EMMA: Ephemeris-Assisted LEO Inter-Satellite Link Spoof Detection Dataset and ML Framework

Varun Kohli, *Graduate Student Member, IEEE*, Arijit Bhattacharjee, *Member, IEEE*, Samar Shailendra, *Senior Member, IEEE*, and Biplab Sikdar, *Fellow, IEEE*

***Abstract*— Inter-satellite link (ISL) spoofing in Low Earth Orbit (LEO) constellations poses a growing yet underexplored security challenge. This letter presents *EMMA*, a lightweight Multi-task Learning (MTL)-based framework for ISL authentication and spoof detection using simple physical behavior-based features. It also presents the first benchmark simulation dataset on ISL security. *EMMA* achieves detection rates of 94.75% and 94.46% on registered and spoofing satellite samples, respectively, at low inference latency (10^{-4} seconds) and low memory overhead (3.5 KB). We also compare *EMMA* with existing works to demonstrate its comparative effectiveness.**

***Index Terms*— Non-terrestrial networks (NTN), inter-satellite links (ISL), multi-task learning (MTL), authentication, spoof detection, 5G/6G.**

I. Introduction

THE evolution of communication networks toward fifth-generation (5G) standards and beyond has generated unprecedented demand for improved connectivity. Specifically, the need for networks to facilitate ubiquitous and reliable connectivity while supporting high data rates at near-zero latency presents a significant challenge. To address these requirements, the integration of non-terrestrial networks (NTNs) with existing terrestrial networks (TNs) to form a vertical heterogeneous network (VHetNet) has emerged as a promising solution [1]. This trend has been further fueled by recent advances in satellite launch platforms and the increasing commercialization of satellite-based services, followed by the latest 3GPP 23.700-19 standard for NTN-TN integration.

Commercial VHetNets generally comprise Low Earth Orbit (LEO) satellite constellations to complement legacy infrastructure and cover remote geographic locations. They are typically managed by a terrestrial operator using a family of geographically distributed Ground Stations (GSs) connected to the core network (Fig. 1a). The open nature of Free Space Optical (FSO) and Radio Frequency (RF)-based Inter-satellite Links (ISL) makes them vulnerable to threats such as spoofing, in which adversaries impersonate legitimate nodes to compromise communication integrity. While there have not yet been reported cases of LEO-ISL spoofing as of the date of this publication, the increasing density of LEO satellites (expected to reach 100,000 by 2030 [2]) will make it increasingly feasible. Further, LEO-ISL spoofing is an important consideration for national defense measures against national-level threats that are agnostic to resource availability for such operations. Therefore, the need for robust ISL authentication and spoof detection measures in real-time is increasing. There has been limited work done in this area. For these reasons, this letter focuses on ISL spoof detection.

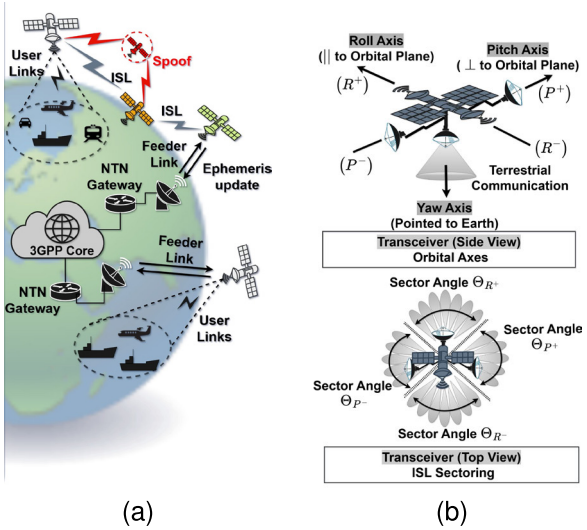

Fig. 1. NTN-TN VHetNet architecture: (a) ISL connectivity between LEO satellites in presence of spoofing satellite(s); (b) Satellite orbital axes and ISL sectoring.

