Electro-optical Signature Generation Using Physic-Aided Deep Neural Networks

Dr. Jorge O'Farrill, Nathan Highsmith
Modern Technology Solutions Incorporated
Isaac Nelson, Elise Theriot
University of Alabama

ABSTRACT

Artificial intelligence and machine learning (AI/ML) are powerful tools for efficiently performing the target characterization needed for Space Domain Awareness (SDA). For SDA, we must characterize objects as well as their intent. This characterization is typically done using remote sensing across multiple phenomenologies. In this paper we will focus on the generation of electro-optical and infrared (EO/IR) data. This is usually done with physics-based models that are computationally intensive and may necessitate the movement of large amounts of data to the ground across communication layers and through ground entry points. These paths are becoming more burdened due to the proliferation of space assets. Unlike physics-based models, AI/ML systems that can exploit hardened GPU architectures are candidates to mitigate the conflict between proliferation and robust SDA by allowing some SDA functions to move to the edge. Although AI/ML systems offer attractive solutions they can be prone to hallucinations which can lead to inaccurate results. We have attempted to mitigate these hallucinations by adapting the generation of our training data using simple reinforcement learning within our training pipeline. This data generation pipeline and the model we develop within are the focus of our paper. We will detail the data generation process, specify the model architectures and explore the model's accuracy.

We have created a physics-aided deep neural network (DNN) for generating physically accurate electro-optical and infrared (EOIR) signatures. This model predicts both thermal emissions and reflections for a given target shape and material laydown for a desired sun/earth/observer/target geometry. The model is an extension of a thermal model we trained to support our Space Force's Overhead Persistent Infrared Modeling and Analysis Center (USSF/OMAC). Our previous models were limited to encoding only thermal signatures and producing reflected signals using ray tracing. Our new model expands on this by encoding the physics and geometry of reflections for a given target. The new model consists of two specialized DNN architectures: an attention-based deep neural network for solving the underlying thermal physics of objects in space and a standard deep neural network for modeling reflective optical signatures. The attention-based DNN efficiently learns spatial and temporal thermal dependencies. This architecture is similar to that used by the transformers in popular Large Language Models (LLMs). The reflection model captures complex bidirectional reflectance distribution functions (BRDFs) and target shape to reconstruct realistic EO signatures using a deep convolutional neural network. These networks are trained using a high-speed physics-based data pipeline that uses our physics-based signature generator HiRTSS (High Rate Thermal and Signature Solver), which generates synthetic training data by simulating thermal dynamics and optical reflections under varying illumination conditions.

By leveraging intelligent data generation and state-of-the art architectures, our model is robust; it maintains accuracy across a wide range of geometric configuration. It is also very fast with speeds of 100-1000x that of HiRTSS. It could be a powerful tool for enhancing SDA including detecting changes in target configurations or classifying different classes of satellites. This class of models also has the potential to be deployed on space assets (the edge) since they can run on modest GPUs.

1. INTRODUCTION

Space Domain Awareness (SDA) requires estimating target attributes (target characterization) and inferring intent from multi-phenomenology remote sensor data. High-fidelity physics-based models are accurate but computationally

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

expensive, limiting throughput and timeliness. We propose Narrow Digital Twins (NDTs)—task-specific, physics-informed deep neural networks—as fast surrogate models that preserve the fidelity needed for confident inference while dramatically reducing compute time. NDTs can augment or replace traditional physics-based kernels for target characterization and are well suited to edge processing on modest GPUs. This is an important consideration for proliferated space architectures (e.g., Golden Dome for America-GD4A) where bandwidth is constrained.

This paper describes NDTs we have built to generate trajectories and electro-optical/infrared (EO/IR) signatures for diverse near-Earth objects. We detail the physics-based tools used to synthesize training data and the algorithm-testing environment employed to validate and harden the resulting AI/ML models. If the reader is interested in only the AI model that is the topic of this paper they can skip to section

2. PHYSICS-BASED MODELS

2.1 Introduction

Our AI/Models use a combination of deep neural networks to emulate the physical response of target models by learning high dimensional distributions connecting the environment, object and viewing geometry to observed signal. NDTs encode the outputs of physics-based models (PBMs) so they can only be as good as those models. In this section we will discuss the EO/IR signature generation PBMs used to generate the training data for the NDTs.

2.2 Signature Model: HiRTSS

Generating high-fidelity target signature data at the rates necessary to train DNNs, support large-scale simulations or perform target characterization is computationally challenging. One must account for the environment as well as target properties as shown in Figure 1. Pre-computed signature data allows for fast signature generation but it does not generalize when used in closed-loop situations where scenarios may not be part of the precomputed set. Simplified physics models lack the fidelity needed to provide confidence in inferences made using said data. To bridge the gap is throughput and maintain confidence in results, Modern Technology Solutions, Inc. (MTSI) has developed a High Rate Thermal and Signatures Solver (HiRTSS) in support of the COILS Small Business Innovative Research program effort.

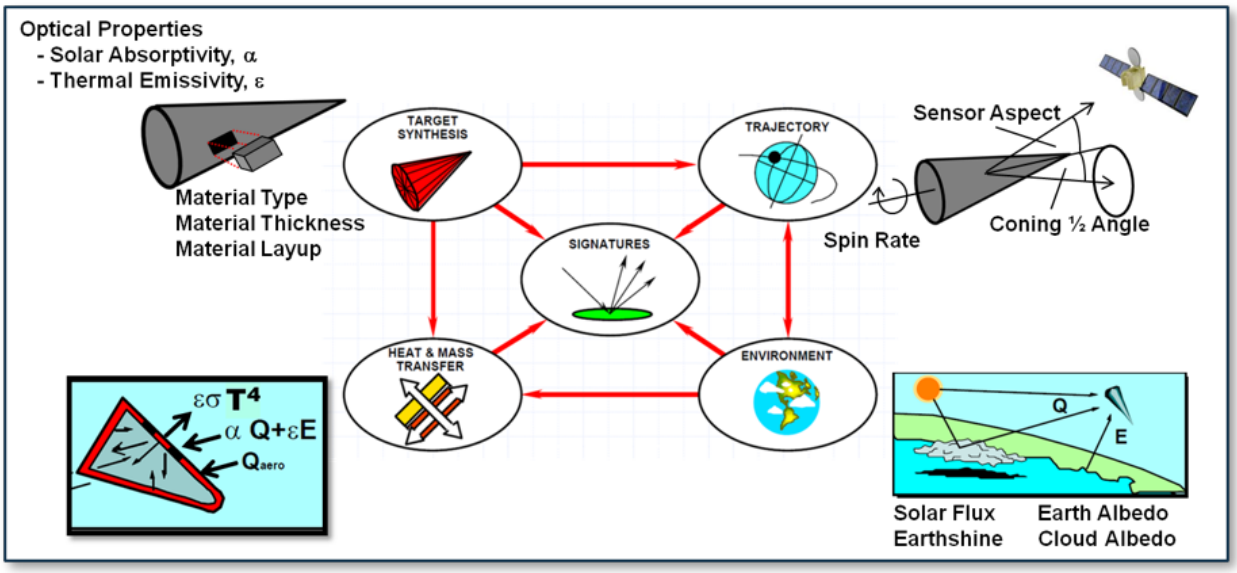

Fig. 1: Signature generation process

This software was developed to meet the needs of in-line signature generation on complex and diverse objects. This

