parameters, all of which exhibit substantial overlaps across the population of objects (see Figure 5). While we can observe certain distribution biases in favor of a particular class, such as the many near-zero n values belonging to GEO communications satellites, these distributions hint at the inherently multidimensional nature of our defender's problem. It makes sense that such a task is non-trivial as it is unlikely that states would continue to invest billions of dollars in laser/optical/radar RSO classification systems if the matter could be trivially resolved by, for example, considering the altitude of an orbit or the separation of a given object from others around it.

One feature in Figure 5 appears promising: *revs*. This is misleading. As *revs* is simply the count of revolutions an object has made since it first appeared in the database, the feature lacks explanatory power for newer objects. The clustering observed is thus reflective of increased satellite launches in recent years coupled with a handful of major debris generating collisions (e.g. the 2007 Chinese ASAT demonstration). While a defender might use *revs* as a components of their classification system, for example to exclude objects from these events, it alone is insufficient for differentiation.

5.2 Why Machine Learning?

An alternative to manually devising a complex rules-based system for orbit differentiation would be the use of machine learning classifiers. The intuition here is that classical machine learning methods may allow us to effectively grapple with the astrodynamic complexity and multidimensionality of RSO classification. The main focus of research to date has been on improving the accuracy of orbital forecasting or improving the RSO classification value of sensor data [1, 22, 31, 56]. To our knowledge, no attempt has been made to tackle the RSO task with a feature set as limited as the one available to defenders in our threat model.

Nevertheless, this research offers valuable insights for our own. For example, Furfaro et al. found that the light curves generated in simulated telescope images were differentiable using a convolutional neural network (CNN) [13]. Using photographs from one Russian SSA telescope, they managed to classify three object classes (Rockets, Satellites, Debris) with roughly 77-85% accuracy. Similarly, Liu et al. proposed an ontology-based classification model which combines a holistic rules-based approach and classical machine learning [32]. This model is used to identify space mission purposes (e.g. commercial vs remote sensing). Again, they rely on supplemental data not available in our threat model, such as telescope brightness, radar cross section sizes, power supply, object mass, and object ownership records. They conclude with the finding that their approach matches the performance of a random forest (80-90% accuracy), but with reduced training times.

These studies suggest that machine learning can solve RSO classification tasks that have proven too costly or difficult using manually-tailored physical models. However, prior work assumes access to data which our defender is unable to trust. One area where solutions

to this issue may be found is the systems-security domain. There, researchers have leveraged machine learning techniques to detect implausible relationships between elements, even in untrusted data.

For example, decision trees have been used to identify fishing vessels pretending to be vessels of different types on the basis of their location and motion information as reported in Automatic Identification System (AIS) messages [27]. The underlying concept here - detecting incongruities between tampered object labels and position data concerning those objects - is quite similar to our own task. Likewise, decision trees have been used to detect covert channel communications disguised as multimedia applications and various network identifier spoofing threats [4, 7].

In short, existing systems security research makes the case that machine learning approaches can help draw out complex inter-relationships between immutable physical characteristics. These relationships can be used to assess the plausibility of related, but more mutable, claims in the presence of an attacker.

6 DEFENSE IMPLEMENTATION AND EVALUATION

In designing our defensive system, it is important to highlight that machine learning is not, itself, the focus of our research. Rather it is a tool for evaluating a more strategically relevant hypothesis: that even basic positional data about a space object contains signatures which may expose concealed satellites. That said, one of our stated objectives is to facilitate future work in the domain. To this end, we start by briefly highlighting feature-engineering challenges and pitfalls that arise from the unique nature of SSA data.

6.1 Feature Engineering Pitfalls

Initially, one might treat this as a straightforward binary classification problem with conveniently pre-labeled data. Running a trivial k-nearest neighbors classifier trained on randomly sampled TLEs detects satellites astonishingly well (*f1-score*: 0.92, *precision*: 0.90, *recall*: 0.94). However, these results are deceptive. This is because there are spatio-temporal interrelationships between TLE entries which "leak" information.

For example, a naive classifier can over-fit if training data contains prior (or future) TLEs belonging to an object in the test data. If we restrict the evaluation data only to *objects* not included in the test data, performance plummets (Figure 6). Moreover, information about related objects, such as the future location of other satellites from the same launch, can be used to forsee clustering patterns unavailable to a defender who lacks time-traveling capabilities. After restricting training sets to TLEs before a certain epoch, we find our initially promising classifier now performs only superficially better than random guessing (*f1-score*: 0.58, *precision*: 0.63, *recall*: 0.55).