In conventional wireless networks, spoofing attacks and access control regulation are managed using cryptographic and application-layer authentication solutions. However, such approaches incur high latency in VHetNets, owing to their complex key generation, distribution, and management methods, making them less effective. Consequently, various physical-layer attributes such as Channel State Information (CSI) [3], doppler shift [1], [4], [5], [6] received power [1], [5], [6], Carrier-to-noise Ratio (CNR) [7], and hardware characteristics-based signal impairments [8], have been used to design authentication solutions. Table I compiles the related work for the reader's reference. Global Navigation Satellite System (GNSS) spoofing [9] and satellite-GS spoofing [3], [6], [7], [8], [10] have been widely studied in the literature with various publicly available real-world datasets and are not the focus of this letter. Furthermore, the limited work on

Received 18 February 2026; accepted 11 March 2026. Date of publication 17 March 2026; date of current version 20 March 2026. The associate editor coordinating the review of this article and approving it for publication was N. Saeed. *(Corresponding author: Varun Kohli.)*

Varun Kohli, Arijit Bhattacharjee, and Biplab Sikdar are with the Department of Electrical and Computer Engineering, National University of Singapore, Singapore 117417 (e-mail: varun.kohli@u.nus.edu; b.arijit@nus.edu.sg; bsikdar@nus.edu.sg).

Samar Shailendra is with the School of IT and Engineering, Melbourne Institute of Technology, Melbourne, Australia (e-mail: sshailendra@mit.edu.au).

TABLE I

QUALITATIVE COMPARISON OF AUTHENTICATION APPROACHES IN VHETNET ARCHITECTURE

Ref	Link	Fingerprint feature	Method	Spoof detection	Dataset
[3]	Sat-GS	Raw satellite signal components	ML	Yes	NA
[6]	Sat-GS	Doppler shift, received power	ML	Yes	NA
[7]	Sat-GS	Carrier-to-noise ratio	ML	Yes	NA
[8]	Sat-GS	Transmitter characteristics	ML	Yes	NA
[1]	ISL	Doppler shift, received power	ML	No	No
[4]	ISL	Doppler shift	Crypto	No	No
[5]	ISL	Doppler shift, received power	Stats	Yes	No
EMMA	LEO-ISL	Latency, doppler shift, sector, azimuth, elevation, range	ML	Yes	Yes

ISL authentication [1], [4], [5] does not cover spoofing or is limited by complexity. Finally, there is no available simulation or real-world benchmark dataset on ISL security.

To overcome these limitations, we present the first ISL security benchmark dataset and propose a lightweight, Multi-task Learning (MTL)-based authentication framework. The proposed approach employs a single feature vector (consisting of observed features of the ISL request and telemetry ephemeris parameters of the claimed identity) for instantaneous inference. This letter makes the following contributions:

- A novel, low-latency, physical behavior-based instantaneous ISL authentication and spoof detection framework called **E**phemeris-assisted **M**ulti-task **M**achine-learning **A**uthentication (*EMMA*) is proposed to ensure secure ISL connections for latency-sensitive applications.
- A lightweight Multi-task Learning (MTL) model is proposed to ensure high authentication and spoof detection accuracy while using low latency and memory overheads.
- The first benchmark dataset for ISL authentication and spoof detection is generated through a three-day simulation campaign of a real-world 60-satellite Iridium constellation and is publicly available on IEEE Dataport [11] to support future research efforts.

II. BACKGROUND

This section provides a brief background of the NTN-TN VHetNet architecture, feature extraction, and MTL to help readers better understand the proposed framework.

A. Architecture of NTN-TN VHetNet

We consider a VHetNet architecture as shown in Fig. 1a. It consists of a family of LEO satellites circling the Earth in specific paths along the orbital plane. The satellites interact with the terrestrial 3GPP core network (TN) through a series of ground stations (GSs) using *feeder links*, and with their space-domain peers using RF-based ISLs following the Free Space Path Loss (FSPL) model. All links are managed and regulated by an NTN operator in coordination with the TN operator to provide a unified VHetNet coverage.