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

tool has been used to support programs within the Missile Defense Agency (MDA) as well as USSF/SSC. The baseline features are:

1. In-line rapid Thermal computations supporting emissive signature calculations for both thermally massive and thermally light objects - Supports all dynamics types and shapes of objects
2. In-line Reflection computations for solar and earth reflections
3. In-line obscuration/shadowing computations
4. Ingesting of Missile Defense Agency Threat models into HiRTSS target model libraries to maintain rigor in object definitions
5. Ingesting of mesh files from external computer-aided design (CAD) sources: Satellites, Missiles, Aircraft, Ground and Naval Targets
6. High accuracy and high throughput, typically running at 50x real time in the current MATLAB prototype.
7. Full Monte Carlo capabilities over all major optical signature drivers including the ability to morph targets and create new objects without having to create new Target Model Libraries (TMLs)
8. The code can be run stand-alone through a graphical user interface (GUI) or command-line interface or it can be configured to run in-line within a larger simulation to provide on demand stimulation for sensor models or algorithms.
9. Morphing capabilities include:
   (a) Target surface deformation (wrinkles)
   (b) Shape and size changes (e.g. sphere to ellipse, radius increase or decrease, length increase or decrease etc.)
   (c) Orientation changes of portions of the body including rotations of parts of the target, retractions of solar panels, modeling of transient events etc.

HiRTSS achieves this throughput, not by sacrificing accuracy but by a combination of dimensional reduction of key functions and precomputations based on full spectral characteristics of the target materials and geometry. HiRTSS has been used to train and test combat ID algorithms and to support AI-driven algorithm testing and reverse engineering of remotely sensed objects in space. The HiRTSS signature solver uses a set of TMLs which contain geometric, material, environmental and reflection information. There are two paths to create these TMLs. Since HiRTSS was originally developed to complement the standard EO/IR threat signature generation tool used by the Missile Defense Agency (as well as various Intel centers), one path leverages the databases and input files that are part of the Optical Signatures Code (OSC) toolkit. This gives us access to the substantial databases of material properties and environmental models as well as direct traceability to the legacy and current approved threat models. The second path for creating TMLs involves the use of Blender.

The HiRTSS signature solver uses a set of TMLs which contain geometric, material, environmental and reflection information. There are two paths to create these TMLs. Since HiRTSS was originally developed to complement the standard EO/IR threat signature generation tool used by the Missile Defense Agency (as well as various Intel centers), one path leverages the databases and input files that are part of the Optical Signatures Code (OSC) toolkit. This gives us access to the substantial databases of material properties and environmental models as well as direct traceability to the legacy and current approved threat models. The second path to creating TMLs will be discussed in a following section. The pre-processor has the job of extracting the target model information from the Optical Signatures Code (OSC) input and geometry files and translating that into HiRTSS TMLs. There are currently 4 TMLs which contain not only the geometry and material information in the OSC input file but also the absorbed fluxes for each material, the target obscuration (shadowing) information and the full reflection computation for the object. Once these TMLs are created, then the solver can execute on any scenario without having to recompute any of these quantities.
To create the target model libraries there are several steps:

1. The target model geometry is read in from the OSC geometry file, geomb

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

2. The materials designated in the input file are assigned to the appropriate parts of the target
3. The material thermos physical and optical properties are read in from the OSC databases
4. The spectra of the earth, albedo and sun and then convolved with the reflection properties of each material as a function of wavelength and angle to create an 'absorbed flux' for each material as a function of angle
5. The angle dependent solar absorptivity and temperature dependent emissivity of each material is derived from the optical data and stored
6. The target obscuration is computed as a function of viewing angle to the body, each facet gets a viewing factor from 0 to 1 indicating its relative visibility to the source or sensor
7. Finally, the reflected signature (including obscuration) is computed for 9 million combinations of sensor and illuminator viewing to allow for on-the-fly reflection calculation in the solver

A key benefit of HiRTSS is contained here: we leverage the highest fidelity data and methods so that no accuracy is sacrificed, and the precomputations of obscuration and reflection provide dramatic increases in throughput over traditional methods.

2.2.1 Overview of Solver

Once a TML is created, the HiRTSS solver can be run using that TML for any scenario needed. A scenario would include the time of day, day of year, trajectory, sensor location and initial temperature. There are essentially two sets of computations done by HiRTSS during a run.

First, the trajectory is used to generate an incident flux history for every thermal node on the target. This flux is summed at each time point for all the sources (Earth, Albedo, Solar, internal interactions and Aerothermal) and a 1-D thermal calculation is performed for each node (a 3-D conduction model for nosetips and leading edges has been prototyped but not yet implemented). From that temperature the emissive signature is directly computed from Plank's law using the measured reflectance properties of the materials. The second set of calculations is done to interpolate the total reflected signature for the target for each time point. Since the target reflection depends on both the relative position of the sensor and observer as well as the orientation of the body, it is necessary to reduce the dimensionality of the computation to create a matrix which is compact enough for rapid interpolation and finely gridded enough for accuracy. Those time-consuming calculations are in the pre-processor so that only a one-dimensional interpolation is done within the solver proper. This process contrasts with the standard serial process where each item is computed from geometry meshing to intensity summation. Figure 2 shows the flow of typical signature generation.

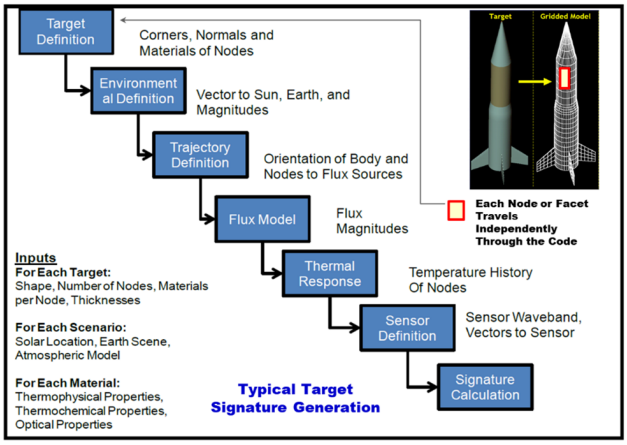

Fig. 2: Typical flow of signature generation

2.2.2 Model Validation

Since HiRTSS is designed to provide a fast-running complement to the high-fidelity tools used by the Missile Defense Agency (MDA), the initial implementation leveraged the same input files to maintain a 1-to-1 relationship with the approved threat models. As a result, validating HiRTSS has taken the form of validating the output against the high-fidelity source (typically the OSC) from which the inputs were taken. The 'validation' is essentially inherited from historical and current validation of the OSC against measured data from flight and ground testing. There are currently

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

well over 100 validation comparisons available between HiRTSS and OSC in wavebands ranging from visible to Infra-Red (IR). Many of these are classified and have been presented in various briefings across MDA. Currently, we are planning to collate all the validation examples in a separate report. Figures 3-4 are simple examples showing both thermally massive and thermally light targets being compared between the OSC and HiRTSS. These target signatures are comprised of emitted and reflected components. The differences seen are typical of those seen across the full test set (classified and unclassified).

