Multi-Phenomenology Characterization of Space Objects Using Reinforcement Learning
Paper Preparation for the 2022 AMOS Conference Proceedings

Dr. Jorge A. O'Farrill
Modern Technology Solutions Inc.
Tracy Mims[1], Dr. Tad Janik[1], Ivan A. Fernandez[1,2]
[1]*Modern Technology Solutions Inc.,* [2]*Mississippi State University*

ABSTRACT

Advances in computing power and Artificial Intelligence have led to new areas of research with application to space situational awareness (SSA). We propose to use reinforcement learning coupled with a fast signature generator to characterize space objects by ingesting their electro-optical/infrared (EO/IR) signals and estimating their physical configuration. We will further determine if that physical configuration has changed since the last observation. We will characterize the motion and material properties of space objects by fusing unresolved infrared and visible data collected by remote sensors. This will be done by leveraging previous work performed by our team at Modern Technology Solutions Inc. (MTSI): Small Business Innovative Research (SBIR) for multi-band IR characterization of signal modulation, multiple Internal Research and Design (IRAD) and SBIR projects for faster than real-time high fidelity EO/IR data generation, an IRAD project for the use of reinforcement learning to augment test and evaluation (TE) and MTSI support of various space systems. We will show that an object's physical configuration can be estimated and that information can be used to characterize changes in that configuration. The characterization of changes based on target physical configuration should be more informative and robust than using the EO/IR signals alone.

This paper will cover the methodology and the models employed to extract shape and motion characteristics of space objects. First, we summarize our prior work in the area of model free motion estimation and synthetic optical imaging - characterization of targets using unresolved EO/IR data. This will be followed by a short discussion in the advances of custom signature prediction tools which have allowed us to perform the current work. The final section will contain results that show we can characterize targets given sufficient coverage and adequate models by employing reinforcement learning in the form of a fractal Multi-Armed Bandit which searches target configuration space to produce signals that best match observed data.

1. INTRODUCTION

Geospace and cislunar space are becoming denser and more perilous with tens of thousands of objects that threaten commercial and military satellites and launches. There are a number of organizations around the globe devoted to tracking and characterizing as many space objects as possible in order to predict future collisions with valuable space assets. An object possesses other attributes besides trajectory that are valuable for space situational awareness (SSA). In this paper we will focus on the target attributes of body motion and shape. These two attributes are connected to each other by Euler's equations and typically decoupled from the orbit for objects outside of the atmosphere. Our current research is focused on model based estimation of motion and shape but there is valuable research that is applicable to this problem that does not assume a shape model but can be used to image certain kinds of objects. We will discuss this technology first as a primer to socialize the concepts and terms we employ later.

[1]Modern Technology Solutions Inc., 360C Quality Cir NW, Suite 310, Huntsville, AL 35806
[2]Department of Computer Science and Engineering, Mississippi State University

2. PREVIOUS WORK - BACKGROUND

We now discuss our previous work in the area of target characterization through the use of unresolved electro-optical/infrared (EO/IR) data. This material is more fully covered in the References [5] [2].

2.1 MFME: Model Free Motion Estimation

We consider a general case of torque-free rigid body motion including precession or tumbling. For this we assume the target trajectory, sensors and solar positions are known. Further we do not consider flexing or deforming targets. Finally, we assume that the body of the target is axi-symmetric with its spin axis coincident with its axis of symmetry and the body has equal cross-spin normal of inertia. Since we do not make any other assumptions about the shape of the target, this approach is called a Model Free Motion Extraction (MFME) algorithm.[5]

The outputs of MFME are estimates of the target's main body axis at a referenced time with respect to the extracted target's Angular Momentum Direction (AMD) expressed through its azimuth L_{az} and elevation L_{el} in some coordinate system (e.g., ECI). The dynamic motion of the target is represented by the precession angle θ, precession frequency ω and reference phase $\phi_R(\phi(t) = \omega t + \phi_R)$ is the target's phase angle (Fig. 1). Denoting AMD by $\overline{T} = [T_x, T_y, T_z]$ the target coordinate system in ECI (3x3 matrix), we express the ECI coordinates of the main body axis B(t) = $[B_x, B_y, B_z]^T$ at time t by using a matrix-vector multiplication convention as:

$$B(t) = \overline{T} * [sin(\theta)cos(\omega t + \phi_R), sin(\theta)cos(\omega t + \phi_R), cos(\theta)]^T \quad (1)$$

