ML-Based Photometric Fingerprinting at Scale for LEO Satellite Monitoring

T. Olson, C. Reid, C. Key, B. Williams, D. Witman, J. Shaddix, D. Kesler, B. Marchand
Slingshot Aerospace, 5475 Tech Center Dr. Suite 225 Colorado Springs, CO 80919

ABSTRACT

Space situational awareness (SSA) is becoming increasingly important as the number of objects in low Earth orbit (LEO) grows exponentially. Due to the congested and dynamic background of the LEO population, objects can be difficult to properly differentiate, and potential hazards can go unnoticed as satellites undertake activities that are not discernable from astrometric observations alone. These challenges can be addressed by exploiting photometric (brightness) information from optical measurements to enhance orbital intelligence by providing complementary signals about individual resident space object (RSO) characteristics. This work demonstrates a two-step machine learning (ML)-based approach developed by Slingshot Aerospace for characterizing satellites in LEO via optical sensors. First, we present a novel photometric fingerprinting algorithm that enables object characterization, identification, and classification. Second, we examine methods for embedding and comparing fingerprints with potential applications to attitude change detection using photometry. We demonstrate performance using a wealth of real LEO observations.

1. INTRODUCTION

Space domain awareness (SDA) has become an increasingly important area of development as investment in space skyrockets. The accurate and timely understanding of the space environment that is generated by SDA activities is required to reduce risk, optimize expenditure, and pave the way for future space operations. Conjunction risk assessment, which evaluates the likelihood of two objects colliding, is one area of SDA research that has received substantial focus in recent years. Projects that aim to capture and remove debris require SDA information about materials, spin rate, and orbital velocity. Anomaly detection algorithms use SDA information to monitor and protect precious space resources.

A myriad of different modalities can be deployed in sensor networks around the world to capture SDA information, ranging from radar and radio frequency (RF) systems to optical imaging cameras. Terrestrial telescopes often collect unresolved detections of light reflected off of the space objects they track. When recorded over a period of time, the reflections form a light curve, which can be used together to characterize various features of the tracked object. Features derived from light curves can provide valuable information to telescope systems that complement traditional orbital state estimates, helping to better identify objects, and surface characteristics such as stability, tumble rate, material, or object geometry.

Herein, we discuss a novel approach to object fingerprinting developed by Slingshot Aerospace that is based on light curves, and that has been tested across a large number of objects and multiple geographic locations. Our approach could be applied to all of the objects in low earth orbit (LEO), but for the purposes of this exploration we focus on differentiating three Starlink bus type variants in particular. Fingerprints and the best methods for comparing them are evaluated across multiple observation sites using a variety of different sampling techniques.

2. BACKGROUND

2.1 Light Curve Inversion

Light curve inversion, or modeling the surface of a rotating object given its brightness variation, is a problem that has long been considered. The heart of the challenge is that, given a set of known variables (e.g. object position and orientation), and a set of observed signals (e.g. object brightness over time), one can attempt to reconstruct the geometry and features of the object. Early works include estimations of the shape, rotation period, pole direction,

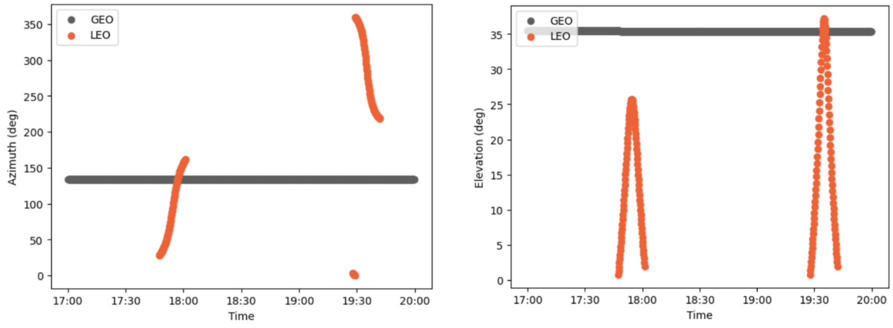

Fig. 1: Difference in satellite to sensor angles over time by regime

and scattering parameter of an asteroid [1]. Over time, it became apparent that the same approach could be used to estimate the shape and stability of nearer objects in geosynchronous orbit (GEO), and more recently, in LEO. Although direct reconstruction of detailed satellite geometries and material properties remains out of reach, the prior work provides a solid foundation for feature set creation and the application of advanced machine learning (ML) methods to light curve data to enable effective characterization, classification, and differentiation of orbiting bodies. Three primary considerations are discussed below for creating feature sets to support ML-based object classification using photometric data.

2.2 Object Position and Light Sources

There are a number of aspects that influence the particular light curve created by an object passing through the field of view of a sensor. The first set of factors affecting the light curve are the orbital positions and angles relating the object position, the sensor position, light source positions, and the relationships between them. The angle between the sensor and the sun, centered on the reflecting object, is defined as the solar scatter angle (or solar phase angle [2]), and it represents the reflection angle of incident light from the observer's perspective. The angle of the object relative to the sensor, often represented as an azimuth and elevation, describes the direction and distance over the horizon an object is positioned. In addition to direct sunlight scattering, reflected sunlight off Earth's surface, called "Earthshine", also contributes a non-negligible illumination component whose strength depends on the Earth surface materials' reflectivity and the position of the satellite over Earth's surface, relative to the solar light source. As an object travels through the sky relative to the sensor, the scattering angles shift, and so the observed apparent magnitude (a measure of photometric brightness) will not only depend on the physical characteristics of the object, but also the dynamic environmental effects.