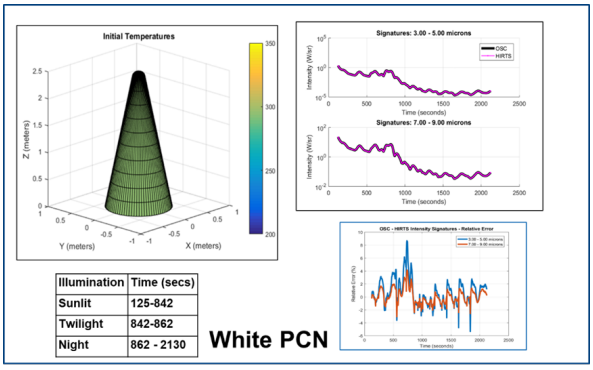

Fig. 3: Comparison for a Thermally Light Cone

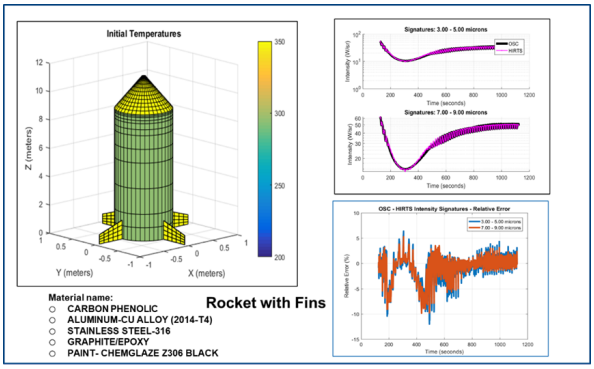

Fig. 4: Comparison for Rocket Body with Fins

HiRTSS is producing viable EO/IR signatures for a variety of targets across a broad set of conditions at speeds of 10-1000X that of OSC. This speed-up has allowed us to generate the data required to train or AI/ML models without the need for large computer clusters,

3. HADES

AI/ML models are prone to inaccuracies and hallucinations. These artifacts typically occur when the desired output is not well represented by the training data. Our NDTs attempt mitigate most hallucinations by learning fundamental target attributes from training on specific target configurations or kinematic models but they can hallucinate nonetheless. To harden our models, MTSI has developed a proven and robust enabling technology to address critical test and evaluation (T&E) engineering challenges which addresses a wide range of critical missions using a process and tool we call HADES (Harnessing AI for the Development and Evaluation of Systems).

Our approach leverages the strengths of AI and machine learning by using them within the testing element where over-training (that is learning the behavior surface of the test subject) is desired rather than an issue. We have successfully adapted this technology to several key DoD challenges. This history of success provides evidence of robustness, confidence in the process, as well as a broad set of modeling capabilities which can be used on addressing key interest areas such as: Warfare Analysis, Test and Evaluation Engineering, AI/ML, providing fast analysis and testing of 'intelligent' systems.

As shown in Fig. 5, the overall HADES structure is modular and straightforward. An Agent drives a Generator to create data specifically to explore the SUT and toxic regions within the performance domain where the SUT behaves in an unexpected or unwanted manner. The ability to identify potential emergent behavior is a core value HADES provides.

Fundamentally, HADES consists of an adversarial 'Agent' that drive the test, evaluation, or training/development process by leveraging modern ML and AI capabilities coupled with a flexible and fast generator to create new data specifically in response to the performance of the system under test (SUT). Regions of poor performance that may be opaque to humans are discovered naturally by the Agent's learning algorithms. In this way, the SUT itself drives the test/training data and regions of unexpected or poor algorithmic performance are mapped into the physical domain. Data from these regions can be used for retraining/hardening of the system or capturing as areas of concern. This fundamental mapping of the high-dimensional and abstract space of the SUT into the physical domain or the real world is a key contribution of the HADES testbed.

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

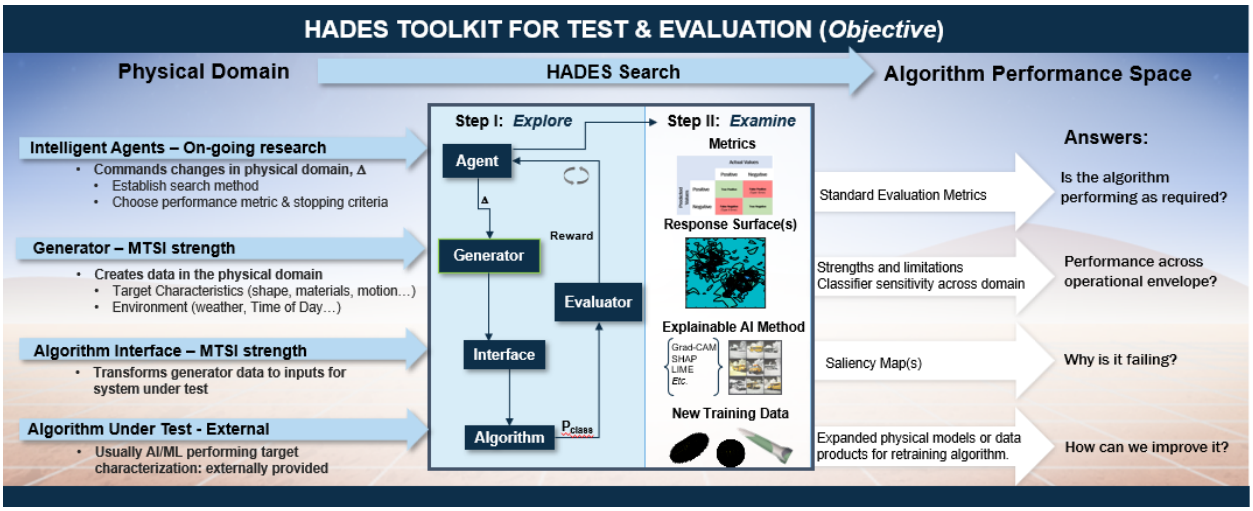

Fig. 5: Overview of the HADES testbed, functions, and benefits

HADES testbeds can be constructed for a wide range of applications if we can develop a Generator to create the high-throughput, high-quality data required to adaptively test the SUT. This points to another strength of our proposal, the generators themselves are powerful modeling and simulation tools which can be delivered as products. Figure 6 shows prior successful applications for MDA, the Space Force (SF), the Missile and Space Intelligence Center (MSIC) and Hypersonic test and evaluation.