6.2 Baseline Performance

In this paper, we focus on decision-tree classifiers as they have proven viable in prior work on both RSO and related anomaly detection tasks (see Section 5.2). While there are a plethora of machine-learning approaches that could be employed, applying them in this paper serves little purpose other than academic diversion. If decision trees can detect object-type signatures in TLE-data

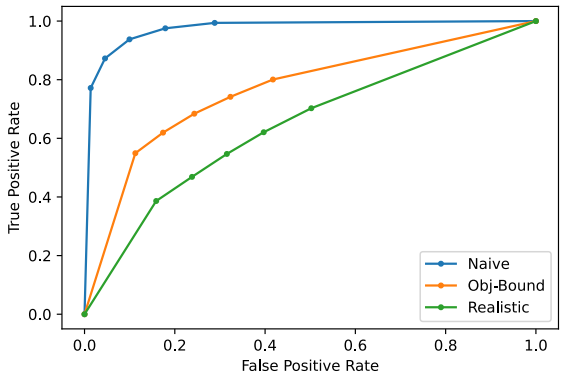

Figure 6: Demonstrative K-nearest neighbors performance comparisons. Note how improper segregation of the training data dramatically overestimates model performance.

that is sufficient to prove the existence of such signatures, even if another technique (e.g. neural networks) could also do so.

On the basis of these intuitions, we tested four popular decision-tree based classifiers against our attack. The first is a basic CART (classification and regression tree) decision tree model which leverages Gini-impurity as a metric for identifying optimal splits [42]. The second was a bagged meta-estimator which applies this decision tree model repeatedly to random subsets of both features and entries within the training data to develop an aggregate classifier (a method sometimes referred to as "random patches") [41]. The third is a basic random forest which also leverages Gini-impurity as a splitting metric [44]. The final is a histogram-based gradient boosting model, which is a relatively complex ensemble decision tree that leverages gradients to correct for errors in individual tree models [43]. Unless otherwise noted, no special effort was made to hypertune these algorithms beyond the reasonable defaults included in the referenced Scikit-learn implementations. While the optimal design of decision-tree classifiers for RSO classification represents a fruitful topic for future aerospace research, it is far beyond the remit of this preliminary security-focused work.

The data used in training these models is described in Section 4.1 and is publicly available on Space-Track.org. The elements n, B^*, Ω, ω, $revs$, $\dot{n}$, i, e and M are extracted from the TLEs with the same precision as the underlying TLE format (see Table 1) and used to predict one of two labels: *Debris* or *Satellite*. TLEs are appropriately segregated on the basis of epoch and object identifier in order to avoid the pitfalls detailed in Section 6.1.

As mentioned in Section 4.2, the scenario we describe as an *attack* occurs when the *Debris* label is falsely applied to a *Satellite* object by the attacker. We further assume the attacker is capable of optimally modifying most TLE metadata (such as the object's ID and name) but cannot modify the orbital elements themselves as they have strong incentives to maintain positional accuracy to avoid conjunctions. In practice, there may be additional constraints (such as a correlation between object identification number and launch

Table 2: Decision Tree Performance Comparison

	Decision Tree	Bagged Tree	Random Forest	Histogram Boosted
Accuracy	0.89	0.91	0.93*	0.92
Precision	0.91	0.94*	0.94	0.91
Recall	0.87	0.86	0.92	0.93*
F1-Score	0.89	0.90	0.93*	0.92
TPR	0.87	0.87	0.92	0.93*
TNR	0.91	0.95*	0.94	0.92
Area Under ROC	0.89	0.91	0.93*	0.92
Train Time (s)	16.09*	30.33	81.77	19.56

* denotes the column with the best value for a metric. The target for true-positive rate (TPR) and true-negative rate (TNR) is "Satellite."

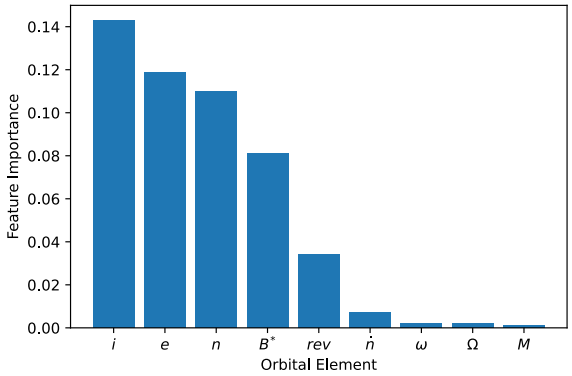

Figure 7: Random Forest Feature Importances

date) that may serve as further signatures and improve defender capabilities beyond this "worst case" scenario.