To facilitate communication between satellites or between a satellite and the TN entities (users, GSs), a satellite relies on its onboard transceivers. As shown in Fig. 1b, a LEO satellite typically has five transceivers. Among these, one of the transceivers is dedicated to communicating with the TN

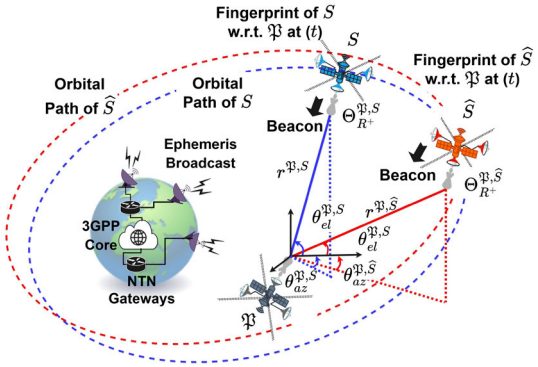

Fig. 2. Features extraction for ISL ingerprinting.

(i.e., feeder and user links), while the remaining four are reserved for ISL communication. The transceiver dedicated to TN communication points along the *Yaw* axis toward the Earth, while the remaining ISL transceivers are aligned with the *Pitch* and *Roll* axes. The roll axis corresponds to the direction of motion of the satellite, while the pitch axis is perpendicular to it. Denoting the positive and negative sides of the pitch and the roll axes using $\{P^+, P^-\}$ and $\{R^+, R^-\}$ respectively, the total coverage angle of each ISL axial transceiver is denoted using $\{\Theta_{P^+}, \Theta_{P^-}, \Theta_{R^+}, \Theta_{R^-}\}$.

B. Extraction of Satellite Features

In this letter, we consider signal latency ($\mathfrak{L}$), doppler shift (f_Δ), sector (Θ), azimuth (θ_{az}), elevation (θ_{el}), and range (r) as ISL fingerprint features due to their simplicity and ease of accessibility. Fig. 2 depicts the feature extraction process. An arbitrary satellite denoted as $\mathfrak{P}$ serves as a *pivot* node, and two neighboring satellites, S and $\widehat{S}$, following slightly different trajectories represent a genuine neighbor and a spoofer, respectively. To launch an ISL spoofing attack on $\mathfrak{P}$, $\widehat{S}$ claims to be the registered neighbor S.

Estimating Θ: The axial transponder at $\mathfrak{P}$ determines the received signal's sector using the signal strength ($X^{\mathfrak{P},S}(t)$) over each sampling period.

Estimating f_Δ: Doppler shift estimation is crucial to compensating for carrier frequency shift and facilitating data exchanges, and is typically performed on satellite transceivers for reliable communication and navigation services [4]. The position ($\rho_S(t)$) and velocity ($V_S(t)$) vectors of S at sampling instant t are used to calculate doppler shift w.r.t. $\mathfrak{P}$ as follows:

$$f_\Delta^{\mathfrak{P},S}(t) = \frac{(V_S(t) - V_{\mathfrak{P}}(t))^T (\rho_S(t) - \rho_{\mathfrak{P}}(t))}{\|\rho_S(t) - \rho_{\mathfrak{P}}(t)\|}. \quad (1)$$

Estimating $\mathfrak{L}$: We estimate $\mathfrak{L}$ evaluated using Two-way Time-transfer (TWTT) with time-stamped messages. Given that the distantly located neighbour satellites are generally unsynchronized with each other, the TWTT approach facilitates the exchange of time-stamped ranging messages between the pivot and the neighbours to estimate the instantaneous latency $\mathfrak{L}^{\mathfrak{P},S}(t)$ as follows:

$$\mathfrak{L}^{\mathfrak{P},S}(t) = \frac{(\Gamma_4 - \Gamma_1) - (\Gamma_3 - \Gamma_2)}{2}, \quad (2)$$