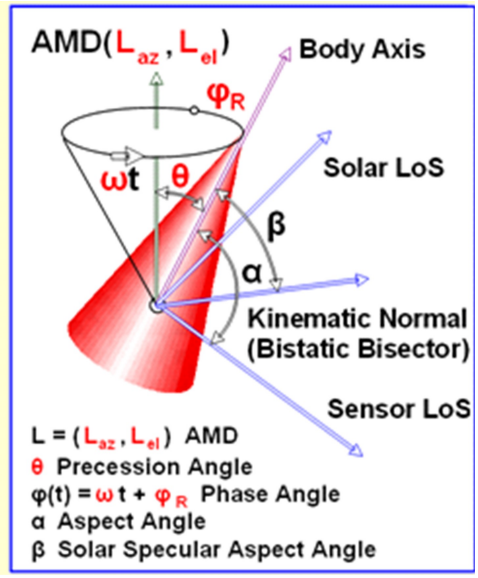

Fig. 1: Geometry of Motion

If we had known rigid body precessional motion and combined it with known sensor line-of-sight, we could have computed certain quantities such as aspect angle or aspect to the specular normal. Then, if we had taken the target signatures and mapped then in these spaces, we would have seen regular patterns. [1]
However, the measured data is not supplied in those coordinates hence the problem is to invert this procedure with the goal to recover unknown underlying signature patterns to obtain body dynamics (5 DOF) expressed in terms of above listed parameters: $L_{az}, L_{el}, \theta, \omega$, and ϕ_R.

In general, the motion of the target is extracted by simultaneously optimizing the orderly grouping of the target signature (EO/IR intensity or RF RCS) vs. a chosen variable for all the observing sensors. This variable is typically either the aspect angle (the angle between the main body axis and the sensor line of sight) or the illumination specular aspect angle (the angle between the main body axis and the body kinematic normal for visible optical signals or aspect angle for monostatic radar measurements) depending on the waveband being considered. This approach requires no model of the target but does assume that the object is executing simple free-body dynamics. Tumbling or spinning targets can be considered.

One measure of fitness for a motion solution is smoothness of the ordered signature (infrared or long wave radar as e.g., UEWR) with respect to aspect angle α. If we define smoothness as:

$$S = \sum_{ij}^{N} w_{ij}(J_i - J_j)^2 \quad (2)$$

with w_{ij} the distance weighting between points i and j ($w_{ij} = 0$ for $|\alpha_i - \alpha_j| > \delta$), J_i is the signature at α_i, then the ordering with the lowest smoothness score is the desired solution in Fig. 2. Another measure of motion solution fitness, appropriate for glint/reflection-dominated signatures (visible or high frequency radar data) is an entropy calculation.

For each hypothesized motion solution and for each sensor, we group the occurrence of glints into their appropriate bins in illumination specular aspect angle β space and calculate the entropy $E = -\sum_i p_i \ln(p_i)$, where p_i is the probability of a glint being in the i_{th} bin (Fig. 2) The entropy minimization (the most regular glint distribution with respect to illumination specular aspect angle) yields the binning of the visible or radar glints about the aspects which corresponds to the characteristic body side angles. From this approach we can extract not only the body motion but also the body shape characteristics. This methodology allows us to extract the motion solution of an axially symmetric body and infer some size/shape characteristic without assuming a target geometry. The motion solution does not include the spin of the body, however. We have been able to further refine this procedure for targets that have hot spots of surface features that produce signal variations as the body spins. This algorithm suite, called Synthetic Optical Imaging (SOPTI), will be discussed in the next section.