These angles and effects are substantially different between orbital regimes. The first and most obvious difference between orbiting bodies in LEO and GEO is in the length of the period of time an object is in the sky. Objects in GEO are placed in an orbit designed to remain steady relative to the sensor. The azimuth and elevation shift very little over time, and so collections of an object's brightness can measure long trends. In LEO, an object's relative position tends to shift much faster, leading to passes that typically span up to only a few minutes. The viewing geometry during the next visible pass will typically differ substantially from the previous, presenting challenges for direct comparisons. Fig. 1 shows the variability in angle of a satellite relative to a sensor over a three hour period for a geosynchronous satellite and a sun-synchronous LEO satellite. Consequently, algorithms designed to build baselines on long stretches of data to recover spin rate or find attitude changes that may work in GEO collapse when applied to LEO data.

A second factor affecting light curves is the variety in solar scatter angle. As objects in LEO are moving much faster than those in GEO, the range in solar scatter angle is substantially wider. This means non-geometry and attitude factors affecting the light curve are changing fast enough that they need to be accounted for in analyses and characterization, especially for LEO objects that tend to cover a significant portion of the field of view over the course of the night.

2.3 Object Geometry, Material, and Attitude

Object geometry, material, and attitude interact with the object position and light source to account for most variation in a light curve. These are features most lightcurve inversion algorithms are attempting to recover. Objects have a large number of potential geometries, which can be broadly grouped into three categories: rocket bodies, debris, and payloads. Rocket bodies tend to be cylindrical and large, whereas debris can range anywhere from large sheets of metal to smaller pieces of solar panels or flecks of paint. This paper focuses on payloads that also are variable in size and structure. They can range from extremely large school bus-sized GEO satellites to small cubesats and nanosats, which are small boxes deployed predominantly in LEO. The body can be symmetrical or asymmetrical, and the placement of solar panels and paint on the payload affects the reflectivity of light as orientation changes.

The third factor tied to unknowns about the object is the attitude and rotation rate. Debris, rocket bodies, and uncontrolled payloads can spin along an unknown axis at an unknown rate, since they have no thrusters or reaction wheels to help stabilize the body as it travels along its orbit. Controlled payloads are different in that they often are equipped with thrusters and reaction wheels designed to hold the body steady[1] – payloads often maintain a specified attitude to achieve their mission. For example, communications satellites must be pointing in the correct direction to orient antennas and properly provide internet service, whereas imaging satellites must be aiming cameras in the direction of the desired image subject. When a vehicle changes attitude, scattering surfaces on the payload rotate relative to the illumination sources, and the apparent magnitude may shift. Additionally, some attitude changes may cause the dominant scattering contributions to migrate from a less reflective surface, like the side of a satellite, to a more reflective surface, like a solar panel, with corresponding modulations to the observed magnitude. This paper will address payloads that are assumed to be holding a constant attitude, to focus explicitly on recovery of satellite geometry and potentially on occasional attitude changes rather than recovery of spin rate or related fields.

The full physical scattering profile from an object is often encoded quantitatively as the bidirectional reflectance distribution function (BRDF). Prior research on the photometry of LEO communications satellites has shown that their operational configurations yield stable BDRFs that can be reconstructed from large data sets of photometric observations, which is an important requirement for the ML methods described below [3; 4].

2.4 Sensor Characteristics and Environmental Noise

The final factor with potential to influence light curve characteristics comes from the surrounding environment and the sensor system. Atmospheric transparency varies both spectrally and in total extinction, as a function of elevation angle (i.e., air mass) and the distribution of water vapor, droplets, and other particles within the light path. These local environmental features can affect the observed apparent magnitude of an object over time. As such, sky transparency effects are mitigated by comparing star apparent magnitudes with the Gaia G, Bp, and Rp band-passes for calibration, computing color correction terms for solar-illuminated objects along with a corresponding photometric zeropoint. Sensor systems also have a number of components that require addressing, including the optical assembly's spectral transmission, stray or background light in the optics path, satellite proximity to other light sources (e.g., stars), and the camera's saturation limit. Finally, optical sensor systems must also account for vibrations caused by wind or gimbal imperfections, local atmospheric seeing conditions, and any satellite/sensor relative angular rates that blur or smear the point spread function.

2.5 Machine Learning for Geometry Classification

Machine learning is well-positioned to provide more insight into extraction of information from light curves, especially as the processing power and amount of available data grows. Substantial effort has been aimed at evaluation of light curves and other optical data in GEO orbits, ranging across tumbling detection, object classification, and attitude derivation. Statistical and feature-based techniques have been applied to assess optical fingerprinting approaches for comparing RSOs under various conditions in GEO, medium Earth orbit (MEO) and high Earth orbit (HEO) [5; 6], and lightcurve data has been combined with astrometric data to estimate the mass of inactive GEO RSOs [7]. Each of these approaches served to prove that a large amount of information about GEO RSOs can be derived from optical data starting with strong statistical models.

A number of machine learning approaches to lightcurve analysis also have been studied. Random forest algorithms have been used against GEO objects with long optical histories to distinguish tumbling objects from stabilized spacecraft [8]. Recurrent neural networks (RNNs), convolutional neural networks (CNNs), multiple kernel learning (MKL), and meta-learning methods have been applied to shape, attitude, and type classification for near-GEO objects based upon lightcurves [9–11].

There has been less work addressing RSOs in LEO orbits, though common ML approaches like Support Vector Machines and Long Short-Term Memory Networks have been applied to discriminate between stable satellites, tumbling

---

[1] There are notable exceptions to this generalization, such as spin-stabilized systems that maintain their orientation in a subset of directions while intentionally rotating around a spacecraft axis.