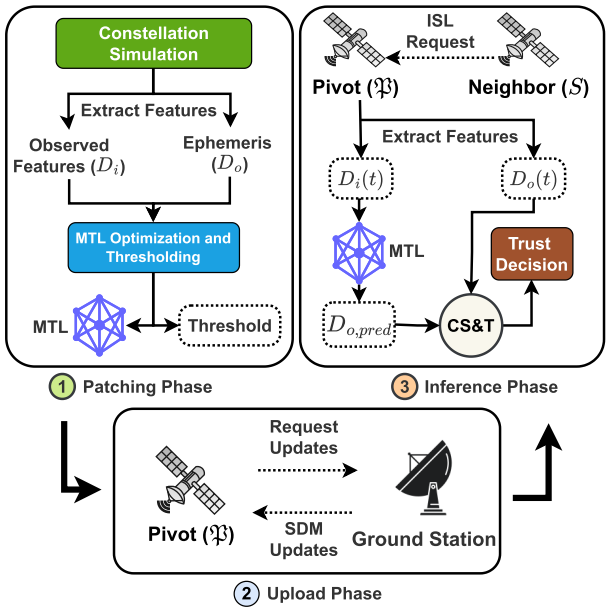

Fig. 3. Stages of the proposed authentication model.

where Γ_1 and Γ_4 are the request transmission and reply reception timestamps, respectively, at $\mathfrak{P}$, and Γ_2 and Γ_3 are the request reception and reply transmission timestamps, respectively, at S.

Ephemeris: In addition to the above, $\mathfrak{P}$ receives its neighbors' reference parameters through periodic ephemeris broadcasts from the network of terrestrial GSs. These features include the instantaneous azimuth angle $\theta_{az}^{\mathfrak{P},S}(t)$, the instantaneous angle of elevation $\theta_{el}^{\mathfrak{P},S}(t)$, and the range $r^{\mathfrak{P},S}(t)$.

The observed and reference parameters are compiled into the input and output feature vectors $D_i^{\mathfrak{P},S}(t) = [\mathfrak{L}^{\mathfrak{P},S}(t), f_\Delta^{\mathfrak{P},S}(t), \Theta^{\mathfrak{P},S}(t)]$ and $D_o^{\mathfrak{P},S}(t) = [\theta_{az}^{\mathfrak{P},S}(t), \theta_{el}^{\mathfrak{P},S}(t), r^{\mathfrak{P},S}(t)]$, respectively for ML training and evaluation.

C. Multi-Task Learning (MTL)

MTL is an ML approach where a model learns multiple tasks simultaneously [12]. MTL is effective when tasks share structure or information, which is the case in the observed physical behaviors of satellites discussed in the previous subsection. Given K tasks, each with dataset $D_j = \{(x_i^j, y_i^j)\}$, MTL minimizes the joint loss:

$$\mathcal{L}_{MTL} = \sum_{k=1}^{K} \lambda_t \mathcal{L}_k(f_k(x; \phi_o, \phi_k)), \tag{3}$$

where ϕ_o denotes shared parameters, ϕ_k are task-specific parameters, $\mathcal{L}_k$ is the loss for task k, and λ_k the parameter to balance a task's importance. MTL reduces overfitting, promotes the reuse of representations, and often results in better performance compared to training separate models.

Given the size of our feature set and the task to instantaneously predict the trustworthiness of an ISL request using a single sample, we use Multi-layer Perceptron (MLP)-based MTL models that takes an instantaneous observed parameter

Algorithm 1 Algorithm for Spoof Detection on $\mathfrak{P}$

$S \leftarrow ReadID(beacon)$
$D_o(t) \leftarrow FetchEphemeris(S, t)$
$D_i(t) \leftarrow ObservedFeatures(beacon)$
$D_i(t) = \frac{D_i(t) - min(D_i^{\mathfrak{P},S})}{max(D_i^{\mathfrak{P},S}) - min(D_i^{\mathfrak{P},S})}$ // normalize
$D_{o,pred}(t) = \mathcal{M}_S^{\mathfrak{P}}(D_i(t))$
$similarity = CS(D_{o,pred}(t), D_o(t))$
$decision := similarity \geq \mathcal{T}_S^{\mathfrak{P}}\ ?\ accept : reject$

set $D_i^{\mathfrak{P},S}(t)$ and predicts reference ephemeris parameters $D_o^{\mathfrak{P},S}(t)$. Each task is assigned equal importance. Implementation details are provided in Sections V and III-.3.

III. EMMA: ISL Spoof Detection Framework

This section presents the proposed three-stage solution, as encapsulated in Fig. 3.