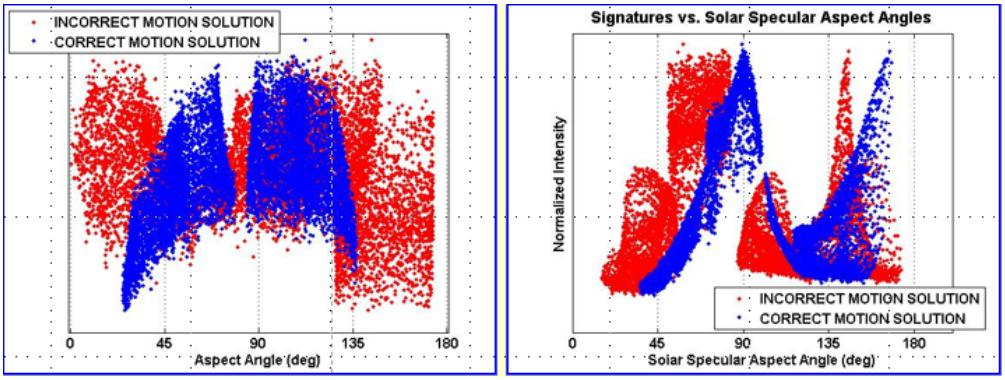

Fig. 2: Signals vs Aspect IR on Right Visible on Left

2.2 SOPTI: Synthetic Optical Imaging

This algorithm ingests unresolved visible and IR signals and outputs an image of the target. The algorithm is outlined in Fig. 3. The same assumption on rigid body dynamics apply to SOPTI as the first step of the algorithm is to generate the 5 DOF solution using MFME. This is typically done with visible signatures by finding minimum entropy of the signatures across the appropriately chosen aspect domain. The body can have any form of azimuthal asymmetry but it must posses some form so that there is signal variation from which the spin can be extracted. We've developed a regularity technique that also employs minimal entropy to extract this frequency from the IR signal. This produces a full 6 DOF motion solution. We must now extract the shape and this is done by postulating three possible shapes: cone, cylinder and sphere. Each shape is used to predict a signal using basic physics models and derive a fitness score.

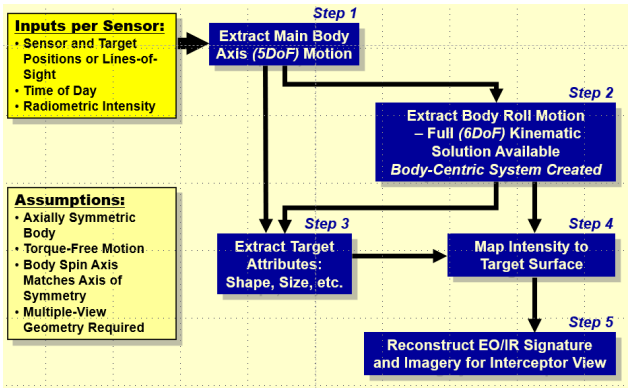

Fig. 3: Process for Synthetic Target Imaging

The final step of the algorithm is to map intensity to the surface of the body. N hot stripes are assigned to the body. These are placed in a manner consistent with the solution extracted from MFME and the shape hypothesis moments of inertia. The number of stripes is informed by the shape fitness scores. Tikhonov Regularization is used to form this final mapping. Full details of this algorithm are beyond the scope of this paper but can be found in [2].

To conclude this section, we show results of SOPTI running on the simulated data shown in Fig. 4. Four sensors are collecting on the object shown on the left of Fig. 5 and the solution produced by SOPTI is shown on the right side of that figure. The correct class of shape is chosen and the correct number of hot stripes is found. This solution is a synthetic image of the target derived from unresolved data and illustrates the power of introducing a target model. The solution required layers of regularization to overcome some of the ambiguities associated with using a generic model but the results prove the value of the method to characterize targets using unresolved EO/IR signals.