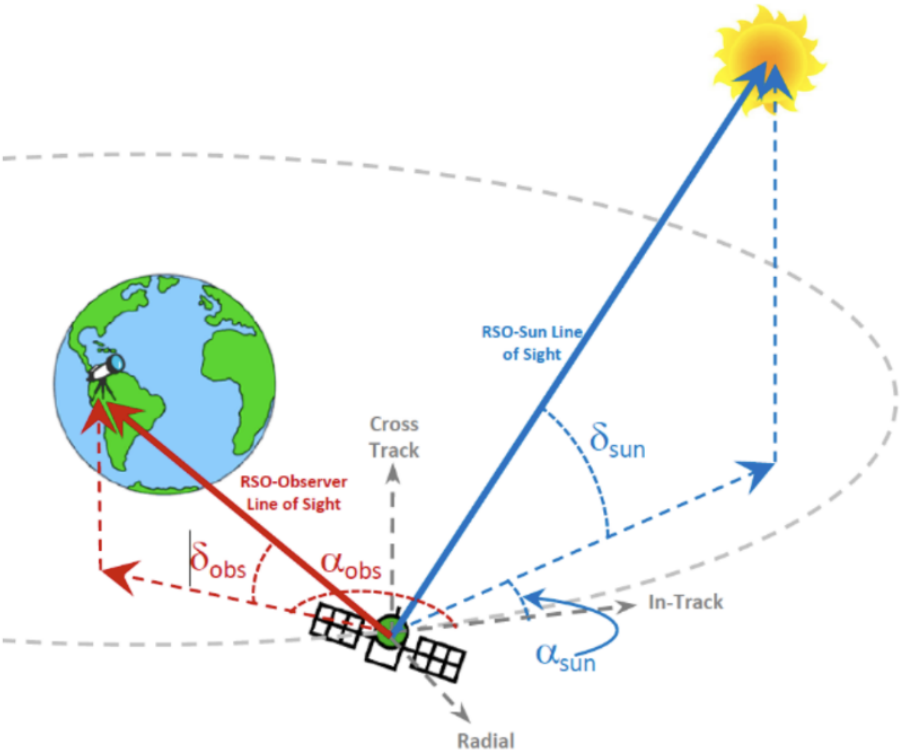

Fig. 2: Illustration of pointing vectors oriented from the satellite body to the primary illumination source (sun-RSO) and to the observer (sensor-RSO).

satellites, and rocket bodies using lightcurves [12]. Perhaps most relevant to the work discussed in this paper is the application of boosted decision trees to predict satellite platforms from light curves of LEO RSOs, which showed that there is sufficient information contained in LEO lightcurves to detect characteristics of different satellite platforms [13].

3. PHOTOMETRIC FINGERPRINT DEFINITION

Herein we define the central mathematical construction that will be processed through ML models for RSO classification and differentiation: the *photometric fingerprint*. It is a representation of an RSO via a numerical function of observed apparent magnitude over a collection of points in a four-angular-dimensional space. Mathematically, it is closely related to an optical BDRF, though it does not directly encode the same information as a BDRF.

Let us assume that photometric measurements have been collected for a number of distinct orbiting bodies, and that approximate position and velocity state estimates were obtained for those objects at the same times that the photometric observations were collected. The position and velocity vectors from the orbital state estimates define a co-moving RSW/RIC (radial-intrack-crosstrack) frame for the RSO. For spacecraft in controlled, stable attitude configurations, that frame is expected to coincide with the body frame of the object. Both the sun and observer follow distinct paths from the point of view of the RSO; at any instant in time, one can define two pointing vectors, one toward the sensor, and another toward the sun, see Fig. 2. For photometry, this perspective decouples the motion of the illumination source(s) – direct and reflected solar irradiance – from the relative trajectory of the observer, with respect to the static body geometry of the object. For a given arrangement of sun and RSO (i.e. a fixed sun-pointing vector in the body frame), different observer-pointing vectors induce reflections off different facets, structures, and materials of the body. Conversely, given a static sensor-RSO vector in the body frame, various arrangements of the illumination sources highlight different geometric features of the spacecraft.

Over the course of a single satellite pass over a ground sensor or through the field of view of a spaceborne sensor, the co-moving frame translates and rotates along with the orientation of the spacecraft, which in this case we assumed stable with respect to the Earth. Consequently the sensor-RSO and sun-RSO pointing vectors both evolve in time, and the light curve collected throughout the pass consists of reflections off multiple facets under varying illumination conditions. While this phenomenon leads to challenges for comparing individual light curves created by different facets of the spacecraft, it lends itself very well to a fingerprinting approach.

Each photometric sample includes an observation of the apparent magnitude at a point in the combined four-angular-dimensional space. Under the aforementioned condition of an RSO with perfectly stable relative orientation, a sparse fingerprint representation encompassing a diverse set of angles can be assembled for each object by combining observations across one or more satellite passes. This is much like sampling different points in a two-dimensional human fingerprint. In practice, however, it may be possible to collect multiple (possibly conflicting) observations at identical or nearby positions in parameter space,[2] so the final step in constructing a photometric fingerprint is to aggregate those multiple observations into a unified representation of the object over the entire four-dimensional space. We will discuss our approach to aggregating multiple measurements for stable objects in the Methods section below, and present alternative extensions in the Discussion section at the end.