1) Patching: The patching phase is initiated on the ground infrastructure whenever a change (such as neighborhood membership, orbital paths, deviation due to environmental factors, decommissioning, etc.) is made by the operator to the constellation. The neighborhood of each $\mathfrak{P}$ is determined using a maximum range-based neighborhood criterion. A training dataset $\mathcal{F}^{\mathfrak{P},S} = \{D_i^{\mathfrak{P},S}(t) || D_o^{\mathfrak{P},S}(t)\}_{24hrs} = \{D_i^{\mathfrak{P},S} || D_o^{\mathfrak{P},S}\}$ is created by sampling observed and ephemeris features every 10 seconds over a 24-hour simulation for all $\mathfrak{P}$ and their expected neighborhoods S. The dataset is normalized using the minimum and maximum values of each feature, and an MLP-MTL model ($\mathcal{M}_S^{\mathfrak{P}}$) is trained to minimize the mean squared error of predicting the reference parameters. Then, a scaled Cosine Similarity (CS) threshold ($\mathcal{T}_S^{\mathfrak{P}}$) is evaluated using the following equations:

$$\text{CS}(x, y) = \frac{x \cdot y}{|x||y|}, \tag{4}$$

$$\mathcal{T}_S^{\mathfrak{P}} = \Psi \cdot \min\{\text{CS}(D_o, \mathcal{M}_S^{\mathfrak{P}}(D_i))\}, \tag{5}$$

where x and y are vectors, and Ψ is the threshold scaling factor. A CS threshold is used instead of mean squared error because it is suitable for behavior matching tasks and allows for intuitive and flexible scaling. The resultant $\mathcal{M}_S^{\mathfrak{P}}$ and $\mathcal{T}_S^{\mathfrak{P}}$ are compiled into a Spoof Detection Module ($SDM_S^{\mathfrak{P}} = [\mathcal{M}_S^{\mathfrak{P}}, \mathcal{T}_S^{\mathfrak{P}}]$) and uploaded to the satellite during Phase 2.

2) Upload: After creating an SDM update, the satellite operator transmits the updates (new SDMs or decommissioned SDMs) to the satellites through a feeder link between the pivots and GS. We assume the *feeder* links to be secure, and their security out of the scope of this letter.

3) Inference: Each $\mathfrak{P}$ executes Algorithm 1 upon receiving an ISL request from a claimed registered neighbor S: $\mathfrak{P}$ extracts the sender's identifier from the pilot message (say, the neighbor claims to be a registered satellite S). It evaluates the observed parameters from the beacon signal $D_i(t) = [\mathfrak{L}^{\mathfrak{P},S}(t), f_\Delta^{\mathfrak{P},S}(t), \Theta^{\mathfrak{P},S}(t)]$ and fetches the expected ephemeris of S at that time, $D_o(t) = [\theta_{az}^{\mathfrak{P},S}(t), \theta_{el}^{\mathfrak{P},S}(t), r^{\mathfrak{P},S}(t)]$. It then calculates the predicted ephemeris set $D_{o,pred} = \mathcal{M}_S^{\mathfrak{P}}(D_i(t))$ and computes $similarity = CS(D_{o,pred}(t), D_o(t))$. The null hypothesis

($\mathcal{H}_0$) is that the ISL request's sender is truly S. The alternate hypothesis ($\mathcal{H}_1$), however, is that the sender is a spoofer ($\hat{S}$). We define the evaluation outcome ($\mathcal{O}$) as follows:

$$\mathcal{O} = \begin{cases} \mathcal{H}_0 & : similarity \geq \mathcal{T}_S^{\mathfrak{P}} \\ \mathcal{H}_1 & : similarity < \mathcal{T}_S^{\mathfrak{P}} \end{cases} . \quad (6)$$

If $\mathcal{O} = \mathcal{H}_0$, the ISL request is accepted; otherwise, the request is rejected. Such an evaluation works for the following reason: since the outputs of the MTL models are continuous variables (and not fixed one-hot encoded vectors), changes in input distributions may lead to an observable change in the output distribution.

IV. Simulation Dataset and Results

This section discusses the simulation, dataset, hyperparameters, results, comparison with prior work, and limitations.