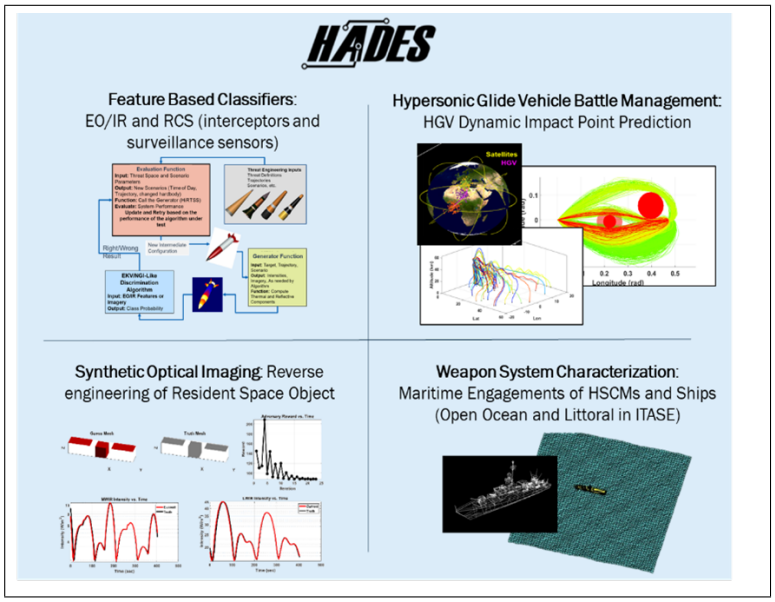

Fig. 6: Typical flow of signature generation

4. AI/ML MODELS

We have trained NDTs to learn complex physics associated with kinematics and thermodynamics of a variety of target types. Our latest network also encodes the reflective properties of objects including geometry to observer. These networks are trained on a specific set of target configurations which include: target geometry, materials and kinematics. This target specificity allows for a more complete set of training data which yields a Narrow Digital Twin that is less likely to hallucinate and can be used as a substitute for the slower running physics-based model. Since our models are encoding time series data they all exploit attention mechanism used in Transformer [1]. In this section will provide a brief overview of legacy models and their architectures along with some sample outputs. In the last section we will provide details on our latest EO/IR signature model which expands upon previous models to include DNN encoded reflections These models run 100-1000x faster than physics-bases code which makes them useful not only for generating data in real time in simulations but also as kernels in optimizers for target characterization. These fast-running kernels allow the user to search larger dimensional spaces more efficiently and are amenable with edge processing since they run on modest GPUs.

4.1 HGV Trajectory Model

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

Hypersonic glide vehicles (HGVs) have complex kinematics that couple to the attitude dynamics of the body. We have a simplified kinematic model that allows for the vehicle to maneuver through time-dependent policies - roll and angle of attack(AoA). HGV EoMs are shown in Figure 7. MTSI can also generate high-fidelity, 6DOF, trajectories. These will be used to generate a more generalized AI trajectory model in the future. Our trajectory NDT is a generative algorithm trained on a modest set of trajectories 800k with a test and validation set comprised of 150K trajectories. The model test metrics from this analysis include: MAE: 0.0082, MAPE: 8.81%, % Real Trajectories Over Dataset Bounds: 99.97%, Average Cross Range Error: 15.6 km, Average Ground Range Error: 62.5 km. The trained AI model produces HGV trajectories 1000x faster than trajectories produced with numerical integration. While there are deviations from the user-specified and predicted ground and cross range, indicates that most trajectories remain with physical bounds of energy, time, and survivability. Training took 2 days on single PC with an i9-10900X 3.GHz CPU and a NVIDIA Quadro RTX 5000 GPU.

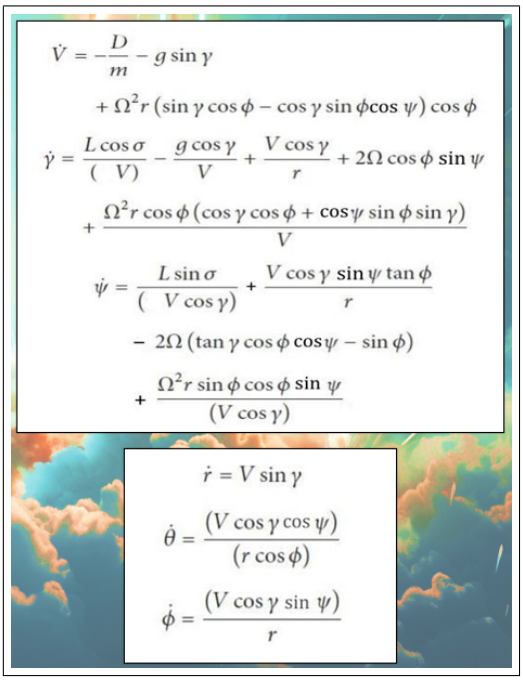

Fig. 7: HGV Kinematic model

This model is a conditional variational autoencoder with attention. The encoder is bottlenecked feed-forward network with a multi-head attention layer. The conditions allow for trajectory shaping through the specification of cross-range and down-range but is generative since a latent space coordinate must also be specified. The network learns kinematic rates (like velocity) which get integrated in a custom layer to produce the final states: position, velocity and policy. This set of outputs provides a full 6DOF solution as our vehicle does not yaw. Example trajectory outputs and mean squared error are shown in Figure 8 This NDT generates trajectories 1000x faster than numerical integration. At these speeds one can solve trajectory optimization problems through brute force. We have also recently used the latent space relationship to attitude policy parameters (roll, angle of attack) to assist in impact point prediction

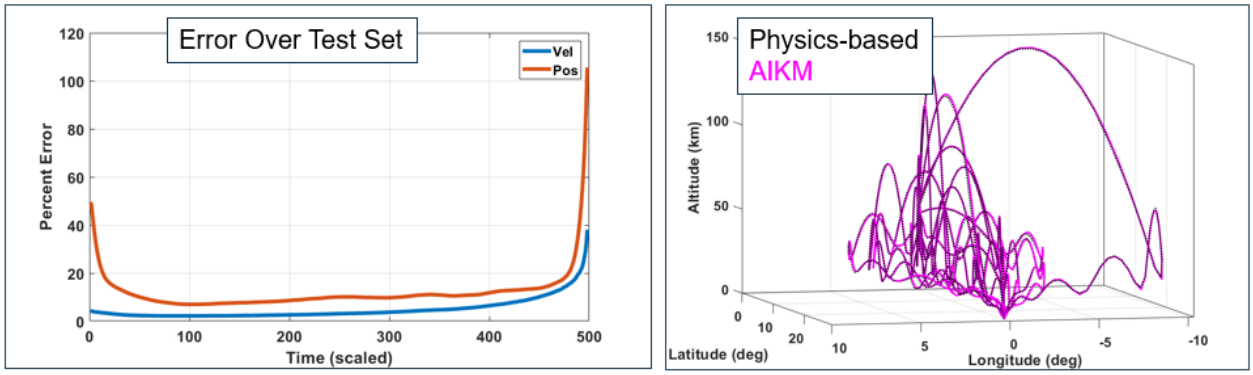

Fig. 8: Trajectory generation model

4.2 HGV Thermal Model

MTSI has trained deep neural networks to produce HGV thermal signatures in support of MDA and USSF/SSC. These models are not included in this paper but their architecture are derived from the models shown herein. The models ingest 6DOF trajectory data with associated time points and return temperature at every facet on the body for each time point. The initial model had a fixed output length of temperature states at 200 time points but we can now return a temperature at every time point specified in the trajectory data. Signatures for this model do not include any reflected radiation as the targets signatures in band of interest are dominated by thermal emissions. This model was trained using 650K trajectories with associated thermal profiles (temperatures on each of the 497 facets). This thermal data