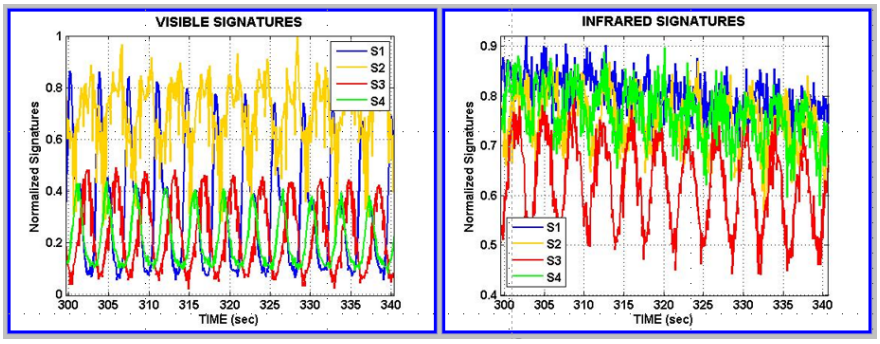

Fig. 4: EO/IR Signature Inputs to SOPTI

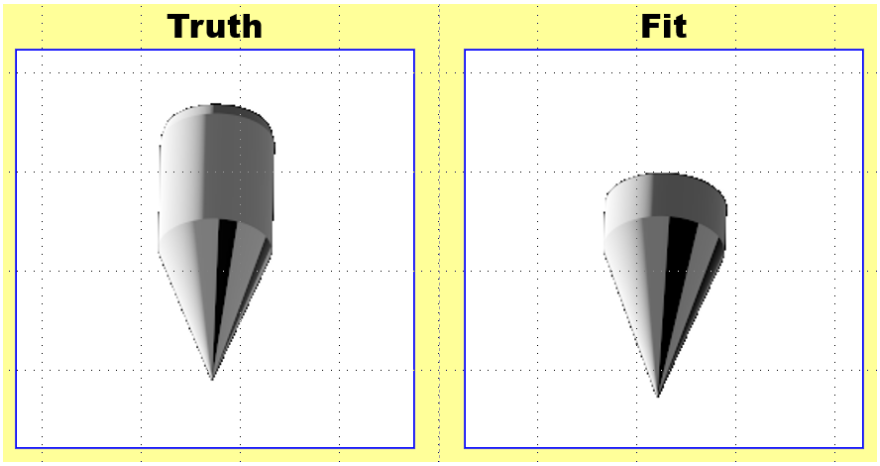

Fig. 5: Truth and Extracted Target

3. CURRENT WORK: MODEL BASED OBJECT CHARACTERIZATION

Our approach relies on the efficient generation of high-fidelity EO/IR signatures coupled with simple reinforcement learning. We use simplified motion models to show that these high fidelity tools can indeed be used to make estimates without the need for massive computer clusters or exorbitant run times. MTSI has developed a fast-running high fidelity EO/IR signature generator called HIRTSS (High Rate Thermal Solver). This tool runs at speeds of 10-1000 times faster than similar models currently employed. In support of this project alone, we ran HiRTSS to produce high fidelity EO/IR signatures over 200,000 times over the course of 4 days on two PCs. It is being used to support both threat data generation and algorithm verification for various customers. This signature solver will be used to fit the signals collected by EO/IR sensors to a class of predefined space object geometries: cubes, cylinders, and paneled configurations. We have built in the ability to morph target meshes and quickly generate new signals. This morphing can deform, scale and even wrinkle targets depending on user input. This will be used in our current effort to generate data and also as the kernel of optimizers that try to best fit signals collected by our space sensors. We have also developed an AI-powered algorithm testbed: Harnessing Artificial Intelligence for Design and Evaluation of Systems (HADES). This construct will

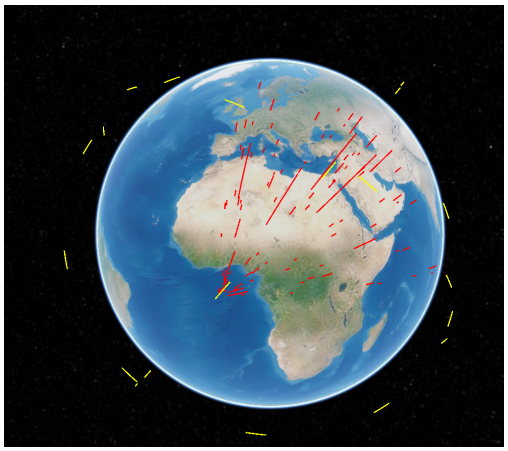

Fig. 6: Targets and Sensors

be used to perform the analysis described in this section. We will present results on three different analyses: Extraction of Motion Solution, Characterization of material properties and quality of shape extraction. We have employed and compared traditional optimizer to a basic reinforcement learner and will contrast their performance. This work is just the beginning of a potentially much larger effort that could include many more target types and enhanced hardware infrastructure for the purpose of providing greater detail on target attributes through the use of EO/IR data. The greatest advancement in this area of target characterization would be the fusion or Radar data with EO/IR data but we have not performed this work.

3.1 High Rate Thermal and Signature Solver: HiRTSS

HiRTSS is a parallelized high fidelity EO/IR signature generator. It is physics based code and operates on user defined geometries and materials. It can be used to generate signals for space objects, as shown herein, ground targets and air targets. It can operate on objects with articulating components or defects. HiRTSS is in Matlab and it still runs faster that industry standards without a loss of fidelity. Fluxes from solar and terrestrial sources are databased against the material properties and used to solve for the temperature at facets as well as reflections given observation geometry. A flow diagram of the process for generating signatures is shown in Fig. 7.