These classifiers are trained using TLEs which have been appropriately segregated on the basis of object identifier and observation time according to the considerations in Section 6.1. The training set consists of ~800,000 TLEs describing ~15,000 space objects. The test set consists of ~200,000 TLEs describing ~4,000 space objects.

The relative performance of these classifiers is summarized in Table 2. We find that all four models far-exceed random-guessing in terms of their predictive ability to determine an object's type on the basis of its TLE. The histogram boosted model offers the greatest performance to training-time ratio. However, the practical difference compared to a basic random forest model is marginal, with the random forest slightly outperforming in some metrics (such as TNR) while slightly under-performing in others (TPR). As a few minutes of additional training time has essentially no impact on testing our hypothesis, we have elected to focus on the random forest approach as it is widely used and well understood in both aerospace and security communities.

The dominant features leveraged by the random forest classifier are summarized in Figure 7 in terms of permutation importance [5]. These features make sense in light of our reasoning on the plausible physical differences that could be used to distinguish debris objects from satellites. i, e, and n roughly express the shape of an object's orbit and its altitude. While Ω, ω and M also play some role in this regard, Figure 5 suggests distribution across the full range of these values, while i, e, and n sit within tighter bounds. Additionally, the importance of B^* aligns well with the supposition that debris orbits tend towards different shapes and altitudes. We find that, contrary to our initial intuition, *revs* is a relatively weak indicator. Indeed, excluding it entirely has no meaningful impact.

One notable feature influence is that the median eccentricity of satellites which successfully evaded detection by our final random forest model was approximately 0.0187, a full order of magnitude greater than the median eccentricity of those satellites which were correctly detected (0.00167). This suggests one weakpoint in our model might be its ability to detect spy satellites in abnormally elliptical orbits. More broadly, that the dominant features are directly tied to the shape and altitude of an object's orbit, rather than its mass or motion, suggests that an attacker seeking to circumvent our model would need to materially alter their satellites' orbits to evade detection. Thus, we expect our approach, at a minimum imposes meaningful costs in terms of the effort and complexity demanded of attackers designing stealthy missions (see Section 7.1).

6.2.1 Final Model. A balanced random forest classifier does not quite meet our defender's requirements. After all, our defender doesn't want to sort isolated TLEs into categories, but rather to determine the nature of the underlying space objects behind them.

We can better serve this need by implementing a voting mechanism which aggregates label probabilities across all of an object's TLEs and applies the most favored label from multiple observations. This improves the Random Forest's performance slightly: TPR increases from 0.92 to 0.93, TNR from 0.94 to 0.95, and area under ROC from 0.93 to 0.94. One might expect a defender could further improve on this by incorporating history as part of the feature space directly. However, grouping the previous 20 TLEs of a given object into a single feature-set actually results in worse performance due to increased dimensionality (*f1-score: 0.91*).

The voting model we use offers an additional benefit by providing a simple mechanism for incorporating defender astrometry capabilities into the classifier's tuning. By decreasing the voting threshold required for classification, a defender can improve their chances of detecting a hidden satellite at the cost of an increase in false positives. For example, reducing the voting threshold from a balanced 50% to 25% increases TPR to 98% with an FPR of 11%. For mid-sized SSA powers, such as China and Russia, these numbers likely represent manageable false positive rates.

6.3 Defense Evaluation

While the machine-learning accuracy metrics tell us about system performance in pure numerical terms, it is difficult to understand what a 93% TPR means for defenders in terms of required astrometry capabilities and investment.

One way to contextualize these values is by simulating a hypothetical attack scenario several thousand times and seeing how the defense performs. For this simulated scenario, our attacker is assumed to launch two covert spy satellites over the course of a single year. The TLEs for these malicious objects are modified according to the process outlined in Section 4.2 and then intermingled with a set of 398 randomly selected legitimate debris objects from

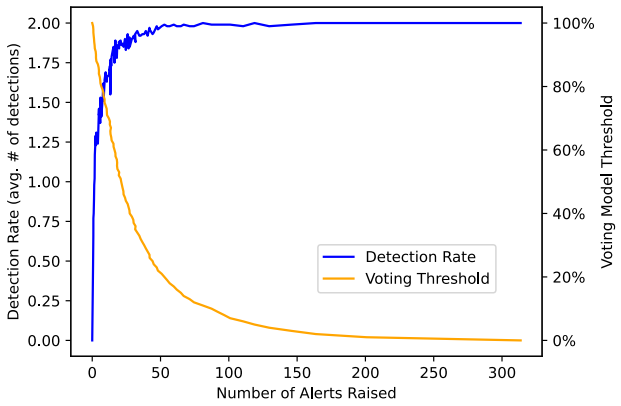

Figure 8: Defense Performance over 10,000 Simulations. At high voting thresholds, a defender receives fewer annual alerts, but may miss the spy satellites.