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

was generated by HiRTSS for an HGV target model with the following characteristics: 474 total facets, body materials feature 20 mm silica phenolic over 10 mm aluminum, leading edges are 50 mm of carbon phenolic over 10 mm of aluminum. The absolute mean error on temperatures across the test set was 2.5% Training was performed on single PC with an i9-10900X 3.GHz CPU and a NVIDIA Quadro RTX 5000 GPU and took about 48 hours.

The AI Thermal model (AITM) is comprised of three DNNs: a transformer-based model (base model) and two convolutional neural networks for upsampling (upsamplers) as shown in 9. The base neural network predicts temperatures for a select number of thermal nodes (approximately 5%) along the geometry surface, while the upsamplers map the temperature to every facet on the body. The base model is based on the original/traditional transformer proposed by [1]. The base model architecture features some modifications to the traditional transformer architecture: convolutional embedding for the inputs and outputs, layer Normalization is used at the start of the transformer encoder and decoder blocks to improve stability and convergence, outputs of the convolutional encoder are put into a 1-D convolutional layer with sigmoid activation, and the architecture consists of 4 encoder blocks and 4 decoder blocks, each with multi-head attention layers consisting of 1 head.

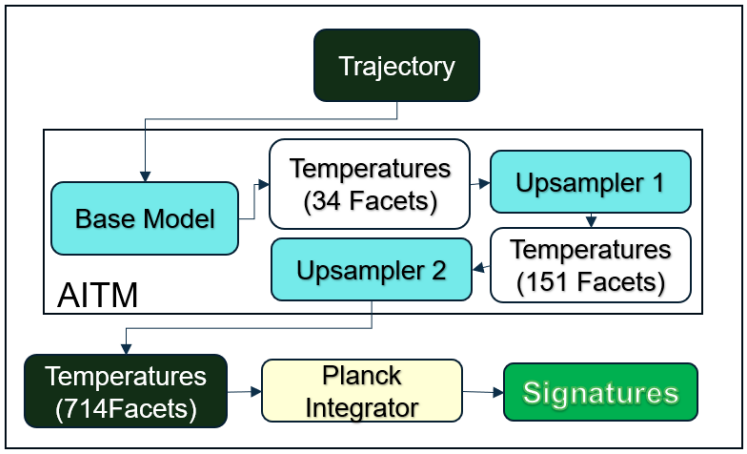

Fig. 9: Thermal model flow diagram

The base model only calculates temperatures at select thermal nodes; the upsampler model interpolates these temperatures across the remaining facets. This methodology is derived from OpenAI's revolutionary work in data generation through the development and use of the Dall-E model, which generates images from user input text [2]. According to OpenAI, their Dall-E model generates small images and then upsamples those images using another AI model. This cuts down on model sizes and potentially improves performance as there is less burden on the base model to produce a dense feature space.

The upsampler model we use differs from that proposed by OpenAI in that, instead of using a diffusion model, we use a Hybrid Attention Transformer [3]. Hybrid Attention Transformers extract shallow features from data using convolutional layers, which are then into multiple transformer blocks to extract deep features. An additive layer is then used to combine the shallow and deep features. The added features are then fed into a convolutional decoder to retrieve upsampled data. Convolutional embedding is used on the inputs to the transformer encoder for data reduction during

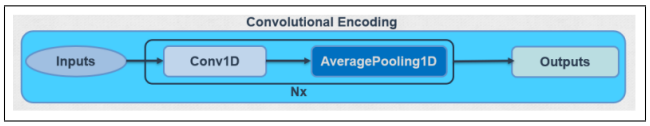

Fig. 10: AI Reflectance Models Breakdown

model training and inference. A series of convolutional layers and pooling/upsampling layers are used to "encode" and "decode" input sequences, which are then put through 1-D convolutional layers and 1-D average pooling layers until key features are extracted from the input data. Convolutional decoding is performed by running decoder inputs through 1-D transposed convolutional layers followed by 1-D upsampling layers.

Once the target facets have been assigned temperatures, we extract signatures using viewing geometry to occult facets hidden from the observer and Planck's law with the assumption each facet is a Lambertian surface to integrate the signature in the band of interest. In **??** the two left panels show the temperature on each facet for each time for the AITM and HiRTSS. The panel to the right shows the percent signature difference over the test set. This model is

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

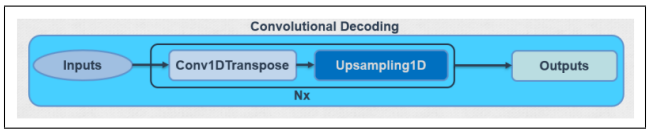

Fig. 11: AI Reflectance Models Breakdown

sufficient for a target whose observed signature is dominated by thermal emission. Most space objects have signatures that are dominated by solar radiation or earth shine for the mid IR and shorter wavelengths. The next model extends the thermal model to include the effects of reflection and absorption of solar/earth radiation on observed EO/IR signal.