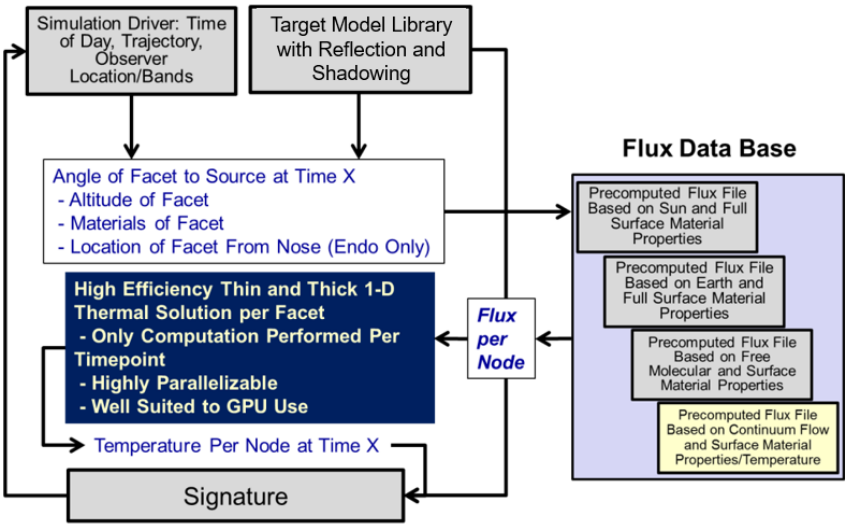

Fig. 7: Signature Generation Process in HiRTSS

3.2 Sensor Models

We are using simple sensor models for this project. Our sensors are non-imaging multi-band EO/IR sensors on satellites in a Walker constellation. They are shown in yellow in Fig. 6. The satellites have an altitude of 1200km and an inclination of 50^o, there are 3 rings with 10 satellites per ring. The sensors report EO/IR signals in the following bands: Visible: 0.4-0.7 micron, medium wave IR (MW): 3-5 micron and long wave (LW) IR : 7-9 micron. Targets are only detectable if they are within 6000km of the sensor with a line-of-sight to the sensor. This results in 1-4 sensors tracking a target for the majority of our test cases. For the purpose of this research, no noise has been added to the signals. An example of the kinds of signals HiRTSS produces for a cone using different material properties is shown in bottom panel of Fig. 8 and the effects of changing shape are shown in top panel of Fig. 8. We used simple models in this research but any model can be run through HiRTSS provided a mesh grid and materials properties are properly defined. Phenomenology manifested through employing physics based models can be exploited to characterize target geometry, materials and motion. Vis and LW are dominated by reflection and emission, respectively, this improves estimation of target at-

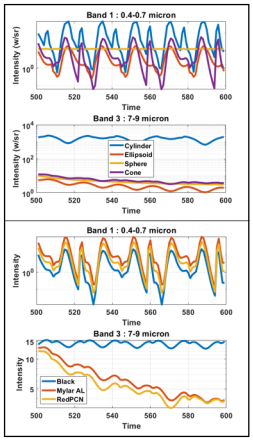

Fig. 8: Material Effects on Signatures

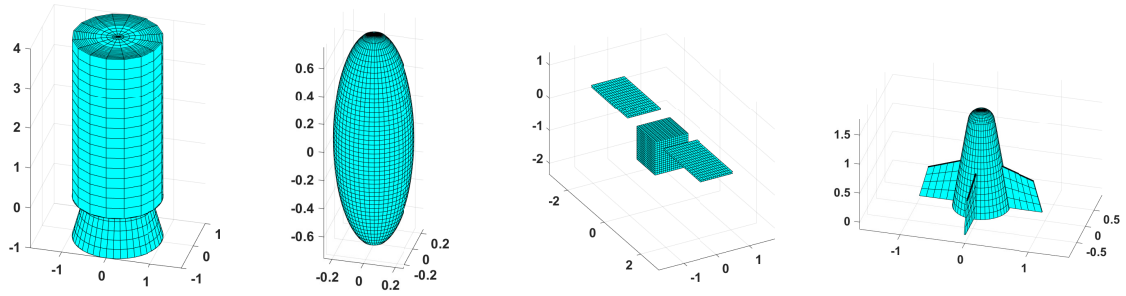

(a) Booster (b) Ellipse (c) Cube Sat (d) Cone w/ fins

Fig. 9: Target Models

tributes.

3.3 Target Models

We have created a set of 100 target trajectories, depicted in red in Fig. 6. The targets trajectories are chosen from random draws on Walker Constellation parameters: altitude $\in$ [600,1300] km, inclination $\in$ [0,90]^o, epoch start time $\in$ [0,1000] sec. We have 5 classes of targets. They are shown in 9. Note that we have used two different balloons of different materials so there are only four unique geometries. The materials for the objects are shown in Table 1.