4,000 contained in the test dataset outlined in Section 6.2. We chose this number as it is roughly equivalent to the number of debris entries added to Space-Track's catalog in 2019. The simulation is thus analogous to the sort of catalog update information a defender would need to be capable of scrutinizing over the course of a single year. To improve simulation fidelity, we repeated this process 10,000 times with randomly selected object combinations.

The defender's task is to identify the spy satellites. However, we can also use this scenario to better understand the relationship between defender capabilities and attack detection rates by considering the voting-threshold discussed in Section 6.2.1 as a second independent variable.

The results are summarized in Figure 8. They suggest that even a small investment in SSA verification capabilities can substantially improve a defender's likelihood of detecting attacks under realistic conditions. For example, a defender with the ability to measure 50 RSOs each year can set their voting probability threshold to 22%, allowing them to detect at least one spy satellite in 100% of the simulations and both spy satellites with a 98% success rate. Even a defender with the ability to only vet one object per month can detect at least one spy satellite 96% of the time and both 70% of the time. In a "worst-case" scenario where the defender has no astrometry capabilities, the defender's best guess of the two most suspicious objects includes at least one disguised satellite 82% of the time and both 35% of the time.

6.3.1 2014-28E. As mentioned in Section 3.2, the closest empirical example of this attack is the case of NORAD object *2014-28E*. We can use the *2014-28E* incident as a demonstrative test of our model by excluding it from the random forest's training data. This is a particularly interesting as it was launched with the conscious intention of evading detection and "behaving" like a piece of space debris. Moreover, the attacker (Russia) and the victim (the US SSN) are the world's two dominant SSA powers.

We find that our model flags *2014-28E* as a disguised satellite from the very first TLE entry in Space-Track.org - six months faster than existing monitoring techniques did. It does so at a > 91% confidence, which equates to 2.55 total anomaly detection alerts per year under the conditions simulated in Section 6.3. This suggests that essentially any implementation of our system - whether or not it assumes defender access to astrometry equipment - would have identified Russia's claim that *2014-28E* was space debris as dubious.

Table 3: "Analyst Objects" Most Like Satellites

NORAD Catalog Identifier	Satellite Classification Confidence
85412U	0.980
87934U	0.962
84006U	0.960
87730U	0.957
82939U	0.945
89177U	0.928
82891U	0.911
85410U	0.910
84353U	0.903
85205U	0.895

6.3.2 Analyst Objects. As a matter of interest, the US SSN maintains a secondary list of objects outside the main catalog. These "Analyst Objects of Unknown Origin" are tracked by the SSN, but the US military has been unable to determine their origin. One would expect the vast majority to be orbital debris and, indeed, the SSN states that "Analyst objects that meet well-tracked criteria are generally debris objects, as is the majority of the space catalog" [47]. By applying the defensive classifier to these TLEs, it may be possible to determine which are most like concealed satellites.

We test this experimentally on a dataset of approximately 23,000 ephemerides relating to 519 analyst objects over a one-year period. At a 50% voting confidence threshold, 43 objects are flagged as potential satellites. Our expected false positive rate at this level is 5.7%. Increasing the threshold to 89% provides a list of the ten analyst objects most like satellites (Table 3). The false positive rate at this threshold was only 0.4%, so one would expect, in absence of deception attacks, only two objects to appear on such a list.

This is, of course, not conclusive proof that the US military either accidentally or deliberately miscategorizes operational satellites as analyst objects. The astrometry equipment necessary to ascertain the nature of these objects is well beyond our means as researchers. However, these ten entries represent an intuitive starting point for states interested in evaluating such a possibility.

7 DISCUSSION AND FUTURE WORK

Under our threat model, attackers initially appeared incredibly powerful, with full control over access to "ground truth" information regarding the nature of objects in orbit. However, a closer analysis shows that strong environmental and strategic considerations bound the extent to which attackers are able or willing to modify SSA claims. Acknowledging the existence of these constraints has allowed us to present an approach by which SSA recipients can significantly increase their chances of detecting satellites masquerading as debris objects.