We finally note for comparison that GEO satellites tend to have effectively static sensor-RSO vectors, and in near-equatorial orbits the solar pointing vector rotates around one axis on a daily period while the yearly period in the orthogonal angular direction leads to slow, low-amplitude variation in that direction. The phenomenon results in GEO RSOs that are well characterized by two-dimensional fingerprints; even a single-dimensional fingerprint (i.e. a light curve) covers almost the entire accessible parameter space, especially after accounting for the slow seasonal variation. LEO satellites, on the other hand, exhibit rapid variation in all four angular directions: two sensor-RSO angles and two sun-RSO angles. This is a principal reason that LEO light curves and fingerprints are much more complex and challenging to process compared to GEO.

4. METHODOLOGY

This paper proposes a new ML-based approach to characterize and classify RSOs in LEO via photometric fingerprinting techniques. Using a large observational dataset collected from Slingshot's global network of Horus optical fence systems [14], we aggregate multiple photometric measurements into numerical representations of individual RSOs and satellite bus families that uniquely distinguish them from others. We call these representations photometric fingerprints, by which we mean concrete, densified versions of the abstract fingerprints introduced in Section 3. We then trained a neural network to quantify the similarity between fingerprints and identify candidate satellite bus classifications for various objects. To demonstrate the effectiveness of the proposed methods, we show that three distinct collections of Starlink satellites can be distinguished by their photometric fingerprints.

For the analysis presented in this paper, we focused on first-generation Starlinks, as they made up the majority of the Starlink population in orbit at the time this study began [15]. The full set of first-generation satellites were further designated into three variants, each with a different configuration of body features that affect the amount of solar radiation reflected toward Earth. The three are detailed below.

- Version 1.0: The original production run of Starlinks
- Version 1.0-visorsat: A modified bus with an additional sun shield “visor” [16]
- Version 1.5: An upgraded iteration of the original version[3] [3]

For the purpose of this study, we assumed all RSOs with a given bus variant had consistent BDRFs, so that all observations of any object in the bus family could be assimilated into a single reference fingerprint for the entire variant – some variations were present in reality, but on average this was found to be a reasonable approximation. A principal finding of our investigation is that the reflectance properties mainly differed between objects in different bus families, presumably due to variations in shape and material. Those differences were substantial enough to yield unique photometric signatures that were captured by the fingerprinting technique.

---

[2] This would be the norm for unstable and tumbling objects.

[3] Starlink v1.5 satellites have undergone additional design evolution through multiple sub-variants, the effects of which were not considered in this study [3].

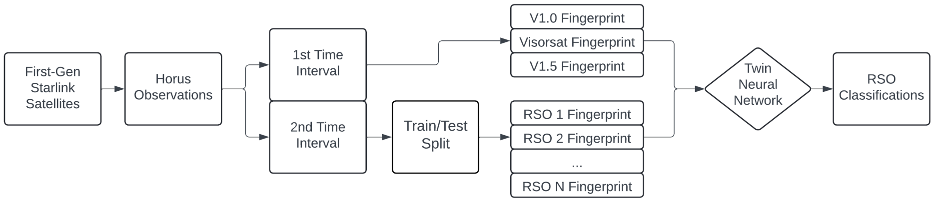

Fig. 3: Classification Workflow

We further demonstrate that fingerprints derived from future observations can be compared against prior established fingerprints to correctly identify the Starlink bus family from which those observations were sourced. This enables a novel approach for identifying and classifying RSOs based on photometric signatures, and it provides a valuable augmentation of traditional trajectory-based object correlation strategies.

4.1 Data Source

4.1.1 Observational Data Collection

Slingshot operates a global network of geographically-diverse Horus optical fence systems, together providing uncued tracking of LEO objects as small as cubesats across both northern and southern hemispheres. Each Horus system continuously images the majority of the sky visible from its ground station, allowing simultaneous long-pass tracking on transiting objects while gathering both astrometric and photometric observations. With uncued processing, Horus routinely detects objects with poor or nonexistent prior orbital state knowledge, allowing regular collection on new foreign launches, newly-released payloads, and objects not present in public catalogs. Thus, the Horus network is ideal for gathering information in support of object characterization and anomaly detection.

4.1.2 Data Preparation

We identified satellites of interest, collected optical and positional data related to those satellites, and then split the observations for experimentation. We selected objects with orbits that lent similar viewing angles, to reduce chances that our model confounded lightcurves and sky position. We filtered to active satellites with one of the three variants described above. The variants were assigned using documentation in the Slingshot Seradata database [15]. We then imposed an orbital constraint by selecting only satellites that were orbiting at 53 degrees inclination. The active Starlinks are kept at generally the same inclination over time, so we were able to detect these by averaging the inclination of the satellite orbits and identifying those station keeping at 53 degrees. Once the selected objects were identified, we gathered all Slingshot observations and state vectors from four Slingshot Horus optical fence locations for one month. Each observation site was processed separately to ensure consistency among the accessible viewing geometries from each sensor location.

The data included the apparent magnitude of an observation and the orbital state for the object at that time. The satellites of interest were typically represented by 1000 to 7000 observations over the time period, with a relatively steady revisit rate. Observation here refers to a set of astrometric and photometric measurements made on a detection of an object within an image or stack of images. The apparent magnitude measurement was computed from the luminous flux collected by each sensor, corrected and normalized by accounting for background sky illumination and known reference star magnitudes. The process yielded comparable apparent magnitude measurements for each observation. Orbital states were obtained by internal orbit determination methods.