Target	Material
Booster	Red PCN
Ellipse	Mylar, Black paint
Cube Sat	Silicon

Table 1: Targets and Materials

Target meshes and material properties couple with external fluxes and initial temperature to produce time varying signals. These signals are modulated by the motion of the target and the relative geometry of the sensors due to projected area as well as potential thermal asymmetries on the body. For motion we use a simple tumbling model where the angular momentum direction is known and we estimate the precession rate and phase. Full 5 degree of freedom (DOF) motion estimation was not explored here as we are showing a proof of principal that high fidelity physics based models can serve as the kernel of optimizers and run without the need of massive computer clusters. In this spirit we have processed 200 runs across each of the 5 target type. Each run is making at least 150 calls to HiRTSS within the optimizer. In previous work we have used different models to generate truth data but these models were too slow running to support this work. Armed with targets and sensors we can generate multi-band EO/IR signals and begin the process of estimating target attributes.

3.4 Algorithms

We employ two optimizers in this research: the simplex method [3] and a fractal multi-armed bandit (MAB) which as generalization of the typical MAB [4]. Both of these algorithm are gradient free so they can deal with mixtures of discrete and continuous spaces which are encountered in target configuration space. Our fractal MAB allows for exploration and precision by spawning a new MAB recursively within the winning bin. Unlike standard optimization techniques, the MAB can hop out of local minima during its exploration phase, and it is possible that more than one target configuration will yield viable results. This information will be valuable for the purposes of detecting changes. Fig. 10 shows this algorithm with three scales. These algorithms produce signature by operating in target configuration space: size, shape, materials, motion parameters. The ability to operate in this space separates this approach from the previous model free approaches. Both optimizers are housed within the aforementioned HADES architecture and exposed to the same data and use very

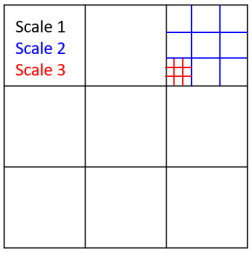

Fig. 10: Fractal MAB for Exploration and Precision

similar scores. The difference in scores is due to the fact that our MAB seeks to maximize reward and the simplex seeks to minimize the cost. The reward σ is function of the variable x, the target attribute vector which we are optimizing over. For simplex optimization we have:

$$\sigma_{simplex} = \sum_{ij}^{NM} n_j \left(1 - \frac{I_{ij}}{\sum_k^N I_k} \rho_{ij}(x)\right) \tag{3}$$

where $\rho_{ij}(x)$ is the correlation coefficient between the signal collected by the jth sensor in the ith band and the signal predicted by HiRTSS given x. I_{ij} is the average of the signature collected in the ith band by the jth sensor and n_j is the number of measurements from said sensor. N = 1 or 3 is the number of bands we are optimizing over and M is the number of sensors that collected for that particular run. The score for the MAB is similar:

$$\sigma_{MAB} = \sum_{ij}^{NM} n_j \left(\frac{I_{ij}}{\sum_k^N I_k} (0.5\rho_{ij}(x) + 0.5)\right) \tag{4}$$

This score is combined with the upper confidence bound (UCB) to form the reward for the MAB. The MAB is drawing parameters from x, this is called a pull, from a user specified set of bins. The UCB for the kth bin is given by:

$$UCB = \sum_{i}^{n} \sigma_{MAB_i}/n + c\sqrt{\log(m)/n} \tag{5}$$

where m is the number of total pulls and n is the number of times the kth bin has been visited. c is constant that is specific to the particular problem at hand, we set it to 0.25. $\sum_i^n \sigma_{MAB}^i/n$ is the average expected reward for the kth bin. Both of these rewards weigh the brightest bands de-weight sensors with few collected data points.

In order to seed these algorithms we employ a Lomb periodogram to initialize the motion solution parameters. The optimizers will return the target attribute vector x which maximizes the correlation between the EO/IR signals collected by the sensors and those predicted by HiRTSS.

3.5 Results

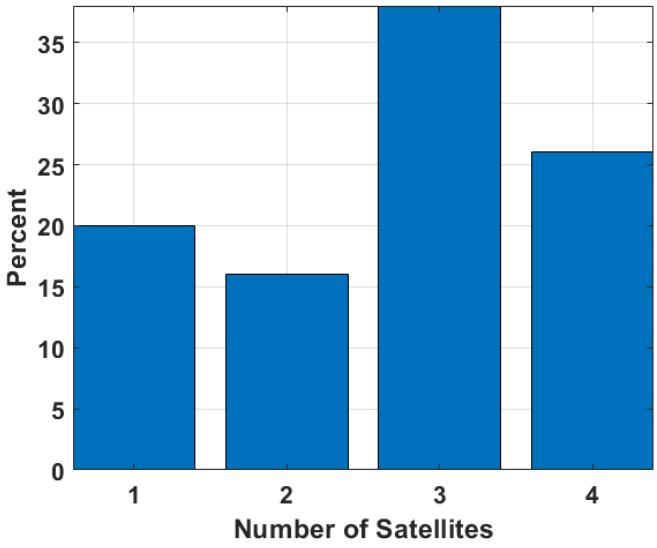

Fig. 11: Distribution of Observers

We conclude the technical section with results obtained by employing the aforementioned algorithms on data created using HiRTSS, the specified models and our space sensor constellation. Quality of fit is encoded into our rewards. It