7.1 Implications for SSA Operators

Given these findings, one might expect attackers to alter their space missions to improve stealthiness. For example, they might launch spy satellites into orbits with properties that cause them to be detected as debris by our trained model as a sort of "adversarial attack." While this is certainly possible, all of the relevant data is public knowledge, it is not without cost. Designing such a mission would almost certainly result in a less-than-idea orbit for other objectives (e.g by shortening mission lifespan or reducing coverage of key surface locations). It may require special maneuvering hardware to correctly inject the satellite into a debris-like orbit increasing complexity and weight whilst decreasing physical stealth. Even where these factors can be mitigated, adding an additional requirement to an already complex optimization space would have negative consequences for mission costs and development time. Efficiently overcoming these barriers could be one avenue for future work.

A more promising tactic for attackers would be to modify the contents of SSA *data* while keeping the mission unaltered, or poisoning the contents of the SSA repository in order to prevent training of classifiers more generally. The reasons we mentioned in Section 4 for assuming the general trustworthiness of SSN data are also confounding factors to consider when implementing such an adversarial strategy. Attackers have strong incentives to keep forecasts generated from this public data accurate to prevent orbital collisions. Even minute modifications to orbital elements can result in changes on the orders of hundreds or thousands of kilometers in magnitude over multi-day forecast windows. As such, striking a balance between adversarial attack effectiveness and TLE usability represents a difficult, but not impossible, option for improving these attacks in the future.

In the short term, the best option for SSA operators who own stealth satellites may be simply not reporting them at all. This is far from ideal since, if a foreign state *does* detect an unlisted space object, its very omission can serve as an indicator of its nature.

7.2 Implications for SSA Recipients

For SSA-recipients who are contemplating the need for upgrades to domestic SSA capabilities, our research shows that even small investments in astrometry can be leveraged effectively when coupled with our anomaly detection model. States do not need to go toe-to-toe with the US military on SSA capabilities in order to impose effective constraints on the ability of third parties to abuse their trust in shared SSA. Likewise, SSA-sharing need not be treated as a zero-sum game. "Trust but verify" systems are possible which allow states to catch deception attempts while also permitting them to reap commercial and diplomatic benefits from continued international SSA cooperation.

7.3 Implications for Security Researchers

This paper represents one of the first attempts to identify and simulate a credible threat to SSA data. We present the case that systems-security thinking can bring valuable and novel solutions to cross-disciplinary problems in the domain. However, our work also puts forward several questions for future research.

Our paper deliberately prioritizes an attack with real-world empirical examples, however, there may be interest in considering previously unseen attacks on SSA data. Attacks which target the orbital motion elements themselves, in order to alter collision forecasts, are more difficult to replicate due to the lack of unclassified information on how such data is generated and used in conjunction forecasting. However, prior work suggests that, if such attacks are feasible, they would be of great strategic importance to both SSA operators and actors seeking to cause physical harm to space systems [40]. Similarly, research on attacks which target the CMOS camera systems or radar sensors used in SSA data collection, although quite expensive to demonstrate on realistic hardware, could answer important strategic questions regarding the capabilities of advanced persistent threat (APT) actors to harm space surveillance missions. Finally, as mentioned in Section 7.1, a variant threat model which considers adversarial machine learning techniques vis-a-vis environmental trade-offs could be of significant utility to spy satellite operators.

Finally, defensive systems which can incorporate data from additional sources represents one avenue for further improvements. These may be low cost measurements, such as radio emission data or hobby telescope photographs. However, they could also be more hybrid metrics, such as the combination of our ephemeris data approach with photographs from a single space observatory.

8 CONCLUSION

The security and authenticity of SSA is critical to the responsible and sustained use of outer space. However, the "blind-trust" nature of SSA sharing creates opportunities for malicious deception attacks. Defenders have limited access to ground truth and, as a result, appear vulnerable to deception.

In this paper, we considered the specific case when an SSA operator would seek to deceive third-party SSA recipients into falsely believing that a satellite is merely a piece of space debris. Through a combination of real-world data and simulation, we present a method for simulating these attacks and use it to train a machine-learning anomaly detection tool. In evaluating our tool, we found that defenders are far from powerless. Our system empowers defenders with little or no astrometry hardware to reliably detect disguised satellites. Within reasonable bounds for the number of false positive alerts, we find defenders can detect 90-98% of deception attempts.