4.2 Model Details

Our classification model was constructed as follows: Horus observations were divided into two time intervals and the Starlink RSOs were divided into training and test sets. All of the observations over the first time interval were used to construct fingerprints of the three variants. The training set observations over the second interval were used to construct individual fingerprints and train a twin neural network. Then classification was performed against the test set RSOs to classify which variant each RSO belonged to. The full workflow is provided in Fig. 3. In the remainder of this section, we describe this workflow in more detail.

4.2.1 Fingerprint Aggregation

The photometric fingerprints introduced in Section 3 required the aggregation of multiple photometric measurements into a unified fingerprint. We addressed this with a sequence of binning and smoothing operations. To reduce compute requirements, we also reduced the dimensionality of the fingerprints by projecting into a three-dimensional subspace. To convert the sparse representation of angles and observations into a dense fingerprint, we first divided the observation space into multi-dimensional angular bins. The photometric observations were then each assigned to their nearest bin, tracking total observation counts separately from the sum of observed magnitudes. We applied a smoothing operation to the sparse data by convolving a Gaussian kernel with both the count and sum metrics. Then, to reduce noise from areas with too few nearby observations, a post-convolution filter was applied to mask out bins with convolved-count values below a minimum threshold. The dense photometric fingerprint was then completed by computing a ratio of the two smoothed metrics, which yielded a numerical representation for each point in the unmasked parameter space.

A photometric fingerprint could have been generated from any collection of observations; there was no strict requirement that the included observations come from a single sensor, or even from the same object. However, observations from objects with substantially different physical characteristics (shape, materials, attitude, etc) were expected to wash out in aggregate because of the smoothing process, yielding bland averages with little effect. On the other hand, multiple objects with very similar physical characteristics yielded consistent photometric patterns, and fingerprints generated from their combined observations yielded useful averages with extensive coverage of the parameter space, due to the diversity of angles at which observations of multiple objects were collected.

This latter detail was valuable for establishing baseline reference fingerprints, against which future observations could be compared. While it was possible to generate a reference-fingerprint for any single object by collecting observations over long periods of time at different solar and sensor angles, combining the observations of multiple similar objects enabled rapid construction of a numerically dense representation for each family of objects. For the analyses presented in this paper, we produced reference-fingerprints for each of the three Starlink variants described above by combining all observations of objects within a family for a particular site and over a selected time period into the reference-fingerprint for that family.

Once the baseline reference-fingerprints were produced with dense coverage of the parameter space, new inference-fingerprints were constructed on a per-RSO basis from collections of additional observations, and those were compared to the reference-fingerprints to determine relative similarity via the comparison methods described in the next sections.

4.2.2 Fingerprint Comparison

We compared fingerprints using two methods: a baseline masked distance calculation and a trained neural network based on the “twin network” architecture [17]. The twin network is a machine learning based pattern for pairwise comparison of complex feature spaces. We trained the network to map fingerprints into a latent embedding space, such that inference-fingerprints from each family were embedded closer to the embedding of their reference-fingerprint than the others. For inference-fingerprints from unknown RSOs, the “best match” reference-fingerprint then provided the classification. The experimental data setup is described below.

4.2.3 Mean Absolute Difference (MAD)

We computed the masked distance metric, or the mean absolute difference (MAD), between two fingerprints, disregarding points that were masked in either fingerprint. To classify an unknown inference-fingerprint, we calculated the MAD from each reference-fingerprint and took the one with the smallest distance as the class selection. This technique was the most straightforward approach for comparing reference fingerprints to inference-fingerprints, and so we used it as a baseline for more sophisticated model results.

4.2.4 Twin Network

We created a model that embedded the fingerprints into a reduced-dimensional latent space representation such that the representations cluster by bus type within the embedding space for comparison of reference-fingerprints to inference-

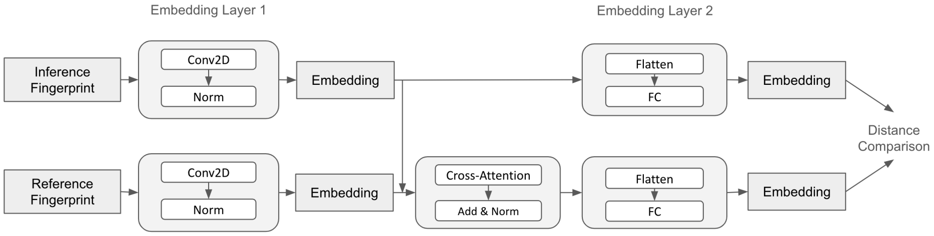

Fig. 4: Twin network architecture

fingerprints. We trained a twin network architecture [17] using a conventional neural network training method with batches of randomly chosen inference- and reference-fingerprints from the training dataset described above. A twin network processes an input (in this case the fingerprints) and reduces the parameter space to a lower-dimensional embedding-space vector. Our network used two stacked embeddings. The first used the full fingerprint and reduced it to a smaller convolved space. The reference-fingerprint embedding was then updated through a cross attention layer, conditioned on the inference-fingerprint. The embedded inference-fingerprint and the embedded and attention-updated reference-fingerprint were then fed into the second encoder, which reduced the samples to vectors of length 8. The model architecture is illustrated in Fig. 4 . Each batch was trained with contrastive loss – the distance between reference-fingerprint and inference-fingerprint belonging to the same (different) class was minimized (maximized) by iterative updates to the network weights.

Reference aggregation data: All observations from a particular site for the first three weeks of the time period were labeled and grouped according to their known Starlink variant. The observations were then aggregated into one reference-fingerprint for each variant.