depends on consistency across bands (if more than one band is being used) and consistency across sensors (there are typically more than one sensor viewing the target). The distribution of number of satellite observers is given in Fig. 11.

A subset of the multi-band signatures generated by HiRTSS for our cube satellite is shown in Fig. 12. The goal is to

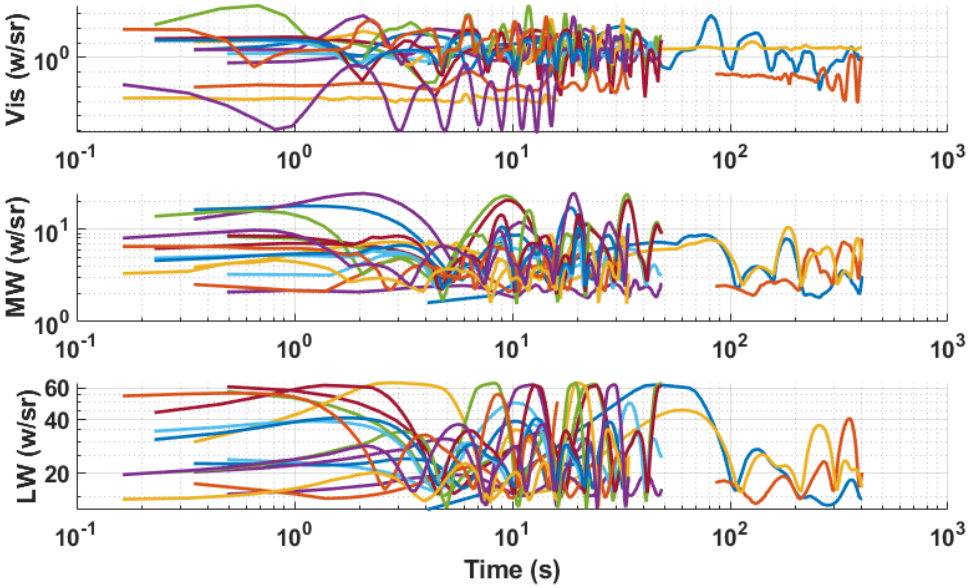

Fig. 12: Subset of Signatures Generated for Analysis

estimate the tumble rate and initial phase (ϕ_R) to best match these signals given a known geometry. We also wanted to compare the performance of the two optimizers to highlight the benefits of reinforcement learning for these kinds of problems. In the last results section we will show that when solving for motion solution one one can disambiguate between true and incorrect target hypotheses by comparing quality of fit.

3.5.1 Motion Estimation

Table 3 summarizes our results and contrast the efficacy of the fits when the optimizers are presented with either three bands of data: vis, MW, and LW or just a single band: MW

	Simplex	Simplex	MAB
Target	Wins (N=3)	Wins (N=1)	Wins (N=1)
Black Balloon	49%	54%	72%
Cube Sat	52%	52%	86%
Booster	59%	54%	90%
Cone	56%	49%	87%
Mylar Balloon	65%	56%	NA

Table 2: Motion Extraction Results

MW (N=1) data or all three bands (N=3) of data. A win in this case is a motion solution whose parameters are within 10% of the truth. For these runs the simplex method was used and one can see that adding more bands typically improves the quality of the fits. We ran the fractal MAB on most of these cases using only one band and got significantly better fits as shoen in the last column of Table 3

3.5.2 Size Estimation

For this analysis we used both the simplex and the fractal MAB to estimate four parameters: target length, target width (assuming axial symmetry), precession rate, and phase of precession. We wanted to investigate how the quality of estimated shape parameters was affected by initialization errors in the motion parameters. Fig. 13 shows the results of this analysis.

12.