These results have broad implications. Our experiments show that even the most basic and essential information about an object's motion - less than 140 characters of TLE-data - can be reliably leveraged to detect disguised satellites. In the status quo, space powers are making critical decisions regarding investments in domestic SSA capabilities and next-generation stealthy nano-satellite platforms. Our findings suggest that many of assumptions underpinning these policy actions are far from certain.

While dependency on third party SSA will always require some degree of trust, this paper shows how the application of proven systems security methods can help balance historical space asymmetries - without the use of a single telescope.

REFERENCES

[1] Roya Afshar and Shuai Lu. 2020. Classification and Recognition of Space Debris and Its Pose Estimation Based on Deep Learning of CNNs. In *HCI International 2020 - Posters*. Springer, 605–613. https://doi.org/10.1007/978-3-030-50726-8_79

[2] I. Agi and L. Gong. 1996. An Empirical Study of Secure MPEG Video Transmissions. In *Proceedings of Internet Society Symposium on Network and Distributed*

Systems Security. 137–144. https://doi.org/10.1109/NDSS.1996.492420
[3] Saika Aida and Michael Kirschner. 2013. Accuracy Assessment of SGP4 Orbit Information Conversion into Osculating Elements. 723 (Aug. 2013), 160. http://adsabs.harvard.edu/abs/2013ESASP.723E.160A.
[4] Bandar Alotaibi and Khaled Elleithy. 2016. A New MAC Address Spoofing Detection Technique Based on Random Forests. *Sensors* 16, 3 (March 2016), 281. https://doi.org/10.3390/s16030281
[5] André Altmann, Laura Toloşi, Oliver Sander, and Thomas Lengauer. 2010. Permutation Importance: A Corrected Feature Importance Measure. *Bioinformatics* 26, 10 (May 2010), 1340–1347. https://doi.org/10.1093/bioinformatics/btq134
[6] Ball Aerospace. 2020. Ball Aerospace - Star Trackers. https://www.ball.com/aerospace/markets-capabilities/capabilities/technologies-components/star-trackers.
[7] Diogo Barradas, Nuno Santos, and Luís Rodrigues. 2018. Effective Detection of Multimedia Protocol Tunneling Using Machine Learning. In *27th {USENIX} Security Symposium ({USENIX} Security 18)*. 169–185. https://www.usenix.org/conference/usenixsecurity18/presentation/barradas.
[8] Niladri Das and Raktim Bhattacharya. 2019. Privacy and Utility Aware Data Sharing for Space Situational Awareness from Ensemble and Unscented Kalman Filtering Perspective. *arXiv:1912.03775 [math]* (Dec. 2019). arXiv:1912.03775 [math] http://arxiv.org/abs/1912.03775.
[9] James David. 2005. Was It Really 'Space Junk'? Us Intelligence Interest in Space Debris That Returned to Earth. *Astropolitics* 3, 1 (April 2005), 43–65. https://doi.org/10.1080/14777620590933584
[10] ExoAnalytic Solutions. 2020. ExoAnalytic Solutions. https://exoanalytic.com/.
[11] Gregory Falco. 2018. The Vacuum of Space Cyber Security. In *2018 AIAA SPACE and Astronautics Forum and Exposition*. American Institute of Aeronautics and Astronautics, Orlando, FL. https://doi.org/10.2514/6.2018-5275
[12] Jeff Foust. 2019. ConsenSys Space Announces Crowdsourced SSA Data System. https://spacenews.com/consensys-space-announces-crowdsourced-ssa-data-system/.
[13] R. Furfaro, T. Campbell, R. Linares, and V. Reddy. 2019. Space Debris Identification and Characterization via Deep Meta-Learning. In *First Int'l. Orbital Debris Conf. (2019)*. https://www.hou.usra.edu/meetings/orbitaldebris2019/orbital2019paper/pdf/6123.pdf.
[14] Forrest Gasdia. 2016. Optical Tracking and Spectral Characterization of Cubesats for Operational Missions. (May 2016). https://commons.erau.edu/edt/212.
[15] Mike Gruss. 2015. U.S. Plans $6 Billion Investment in Space Situational Awareness. https://spacenews.com/planned-u-s-investment-in-space-awareness-is-6-billion-gao-says/.