Model training data: The last week of the time period was partitioned into train and test sets, randomly assigned by NORAD ID; 80% of the RSOs in each family were reserved for model training and the other 20% were reserved for performance evaluation. This separation ensured that no observations in the evaluation dataset were physically associated with any in the training dataset. The observations were then grouped by time into 24-hour batches. At least one inference-fingerprint per batch was generated for each RSO (excluded if no observations were collected for that RSO in that time period). This setup was designed to test the efficacy of the fingerprints as they would likely be used in an online system for attitude change analysis, cross-tag remediation, or other applications (see Discussion section for further considerations).

Model evaluation: A principal application of the fingerprints could be the detection of differences from expected photometric signals, so we evaluated performance by measuring their ability to classify unknown fingerprints into the correct bus model variant. For each reference-fingerprint, the network processed the reserved evaluation RSOs from the test set and compared inference-fingerprints to the reference-fingerprint. The output of the model was a measure of similarity between the inference- and reference-fingerprints; the reference with maximum similarity was then selected as the candidate Starlink variant classification for each RSO. These candidate classifications were compared with the true class labels (Starlink variants) for those objects, yielding statistical performance measures including precision, recall, and f1-score. The results were taken overall, and then broken down by the binned number of observations used to create the inference-fingerprints. RSOs that were tracked with fewer observations on a day would fall into a lower bin, and RSOs with a higher number of observations would go into one of the higher bins. Examination of the scores broken down in this way provided some insight into how many observations would be needed on an RSO in a day to provide a useful inference-fingerprint.

5. RESULTS

We processed the data via the methods described above and found that, given well-established reference-fingerprints, inference-fingerprints generated from small numbers of observations, even for a single RSO, provided effective tools for discriminating between the possible classes.

Fig. 5 illustrates reference-fingerprints generated for each of the four sites used in this analysis. Two of the sites were

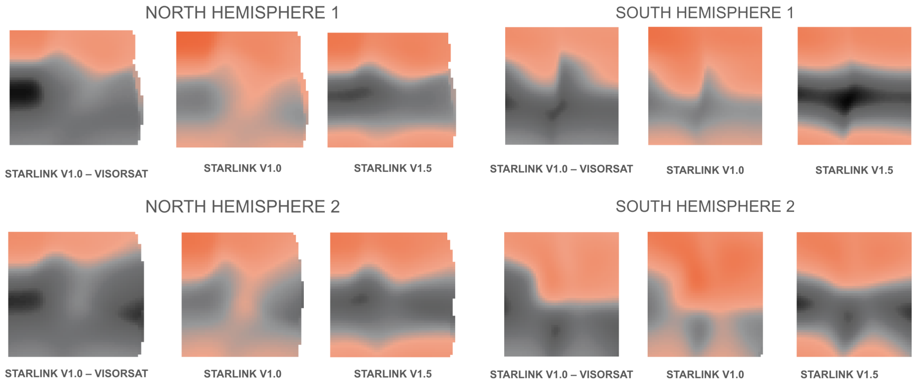

Fig. 5: Reference Fingerprints for four sites

Table 1: Evaluation f1 scores averaged across sites for daily model

		VISORSAT	**v1.0**	**v1.5**	**Average**
<20 Obs	*MAD*	0.55	0.04	0.14	0.24
	Twin	0.35	0.06	0.61	0.34
<40 Obs	*MAD*	0.47	0.19	0.37	0.34
	Twin	0.59	0.19	0.77	0.52
<80 Obs	*MAD*	0.53	0.25	0.44	0.41
	Twin	0.63	0.37	0.8	0.6
<160 Obs	*MAD*	0.63	0.43	0.57	0.54
	Twin	0.84	0.59	0.87	0.77
<320 Obs	*MAD*	0.77	0.64	0.73	0.71
	Twin	0.92	0.79	0.94	0.88
<640 Obs	*MAD*	0.79	0.68	0.77	0.75
	Twin	0.96	0.87	0.96	0.93
>640 Obs	*MAD*	0.94	0.72	0.79	0.82
	Twin	0.99	0.85	0.94	0.93

in the northern hemisphere. The reference-fingerprints were created using all observations from each site and for each variant across a set time period. The visualizations were projected from the true three-dimensional fingerprints, down to two dimensions for visualization purposes. Note that the fingerprints were roughly consistent across hemispheres. The apparent differences may be caused by the variety of latitudes at which the sensors were located, which affected the accessibility of different regions of solar angle parameter space. Since the sample data were collected over a relatively short period of time, compared to the full annual cycle, the higher-dimensional base fingerprints did not cover the full space of possible scattering angles. Reducing the projection to two dimensions exposed some visual artifacts that may reflect the seasonal shift. An analysis of fingerprints from samples collected throughout the year may be interesting for future work.

The classifications generated on the test set by the twin network resulted in a mean f1-score of 0.71. The metric here was reached by calculating the f1-score for each class and then averaging them. This allows us to measure the performance without over-representing the majority class. Table 1 shows the f1-scores broken down by class and by range of observations used to create the inference-fingerprint. Across the variant classes, the model did better when given an inference-fingerprint created on more observations, up to those created with >640 observations. Inference-fingerprints created with more observations than that saw a slight decrease in f1-score for the v1.0 and the v1.5 variants. The Visorsat and v1.5 Starlink variants had higher performance metrics than the v1.0 Starlink for every observation bin. This may be due to similarities between the variant geometries, but also may be due to v1.0 being a minority class. We had the smallest number of examples of Starlink v1.0 variants in our dataset, and so there were fewer examples to learn from and the reference-fingerprint was likely sparser.