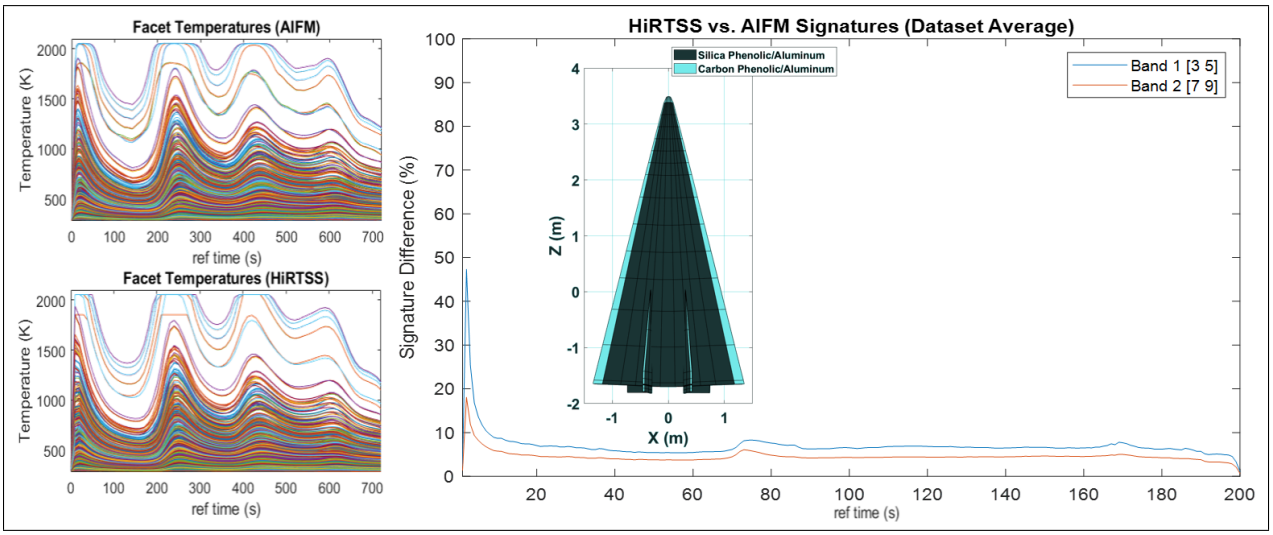

Fig. 12: Thermal model results and target geometry

4.3 Satellite Signature Model

The satellite AI/ML model extends our thermal model with the inclusion of a physics-based reflection component. This reflection algorithm slows down this model but it produces accurate reflected signatures. The thermal model is now required to learn the thermal properties of heating via radiation versus the aero heating. Since the sun is a dominant source of radiative heating, we also include the location of the sun (and earth) as an input to the model.

4.3.1 Thermal Emission Model

This satellite model uses the same architecture to encode the thermal response of the target in a typical near-earth space environment. We trained three satellite models: cube, cube w/panels and cylinder with panels. The model was then used to classify a targets physical geometry using only light curves collected by satellites equipped with EO/IR sensors. Full details of the architecture and associated analysis can be found in [4]

The resultant temperature estimates had a MAE of 3% over the 150k test data samples and trained in 40 hours on the hardware mentioned previously. This model is only 2-10x faster than HiRTSS due to the reflection algorithm employed. This modest speed up does not allow for efficient extraction of target attributes and is not amenable to edge processing. In order to make this class of algorithms faster we have encoded reflections. This fully encoded model still relies on Planck's Law in order to convert temperatures to signatures but provides a significant speed up when coupled with the model described in the next section.

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

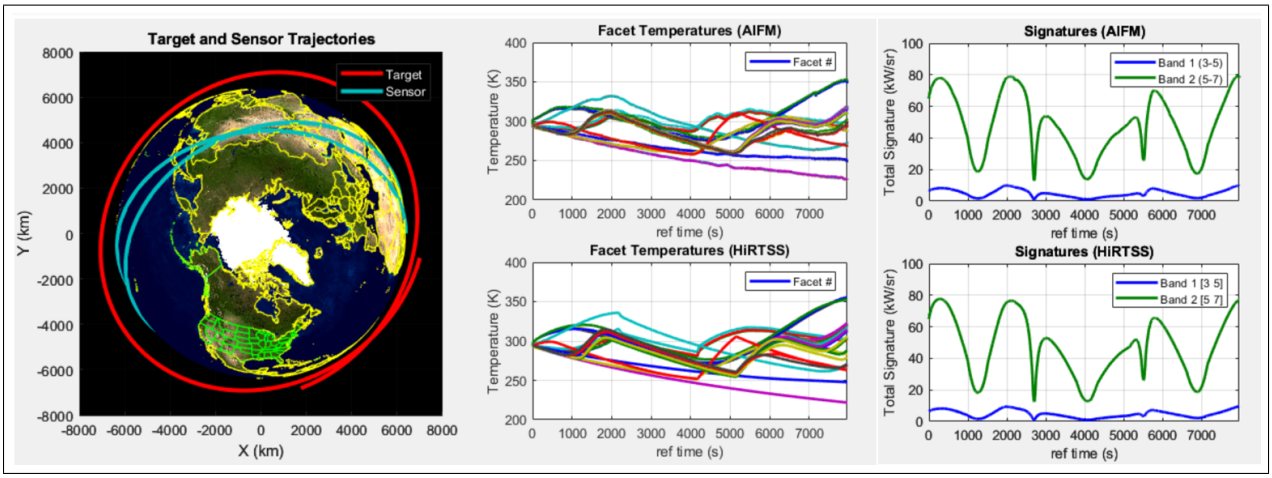

Fig. 13: EO/IR Transformer model with reflection

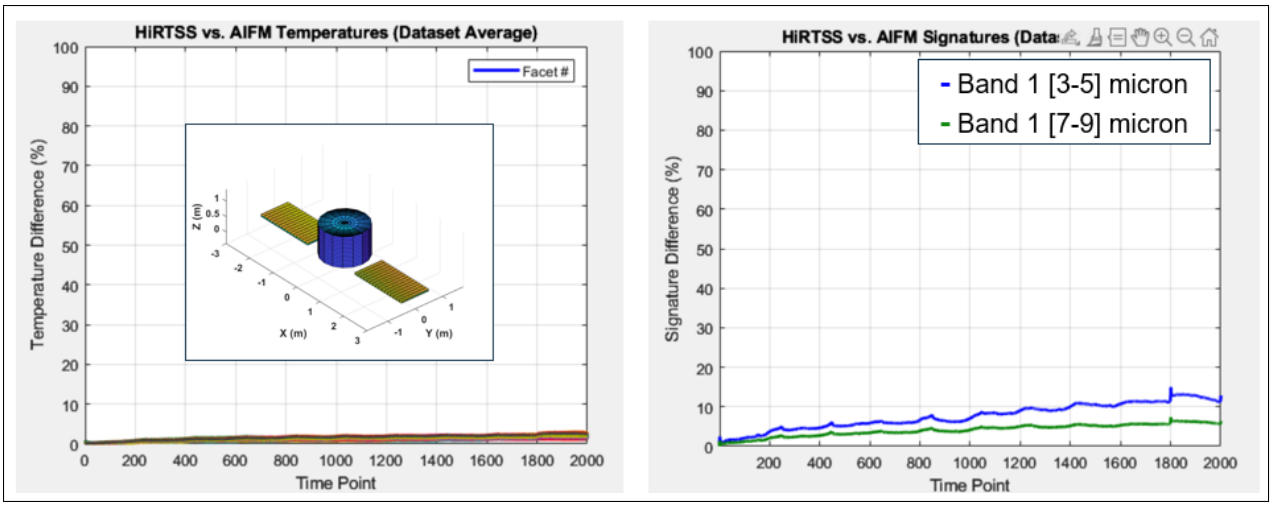

Fig. 14: Model Results and target geometry

4.4 Reflection Model

In order to achieve even higher throughput for data generation, we encoded the reflections of a cone as well. We have trained an AI model to also predict the reflections of a basic cone with exaggerated fins that was simulated to fly along a ballistic trajectory.