[16] W.J. Hennigan. 2020. Russian Craft Shadowing U.S. Spy Satellite, Space Force Commander Says. https://time.com/5779315/russian-spacecraft-spy-satellite-space-force/.
[17] Caleb Henry. 2018. Space Situational Awareness Experts Urge Russia to Join Orbital Neighborhood Watch. https://spacenews.com/space-situational-awareness-experts-urge-russia-to-join-orbital-neighborhood-watch/.
[18] Theresa Hitchens. 2020. New Satellite Imagery Rules Hover In Interagency Limbo. https://breakingdefense.com/2020/03/new-satellite-imagery-rules-hover-in-interagency-limbo/.
[19] Thorsten Holz, Benedikt Driessen, Ralf Hund, Carsten Willems, and Christof Paar. 2012. Don't Trust Satellite Phones: A Security Analysis of Two Satphone Standards. In *2012 IEEE Symposium on Security and Privacy*. IEEE, San Francisco, CA, USA, 128–142. https://doi.org/10.1109/SP.2012.18
[20] Felix R. Hoots and Ronald L. Roehrich. 1988. *Models for Propagation of NORAD Element Sets*. Technical Report. https://apps.dtic.mil/docs/citations/ADA093554.
[21] ISON. 2011. International Scientific Optical Network (ISON) & Low Frequency VLBI Network (LFVN). http://lfvn.astronomer.ru/index.htm.
[22] S. Jahirabadkar, P. Narsay, S. Pharande, G. Deshpande, and A. Kitture. 2020. Space Objects Classification Techniques: A Survey. In *2020 International Conference on Computational Performance Evaluation (ComPE)*. 786–791. https://doi.org/10.1109/ComPE49325.2020.9199996
[23] Thomas M Johnson. 2015. SSA Sensor Calibration Best Practices. https://amostech.com/TechnicalPapers/2015/Poster/JohnsonT.pdf.
[24] Mykola Kaliuzhnyi, Felix Bushuev, Oleksandr Shulga, Yevgeniya Sybiryakova, Leonid Shakun, Vladislavs Bezrukovs, Sergiy Moskalenko, Vladislav Kulishenko, and Yevgen Malynovskyi. 2016. INTERNATIONAL NETWORK OF PASSIVE CORRELATION RANGING FOR ORBIT DETERMINATION OF A GEOSTATIONARY SATELLITE. *Odessa Astronomical Publications* 29, 0 (2016), 203–206. https://doi.org/10.18524/1810-4215.2016.29.85228
[25] T. S. Kelso. 1995. Orbit Determination. https://www.celestrak.com/columns/v01n06/.
[26] T. S. Kelso. 2019. CelesTrak: "FAQs: Two-Line Element Set Format". http://celestrak.com/columns/v04n03/.
[27] M. Krüger. 2019. Detection of AIS Spoofing in Fishery Scenarios. In *2019 22th International Conference on Information Fusion (FUSION)*. 1–7.
[28] Bahavya Lal, Asha Balakrishnan, Becaja Caldwell, Reina Buenconsejo, and Sara Carioscia. 2018. *Global Trends in Space Situational Awareness and Space Traffic Management*. Technical Report. https://www.ida.org/idamedia/Corporate/Files/Publications/STPIPubs/2018/D-9074.pdf.
[29] Lasunncty. 10 October 2007 (original upload date). English: Digram Illustrating and Explaining Various Terms in Relation to Orbits of Celestial Bodies. https://commons.wikimedia.org/wiki/File:Orbit1.svg.
[30] LEOLABS. 2020. LEOLABS. https://www.leolabs.space/.
[31] Bin Li, Jian Huang, Yanming Feng, Fuhong Wang, and Jizhang Sang. 2020. A Machine Learning-Based Approach for Improved Orbit Predictions of LEO Space Debris With Sparse Tracking Data From a Single Station. *IEEE Trans. Aerospace Electron. Systems* (2020), 1–1. https://doi.org/10.1109/TAES.2020.2989067
[32] Bin Liu, Li Yao, and Dapeng Han. 2016. Harnessing Ontology and Machine Learning for RSO Classification. *SpringerPlus* 5, 1 (Sept. 2016). https://doi.org/10.1186/s40064-016-3258-2
[33] Bruce McClintock. 2019. Space Safety Coordination: A Norm for All Nations. https://www.rand.org/blog/2019/04/space-safety-coordination-a-norm-for-all-nations.html.
[34] NASA. 2018. ARES: Orbital Debris Program Office Frequently Asked Questions. https://orbitaldebris.jsc.nasa.gov/faq.html.