A. Simulation, Datasets and Tasks

To evaluate the proposed model, an Iridium constellation with 60 satellites in near-polar orbits was simulated for 72 hours using MATLAB's Satellite Communication Toolbox. The simulation focused on 10 pivot satellites, each with up to 20 neighboring satellites. For each $(\mathfrak{P}, S)$, six features corresponding to the parameter set $\mathcal{F}^{\mathfrak{P},S}$ were sampled every 10 seconds. Fig. 4a shows the temporal variation in the number of neighboring satellites for an arbitrary $(\mathfrak{P}, S)$ over time. The neighborhood changes dynamically with time. Similarly, Fig. 4b depicts the relative distribution of these neighboring satellites. For each $(\mathfrak{P}, S)$, we create datasets for three tasks: (1) The *authentication* ($Auth$) dataset comprises of test samples of the registered S. (2) *Spoof detection* (S_i) datasets are synthesized across the following three synthetically generated spoofing categories:

- S_1: Samples from neighbors of $\mathfrak{P}$ that are *not* S but claim to be S, modeling an adversary attempting to spoof $\mathfrak{P}$ while being incompatible with S's expected orbital state during a valid neighborhood window.
- S_2: Legitimate observed parameters received at incorrect timings within a neighborhood window of S, modeling a spoofing scenario that is difficult to execute and detect.
- S_3: Legitimate observed parameters replayed outside neighborhood windows, modeling adversaries that follow legitimate orbital paths at incorrect timing windows.

For reference, we collectively refer to them as $\mathcal{S}_{avg}$ in the aggregated analyses. Finally, (3) a *sensitivity analysis* dataset is generated by perturbing $Auth$ samples with uniform noise of varying degrees. This category gives an insight into the error bound of the MTL models and spoofing detection for varying degrees of deviation from legitimate orbital paths.

Pre-processing: We use a 33/67 train/test split of registered data. The training set consists of an average of 1,500 samples, while the testing set consists of an average of 80,000 samples (2,500 for $Auth$ and S_2, 20,000 for S_1, and 50,000 for S_3). All features are normalized to a [0, 1] range using the features observed during training.

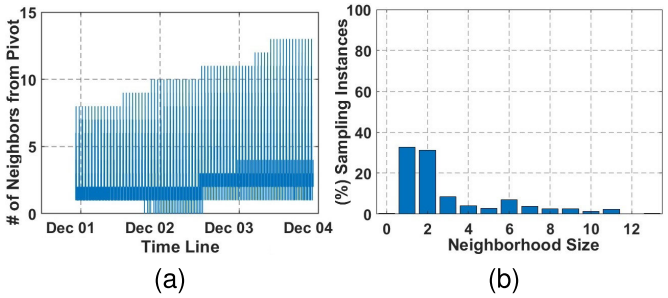

Fig. 4. Temporal changes in $\mathfrak{P}$'s neighborhood: (a) Variation of neighborhood size in 72 Hrs; (b) (%) distribution of neighborhood sizes over 72 Hrs.

TABLE II
Average Detection Rates (%), FPR (%) and Model Sizes of the Best-Performing MLP-MTL Models

Pivot	Neighbors	MLP-MTL=[3,128,3]			MLP-MTL=[3,128,128,3]		
		$Auth$ (%)	$\mathcal{S}_{avg}$ (%)	FPR (%)	$Auth$ (%)	$\mathcal{S}_{avg}$ (%)	FPR (%)
1	20	94.62	93.04	6.17	91.5	92.64	7.93
2	20	94.7	92.95	6.175	89.84	92.71	8.725
3	20	92.17	93.33	7.25	91.08	93.21	7.855
4	20	93.22	94.17	6.305	92.94	94.13	6.465
5	20	93.98	94.23	5.895	93.8	94.07	6.065
6	16	97.04	96.2	3.38	97.4	96.15	3.225
7	20	94.04	94.22	5.87	93.07	94.15	6.39
8	16	96.55	96.13	3.66	97.33	96.16	3.255
9	20	94.37	94.3	5.665	93.19	94.08	6.365
10	16	96.83	96.07	3.55	97.97	96.05	2.99
Average		**94.752**	**94.46**	**5.39**	93.81	94.33	5.93
Total parameters		899			17,411		
Model size (KB)		3.5			68.1		
Inference latency (s)		10^{-4}			10^{-4}		