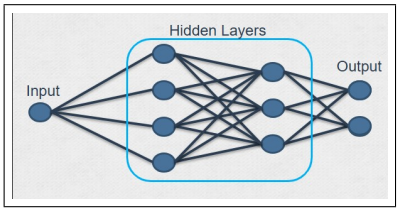

Fig. 16: DNN Architecture

The reflectance was decomposed into 3 different types: albedo, earthshine, and solar. These reflection types were also further decomposed into 3 values: signature, energy scaler, and transmit scaler. 9 total networks are trained and their outputs combined so that signatures for any specified waveband can be calculated given a trajectory, directional cosine matrix, and time (TALO and time of day).

A simple feed-forward neural network was trained for each signature decomposition variable. Each network is small and lightweight, so it is quick to train, load, and run inference. A log transformation is also applied to the signatures and energy scalar, and a custom loss is used that emphasizes losses for higher truth values:

$$|y_true - y_pred| * (y_true + \alpha) \quad (1)$$

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

where α represents a tunable parameter to help the model converge on an optimal solution.

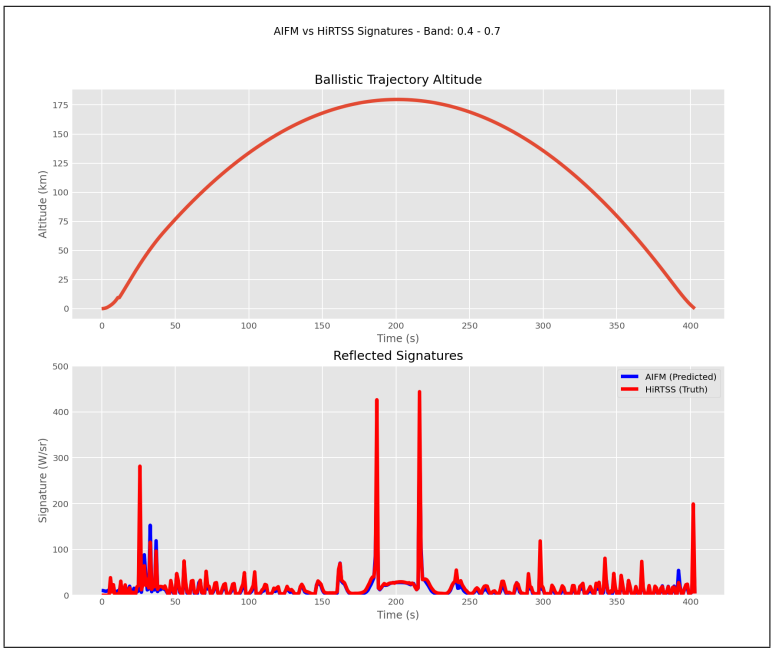

Fig. 17: Model Results: Reflected Signatures

Training for the reflectance model involves spanning all possible observer and signature source positions in a sphere divided into N three-dimensional slices. Other variables that play a role in albedo, solar, and earthshine signatures are also included as input to the model, such as waveband information and altitude. Not only are our trained AI models able to accurately reproduce reflective signatures, but they are able to do it at a much higher rate than our physics-based models. For 1000 trajectories, our physics-based models produced the corresponding signatures in 1.5 hours while the reflectance AI models produced them in 5 minutes.

4.5 Combining HGV Thermal/Reflective Models

By utilizing the methodology outlined in section 4.4 to create a reflection model for an HGV, and combining that reflection model with a thermal model, we are able to achieve an AI functional model that computes the total signature of an HGV object. The flexibility of the model to be able to predict a signature for any given waveband is important as signature waveband specificity is a bottleneck for data generation in this field. For an HGV, given a trajectory, directional cosine matrix, and a time, a reflected and emitted a total signature comprising reflective and emitted signatures can be calculated. An outline of this process is shown in 18.

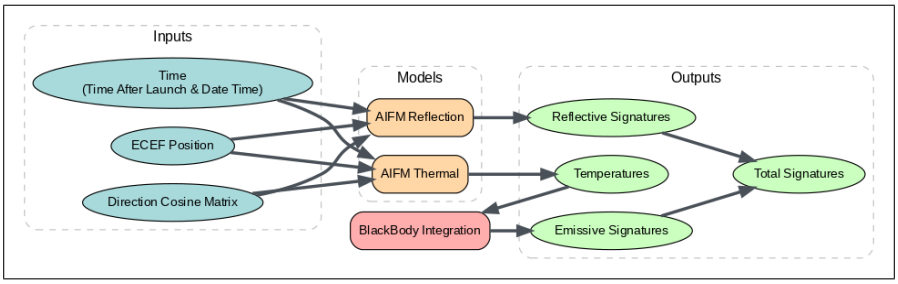

Fig. 18: Outline of Combined Model

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

This thermal/reflective AI signature model is our latest model and showcases the ability of AI to learn narrow domain target configuration.

5. RESULTS

In order to test our latest Thermal/Reflective AI models, we generated 100 HGV trajectories and corresponding signatures using our in-house physics-based simulation code. Using the AI models we were able to achieve a mean error percentage rate of 30.7% for waveband 1 (0.4 - 0.7), 1.37% for waveband 2 (3 - 5), and 0.38% for waveband 3 (7 - 9). In addition, the AI models were able to produce the data 10-100x faster than our in-house physics-based simulation code (number depends on AI hardware and use of multiprocessing for physics-based code). Figures 19, 20, and 21 detail case-by-case example of given trajectories and their signatures in multiple bands.

Fig. 19: Comparison for a Thermally Light Cone

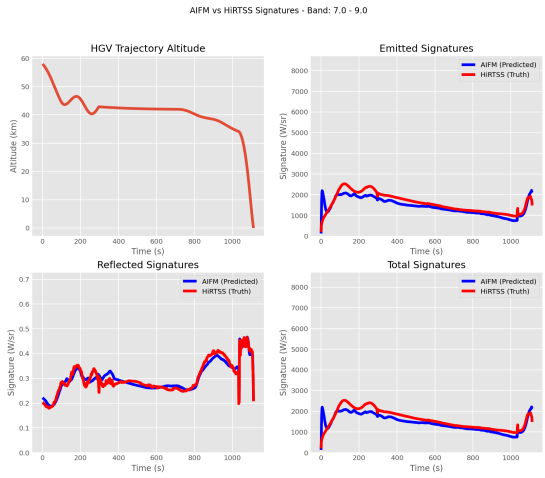

Fig. 20: Comparison for Rocket Body with Fins

6. CONCLUSION

This paper demonstrated that Narrow Digital Twins (NDTs)—task-specific, physics-informed neural networks—can deliver high-fidelity surrogates for key Space Domain Awareness (SDA) workloads at orders-of-magnitude lower computational cost. We coupled fast, traceable physics-based signature generation (HiRTSS) with targeted AI/ML architectures and a hardened test-and-evaluation harness (HADES) to produce NDTs for trajectory generation and EO/IR signature prediction that are suitable for both large-scale simulation and edge deployment.