[35] M. Nayak, J. Beck, and B. Udrea. 2013. Real-Time Attitude Commanding to Detect Coverage Gaps and Generate High Resolution Point Clouds for RSO Shape Characterization with a Laser Rangefinder. In *2013 IEEE Aerospace Conference*. 1–14. https://doi.org/10.1109/AERO.2013.6496861
[36] Numerica. 2020. Space Defense | Tracking & Surveillance | Telescope Network. https://www.numerica.us/space-defense/.
[37] Joe Pappalardo. 2018. America's Next Spy Satellites Will Disappear. Here's How. https://www.popularmechanics.com/military/research/a25349950/nro-satellite-space-junk/.
[38] James Pavur, Daniel Moser, Martin Strohmeier, Vincent Lenders, and Ivan Martinovic. 2020. A Tale of Sea and Sky On the Security of Maritime VSAT Communications. In *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, San Francisco, CA, USA, 1384–1400. https://doi.org/10.1109/SP40000.2020.00056
[39] J. Puig-Suari, C. Turner, and W. Ahlgren. 2001. Development of the Standard CubeSat Deployer and a CubeSat Class PicoSatellite. In *2001 IEEE Aerospace Conference Proceedings (Cat. No.01TH8542)*, Vol. 1. 1/347–1/353 vol.1. https://doi.org/10.1109/AERO.2001.931726
[40] Reference Blinded for Peer Review. [n.d.].
[41] scikit-learn. [n.d.]. BaggingClassifier. https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.BaggingClassifier.html.
[42] scikit-learn. [n.d.]. DecisionTreeClassifier. https://scikit-learn.org/stable/modules/generated/sklearn.tree.DecisionTreeClassifier.html.
[43] scikit-learn. [n.d.]. HistGradientBoostingClassifier. https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.HistGradientBoostingClassifier.html.
[44] scikit-learn. [n.d.]. RandomForest. https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.RandomForestClassifier.html.
[45] Tommaso Sgobba, Firooz A. Allahdadi, and Fernand Alby. 2013. Orbital Operations Safety. In *Safety Design for Space Operations*. Butterworth-Heinemann, 411–431. https://learning.oreilly.com/library/view/safety-design-for/9780080969213/.
[46] N.N. Smirnov, A.I. Nazarenko, and A.B. Kiselev. 2000. Continuum Model for Space Debris Evolution with Account of Collisions and Orbital Breakups. *Space Debris* 2, 4 (Jan. 2000), 249–271. https://doi.org/10.1023/B:SDEB.0000029928.70053.a4
[47] Space Track. 2019. User Agreement. https://www.space-track.org/documentation#user_agree.
[48] Space Track. 2020. SSA Sharing and Orbital Data Requests. https://www.space-track.org/documentation#odr.
[49] SpaceQuest. 2017. GNSS-701 Satellite GNSS Receiver. http://www.spacequest.com/attitude-determination-control/gps12-v1.
[50] Swapnil Anil Surdi. 2019. Space Situational Awareness through Blockchain Technology. In *First Int'l. Orbital Debris Conf.* https://www.hou.usra.edu/meetings/orbitaldebris2019/orbital2019paper/pdf/6192.pdf.
[51] The Guardian. 2014. What Is Object 2014-28E – a Russian Military Satellite or a Piece of Unidentified Debris? *The Guardian* (Nov. 2014). https://www.theguardian.com/science/shortcuts/2014/nov/18/object-2014-28e-space-russian-satellite-unidentified.
[52] The Telegraph. 2015. Russia 'Busts Foreign Satellite Spy Ring'. https://www.telegraph.co.uk/news/worldnews/europe/russia/11531013/Russia-busts-foreign-satellite-spy-ring.html.
[53] Union of Concerned Scientists. 2020. Satellite Database. https://www.ucsusa.org/resources/satellite-database.
[54] David A Vallado and Jacob D Griesbach. 2011. Simulating Space Surveillance Networks. In *AAS/AIAA Astrodynamics Specialist Conference*.
[55] Megan Wallace. 2015. Tracking and Data Relay Satellite (TDRS). http://www.nasa.gov/directorates/heo/scan/services/networks/tdrs_main.
[56] Jiangbo Xi, Yaobing Xiang, Okan K. Ersoy, Ming Cong, Xin Wei, and Junkai Gu. 2020. Space Debris Detection Using Feature Learning of Candidate Regions in Optical Image Sequences. *IEEE Access* 8 (2020), 150864–150877. https://doi.org/10.1109/ACCESS.2020.3016761