B. Models and Hyperparameters

To achieve low latency and memory overhead, and high predictive performance with the simplest design, we performed a grid search over depths, hidden dimensions, and other hyperparameters (dropout, learning rate, and Ψ) of various neural network architectures among MLP, MLP-CNN, and MLP-LSTM, among which the MLP-MTL models showcased the best results. In the interest of space, we highlight the best-performing models among the 3-layer and 4-layer MLP-MTL architectures, which comprise a three-dimensional input (for the observed parameters), up to two hidden layers followed by a 0.1 probability dropout, and a three-dimensional output layer (for the expected ephemeris). These models are $MTL_3^* = [3, 128, 3]$ and $MTL_4^* = [3, 128, 128, 3]$. All layers are activated using the Rectified Linear Unit (ReLU) activation. The models are optimized using Adam with a 0.005 learning rate for 200 epochs. We also find $\Psi = 0.95$ in Eq. (5) to have the balance between authentication and spoof detection. Lastly, the random seed is set to 42.

V. EMMA: ISL Spoof Detection Framework

A. Numerical Results

Table II reports the average authentication, spoof detection and False Positive Rates (FPR) over 10 pivots for the best performing three- and four-layer MLP-MTL models, i.e. $MTL_3^* = [3, 128, 3]$ and $MTL_4^* = [3, 128, 128, 3]$, respectively. MTL_3^* achieves a 94.75% authentication rate, 94.46% spoof detection rate, and 5.39% FPR. MTL_4^* achieves

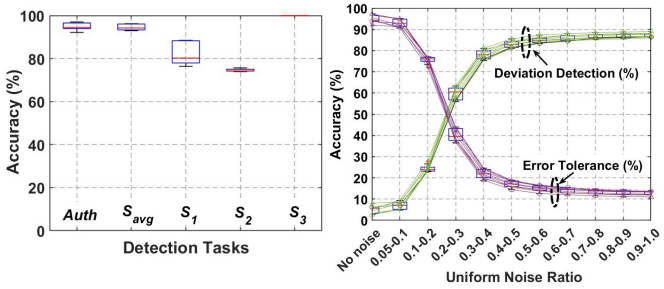

(a) (b)

Fig. 5. Predictive performance of MTL_3^* over 10 pivots. (a) Authentication and spoof detection; (b) Sensitivity analysis for error tolerance and deviation detection.

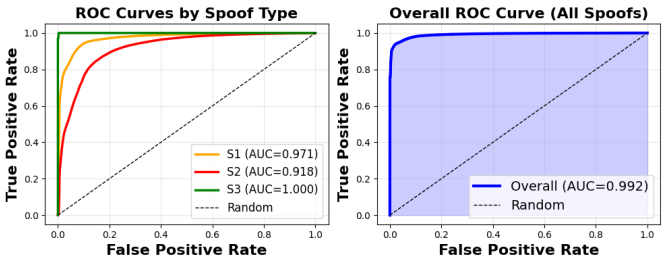

Fig. 6. ROC-AUC for the spoof detection tasks.

a 93.81% authentication rate, 94.33% spoof detection rate, and 5.93% FPR. Evidently, MTL_3^* performs marginally better than MTL_4^* with fewer parameters (899 vs 17,411), and thus, we use it for further analysis. Fig. 5a provides a fine-grained view of MTL_3^*'s predictive performance, and Fig. 5b presents the sensitivity plot of MTL_3^* for all pivots at uniform noise $\sim \mathcal{U}(a, b)$, where $0.05 \leq |a| < |b| \leq 1$. The error tolerance curve illustrates *EMMA*'s capacity to accommodate minor deviations from expected paths, while deviation detection highlights the detection of spoofing satellites that imperfectly mimic the behavior of registered satellites. Fig. 6 presents the Receiver Operating Characteristic Area Under the Curve (ROC-AUC) for all spoofing types. Finally, the proposed model requires less than twenty seconds to train, occupies 3.5 KB of on-satellite memory, and has an inference latency of 10^{-4} seconds on an Intel i7 processor. *EMMA*'s high predictive performance, low latency, and low memory overheads highlight an ease of integration into existing authentication mechanisms in lightweight LEO satellite platforms, such as CubeSat, as an additional lightweight security check to help detect spoofing and support authentication in LEO constellation deployments.