On kinematics, our conditional VAE with attention generated hypersonic glide vehicle (HGV) trajectories $\sim$1000 $\times$ faster than numerical integration while maintaining physical plausibility across a broad envelope (e.g., average ground- and cross-range errors within tens of kilometers over $\sim$800k training and 150k test trajectories). On signatures, the thermal NDT achieved low absolute temperature error (e.g., $\sim$2.5–3% MAE across test sets) and, when combined with a learned reflectance model, produced total EO/IR signatures with mean error of 30.7% (0.4–0.7 μm), 1.37% (3–5 μm), and 0.38% (7–9 μm) while running 10–100$\times$ faster than our in-house physics-based code. These results indicate that NDTs can replace or augment slower kernels in optimization and characterization pipelines, enabling wider, deeper searches over target-attribute spaces and timelier inference under bandwidth and compute constraints.

HADES closed the loop between model development and evaluation by automatically driving test generation toward toxic and emergent-behavior regions, improving robustness and revealing failure modes that are difficult to expose with static datasets. Together, HiRTSS, NDTs, and HADES form a practical pathway to deployable SDA analytics: traceable training data, accelerated inference, and principled hardening.

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

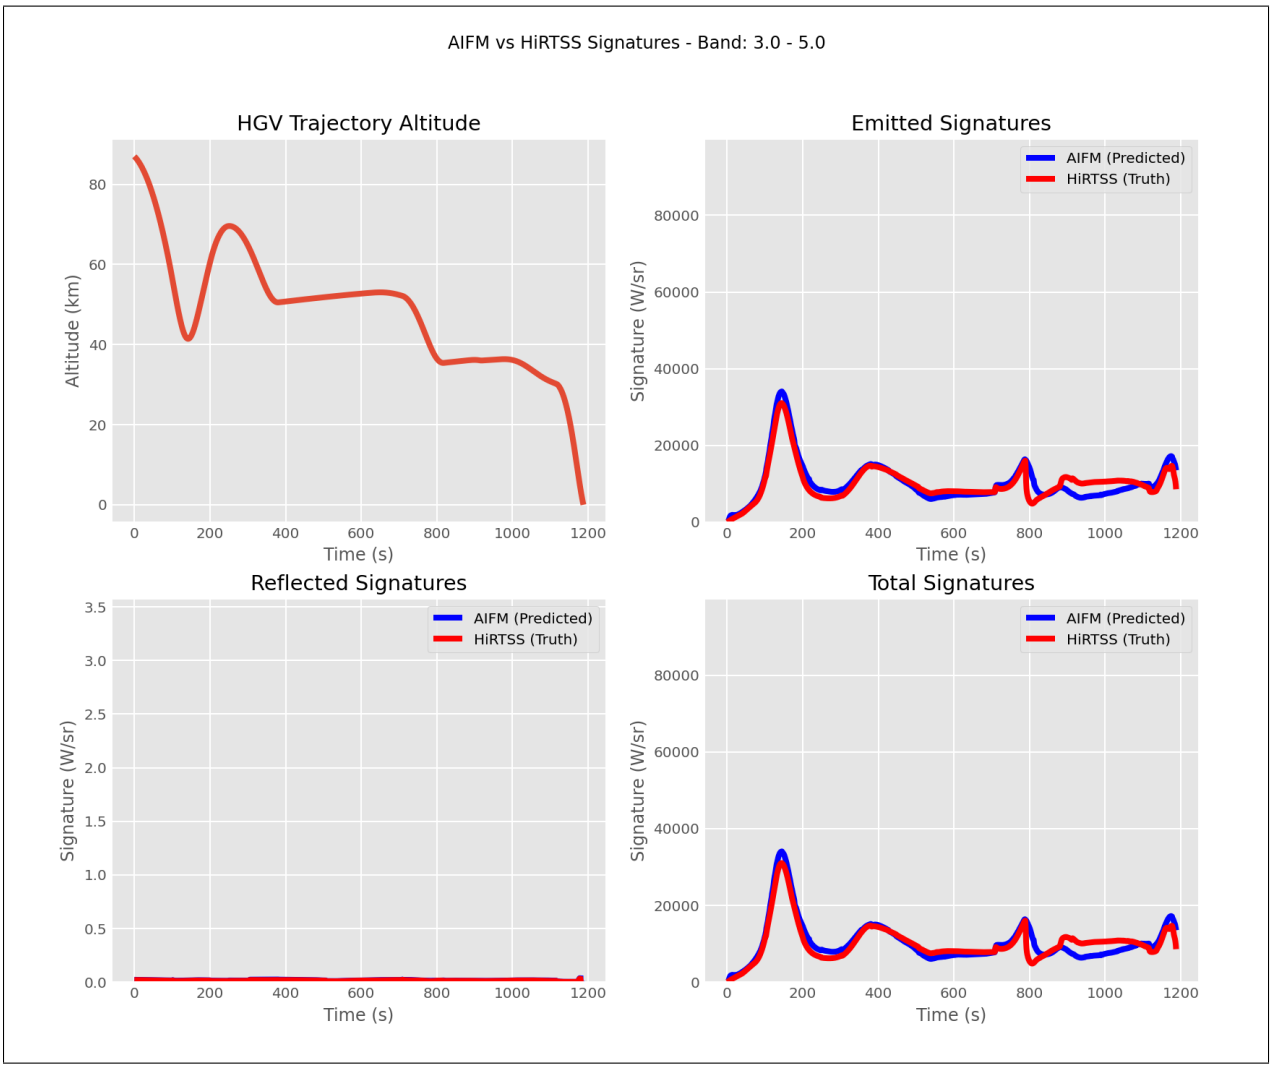

Fig. 21: Model Results and target geometry

Limitations and Future Work. Current NDTs are narrow by design: performance degrades when operating far outside the trained target configurations, material sets, or illumination/view geometries. Fidelity is bounded by fidelity of the physics based models and measured databases used to generate training data. Near-term work will (i) expand geometry/material coverage and incorporate articulated appendages (ii) extend reflectance encoders beyond baseline shapes to full spacecraft classes (iii) validate against on-orbit collections where available; and (iv) package models for embedded execution on modest GPUs with HADES-driven active learning to sustain performance after deployment.

In summary, physics-informed NDTs trained with HiRTSS and hardened with HADES preserve essential physical traceability while delivering the throughput demanded by proliferated space architectures, enabling confident, timely target characterization and intent inference at scale.

7. REFERENCES

REFERENCES

[1] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30.

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

Curran Associates, Inc., 2017.
[2] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision, 2021.
[3] Xiangyu Chen, Xintao Wang, Wenlong Zhang, Xiangtao Kong, Yu Qiao, Jiantao Zhou, and Chao Dong. Hat: Hybrid attention transformer for image restoration, 2024.
[4] Nathan Highsmith and Jorge O'Farrill. Transformer models for efficient eo/ir signature generation. Briefing at the AMOS Conference, 2024.

8. PUBLIC RELEASE CLEARANCE

This paper can be released to the public.

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com