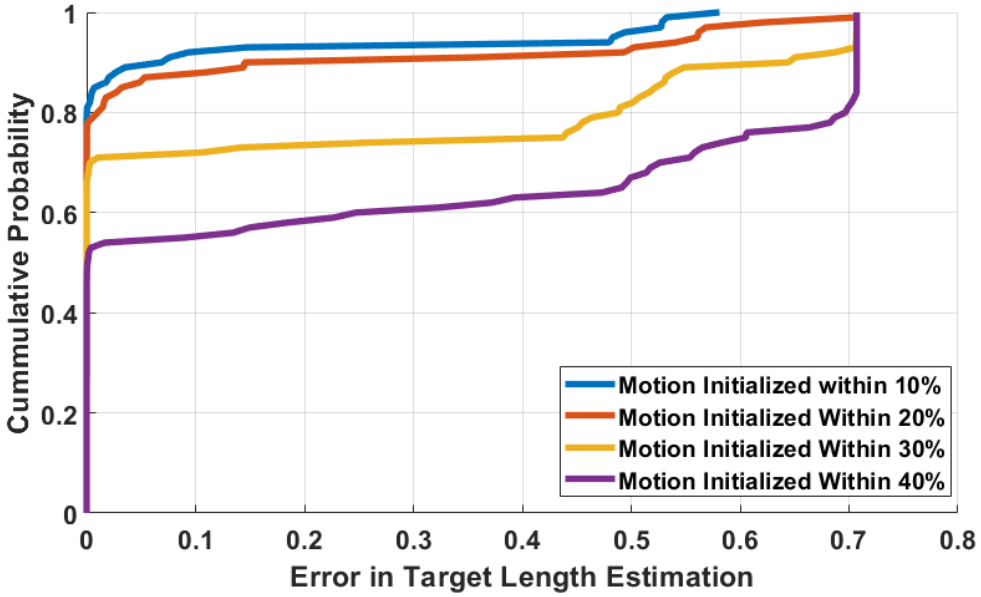

Fig. 13: Subset of Signatures Generated for Analysis

3.5.3 Multi-hypothesis Estimation: Shape and Materials

To determine the efficacy of our multi-hypothesis shape estimation we placed a booster and cube sat in the same orbit with the same motion. We attempted to fit the cube sat data using the booster model and found that 70% of the runs found the correct shape hypothesis.We repeated the same experiment but using the black balloon and the mylar balloon and found that the algorithm was able to correctly identify the proper material on 87% of the runs.

Truth Obj	TML Obj	Good Reward > Bad Reward Time Pct. (MW Only)
Mylar Ellipsoid	Black Ellipsoid	87%
Cube Sat	Booster	72%
Cube Sat	Cone	66%
Cube Sat	Black Ellipsoid	70%

Table 3: Multi-hypothesis Fitting Across Materials and Shapes

4. CONCLUSIONS AND SUMMARY

We've shown that it is feasible to use high fidelity signature tools to characterize space objects using unresolved imagery. We have also shown that the fractal MAB can achieve a much higher convergence rate than the simplex method.

We will continue to increase the fidelity of our models, generalize our morphing capability and explore other optimization algorithms produce to full 5DOF motion solutions and improve our ability to support characterization of space objects.

5. ABBREVIATIONS AND ACRONYMS

DOF - Degree of freedom
EO/IR - Electro-optical/Infrared
HADES - Harnessing Artificial Intelligence fro Design and Test of Systems
HiRTSS - High Rate Thermal and Signature Solver
IRAD - Internal Research and design
MAB - Multi-Armed Bandit
MTSI - Modern Technology Solutions Incorporated
PCN - Polyurethane Coated Nylon
SBIR - Small business innovative Research
SSA - Space Situational Awareness
TE - Test and evaluation
TML - Target Model Library
UEWR - Upgraded Early Warning Radar

6. REFERENCES

[1] D. Choi and J Williamson. Advanced modeling techniques for estimation of dynamics. *MDSEA, Monterey, CA*, 2006.

[2] Watts D. McCoy S. Hartley S. Janik, T. Synthetic optical imaging: A new approach to model-free target characterization. *7th AIAA Missile Defense Conference, Washington DC*, 2009.

[3] J. A. Reeds M. H. Wright Lagarias, J. C. and P. E. Wright. Convergence properties of the nelder-mead simplex method in low dimensions. *SIAM Journal of Optimization*, 9:112–147, 1998.

[4] A. Slivkins. Introduction to multi-armed bandits. *Foundations and Trends in Machine Learning*, 12(1-2):1–286, 2019.

[5] J. Beaupre T. Janik, D Watts. Model-free dynamics estimation of a ballistic target using optical data. *MDSEA, Monterey, CA*, 2007.