6. CONCLUSIONS AND FUTURE WORK

Our results indicate that the novel fingerprinting method developed reference- and inference-fingerprints suitable for Starlink version classification, even when the inference-fingerprints were created using only the observations available in a brief 24 hour time period. The model showed reduced performance when the inference-fingerprint had a lower number of observations available over the course of a day. This supports the concept presented above: photometric fingerprints assimilate scattering characteristics across a range of viewing and illumination angles; RSOs with too few observations (or those only viewed at ambiguously informative geometries) are more difficult to classify because their fingerprints do not adequately populate the fingerprint parameter space.

One way to improve performance here may be to extend the timeframe that observations are taken across for the inference-fingerprints, as that would both increase the number of observations available and would increase the number of angles the observations cover as the RSO orbit shifts overhead. A future study investigating how fingerprint performance evolves with different time-windows would help to shed more light on this aspect of the fingerprints. Our results lent insight into the optimal number of observations required to identify RSO variants. The daily fingerprints required at least 160 observations for the twin network to exceed 50% accuracy on the minority class (Starlink v1.0). This figure may be used as a point of reference for future attempts to classify RSO variants.

Another interesting follow-up study might include the analysis of the photometric fingerprints for each variant, with an in-depth investigation into the reasons for performance differences among variant types. Further areas of focus might include an evaluation of which regions of the fingerprint parameter space are more or less effective at differentiating among variants. An identification of strongly-separated viewing and illumination geometries could then provide insight for prioritization of sensor tasking to more efficiently classify unknown objects.

6.1 Applications for Cross-tag Filtering and Candidate Rejection

Track correlation involves processes for tagging sequences of observations (tracks) with the cataloged objects most likely to be responsible for those observations. Track correlation has broadly presented challenges in the past, especially when multiple cataloged orbits pass within close proximity. Those close proximity passes may lead to cross-tagging, wherein the incorrect object becomes associated with the observations.

Photometric measurements have the potential to enhance track correlation algorithms in several ways. Photometric based fingerprints may provide an independent method for filtering the candidate object list based on comparisons between small inference-fingerprints with known reference-fingerprints for the candidates. In such applications, the observation period used in the inference-fingerprint may span only a few minutes to days. The results of our analysis provide a quantitative measure of candidate classification reliability as a function of the number of available observations from short periods of time.

Moreover, those performance metrics were based on a traditional three-class classification problem, wherein only a completely correct answer counted in favor of the model. In a production system there would be many more classes, and it is expected that even a small number of observations would be able to rule out many candidates. A small example of this is given in Fig. 6, where we have plotted two RSOs and the model-computed distances between their inference-fingerprints and each reference-fingerprint. The RSO on the left had a clear match, with the Starlink v1.0 reference-fingerprint being significantly closer than the other two. The RSO on the right had two close potential matches but one clear rejection – Starlink v1.5. Thus, while additional observations may be required to uniquely determine a single class, even a small collection of observations should significantly reduce the number of possible object families. We leave that study as an exercise for the reader.

6.2 Effects of Seasonality

The data used in the experiments discussed here spans one month. We expect that the fingerprint of an object, both reference and inference, will shift with seasonality. An extension on the experiments discussed to cover windows of time around the year would provide insight into how steady the fingerprinting approach is over time. A multi-year experiment would lend insight into whether the fingerprints are cyclical on an annual scale. Additionally, a series of experiments could test how long after a reference-fingerprint is created it continues to be useful for the classification of inference-fingerprints. The results from this paper compare reference-fingerprints created over three weeks of data to fingerprints from the subsequent week. Although reference-fingerprints should be consistently updatable, it would be worthwhile to explore and see how results degrade over time as the time between reference-fingerprint and inference-fingerprint data increases.

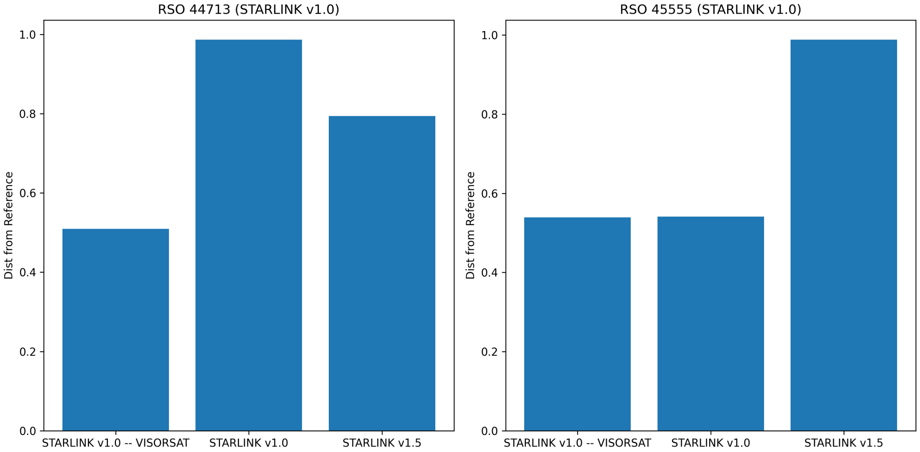

Fig. 6: Twin embedded distances to reference fingerprints (smaller distance means the fingerprints are more similar)

6.3 Applications for Attitude Change Detection

In addition to providing a method for gathering additional information about a satellite and aiding track correlation processes, the daily fingerprints show promise in applications of attitude change detection. The current fingerprinting approach operates under the assumption that a satellite is holding a consistent attitude relative to Earth, for instance to operate a nadir-pointing antenna. As the object passes overhead, ground sensors gather a number of lightcurves that can be unfolded across the fingerprint feature space. If a satellite were to adjust its attitude on or before a pass overhead in a manner that measurably affects its photometric profile, the subsequent fingerprint representation for that pass would be different from prior passes. This difference would appear as an anomaly or a shift in the inference-fingerprint, especially in the embedding space produced by a model like the twin network. Future work investigating the changes in inference-fingerprints over time, perhaps expanded to a newer and more actively maneuvering set of satellites, would help to confirm or deny this theory.

6.4 Applications for Tumbling Object Recognition

Although the methods described above apply only to objects with stable orientations, the concept can be extended to incorporate a temporal component that further captures dynamic effects such as rotation and tumbling. As in the static fingerprint cases, there are numerous potential applications ranging from object classification and characterization systems to dynamic change detection enhancements.

6.5 Parting Thoughts

The photometric fingerprinting methodology developed by Slingshot Aerospace and introduced in this paper has demonstrated effective differentiation between quasi-similar satellite bus families. The applications to broader collections of satellites and data processing pipelines are straightforward, and the potential benefits are easy to imagine. We look forward to future investigation of the extensions described above to support the growing investment in space fleets and associated data processing systems.

7. ACKNOWLEDGEMENTS

The authors are grateful to the leadership at Slingshot Aerospace for supporting this research, and to the many teammates and subject matter experts who provided valuable feedback and advice throughout the project. We very much appreciate the support from the Slingshot Global Sensor Network team who provided data access and helpful suggestions. We would also like to express our gratitude to our colleagues at DARPA and USSF as well as our project manager, Erin Pennington, for helping to refine and craft early explorations of these ideas.

REFERENCES

[1] M. Kaasalainen, J. Torppa, and K. Muinonen, “Optimization methods for asteroid lightcurve inversion: Ii. the complete inverse problem,” *Icarus*, vol. 153, no. 1, pp. 37–51, 2001.

[2] P. W. Kervin, D. Hall, M. Bolden, and J. Toth, “Phase angle: What is it good for,” in *Proceedings of the Advanced Maui Optical Space and Surveillance Technologies Conference, Maui, HI, USA*, pp. 14–17, 2010.

[3] F. Fankhauser, J. A. Tyson, and J. Askari, “Satellite optical brightness,” *The Astronomical Journal*, vol. 166, no. 2, p. 59, 2023.

[4] Y. Lu, “Brdf-based photometric modeling of leo constellation satellite from massive observations,” *Universe*, vol. 10, no. 5, p. 215, 2024.

[5] T. Payne, S. Gregory, J. Tombasco, K. Luu, and L. Durr, “Satellite Monitoring, Change Detection, and Characterization Using Non-Resolved Electro-Optical Data from a Small Aperture Telescope,” in *Advanced Maui Optical and Space Surveillance Technologies Conference* (S. Ryan, ed.), p. E50, Jan. 2007.

[6] A. Chaudhary, T. Payne, S. Gregory, and P. Dao, “Fingerprinting of non-resolved three-axis stabilized space objects using a two-facet analytical model,” in *Proceeding of AMOS 2014 Technical Conference*, 09 2011.

[7] R. Linares, M. K. Jah, J. L. Crassidis, F. A. Leve, and T. Kelecy, “Astrometric and photometric data fusion for inactive space object mass and area estimation,” *Acta Astronautica*, vol. 99, pp. 1–15, 2014.

[8] P. Dao, K. Weasenforth, J. Hollon, T. Payne, K. Kinateder, and A. Kruchten, “Machine Learning-based Stability Assessment and Change Detection for Geosynchronous Satellites,” in *The Advanced Maui Optical and Space Surveillance Technologies Conference* (S. Ryan, ed.), p. 39, Sept. 2018.

[9] Y. Huo, Z. Li, Y. Fang, and F. Zhang, “Classification for geosynchronous satellites with deep learning and multiple kernel learning,” *Appl. Opt.*, vol. 58, pp. 5830–5838, Jul 2019.

[10] R. Furfaro, R. Linares, and V. Reddy, “Space objects classification via light-curve measurements: Deep convolutional neural networks and model-based transfer learning,” 09 2018.

[11] R. Furfaro, T. Campbell, R. Linares, and V. Reddy, “Space Debris Identification and Characterization via Deep Meta-Learning,” in *First International Orbital Debris Conference*, vol. 2109 of *LPI Contributions*, p. 6123, Dec. 2019.

[12] R. Qashoa and R. Lee, “Classification of low earth orbit (leo) resident space objects’ (rso) light curves using a support vector machine (svm) and long short-term memory (lstm),” *Sensors*, vol. 23, no. 14, 2023.

[13] B. Shrive, D. Pollacco, P. Chote, J. A. Blake, B. F. Cooke, J. McCormac, R. West, R. Airey, A. MacManus, and P. Allen, “Classifying LEO satellite platforms with boosted decision trees,” *RAS Techniques and Instruments*, vol. 3, pp. 247–256, 05 2024.

[14] K. Kochis, J. Shaddix, and J. Aristoff, “All-sky electro-optical tracking of mega-constellations in low earth orbit cam key, alex ferris, max geissbuhler, joshua horwood,” 2022.

[15] Slingshot Aerospace, “Seradata.” `https://spacetrak.seradata.com/`.

[16] SpaceX, “Brightness mitigation best practices for satellite operators.” `https://api.starlink.com/public-files/BrightnessMitigationBestPracticesSatelliteOperators.pdf`, 2022.

[17] G. Koch, R. Zemel, R. Salakhutdinov, *et al.*, “Siamese neural networks for one-shot image recognition,” in *ICML deep learning workshop*, vol. 2, pp. 1–30, Lille, 2015.