B. Comparison With Prior Work

There are only three existing works on ISL authentication to the best of our knowledge. An ML-based authentication approach is proposed in [1], which achieves 70-96% authentication accuracy but does not cover spoof detection. The authors of [4] use Doppler shift for key generation, and [5] combines probabilistic decisions made by up to ten satellites to achieve an 80-95% spoof detection probability and 5-20% FPR at a high expected latency. Furthermore, none of these works presents a benchmark dataset for future research. In comparison, *EMMA* achieves nearly 94.5% detection rate and 5.39% FPR on authentication and spoof detection tasks using instantaneous samples collected by a single pivot at low latency. We also present the first simulation dataset for ISL security, which may be used by future research [11].

VI. Conclusion

This letter proposed *EMMA*, a novel framework for real-time authentication and spoof detection in LEO ISL communication. It also presented the first simulation dataset for ISLs collected from a 72-hour simulation of a 60-satellite Iridium constellation to validate the effectiveness of the proposed approach. *EMMA* achieved a 94.75% detection rate on registered satellites, 94.46% spoof detection rate, 5.39% FPR, and incurred low inference latency ($\sim 10^{-4}$ seconds on an Intel i7 CPU) and low memory overhead (3.5 KB). These results indicate that *EMMA* holds strong potential for enhancing ISL communication security for high-speed services. Future works may extend the simulation to denser constellations and multi-operator collaboration, and generate real-world ISL data.

References

[1] M. Abdrabou and T. A. Gulliver, "Authentication for satellite communication systems using physical characteristics," *IEEE Open J. Veh. Technol.*, vol. 4, pp. 48–60, 2023.

[2] (2025). *Space Debris: Is It a Crisis?*. [Online]. Available: https://www.esa.int/ESAMultimedia/Images/2025/04/Around100000satellitesareexpectedtobeinorbitby2030

[3] J. Wigchert, S. Sciancalepore, and G. Oligeri, "Detection of aerial spoofing attacks to LEO satellite systems via deep learning," *Comput. Netw.*, vol. 269, Sep. 2025, Art. no. 111408.

[4] O. A. Topal, G. K. Kurt, and H. Yanikomeroglu, "Securing the inter-spacecraft links: Physical layer key generation from Doppler frequency shift," *IEEE J. Radio Freq. Identificat.*, vol. 5, no. 3, pp. 232–243, Sep. 2021.

[5] O. A. Topal and G. K. Kurt, "Physical layer authentication for LEO satellite constellations," in *Proc. IEEE Wireless Commun. Netw. Conf.*, Austin, TX, USA, Apr. 2022, pp. 1952–1957.

[6] M. Abdrabou and T. A. Gulliver, "Game theoretic spoofing detection for space information networks using physical attributes," *IEEE Trans. Commun.*, vol. 72, no. 7, pp. 3947–3956, Jul. 2024.

[7] G. Oligeri, S. Sciancalepore, and A. Sadighian, "SatPrint: Satellite link fingerprinting," in *Proc. 39th ACM/SIGAPP Symp. Appl. Comput.*, Apr. 2024, pp. 177–185.

[8] J. Smailes et al., "Watch this space: Securing satellite communication through resilient transmitter fingerprinting," in *Proc. SIGSAC Conf. Comput. Commun. Secur.*, 2023, pp. 608–621.

[9] D. Roy, T. Mukherjee, A. Riden, J. Paquet, E. Pasiliao, and E. Blasch, "GANSAT: A GAN and SATellite constellation fingerprint-based framework for GPS spoof-detection and location estimation in GPS deprived environment," *IEEE Access*, vol. 10, pp. 45485–45507, 2022.

[10] Q. Li, M. El-Hajjar, K. Cao, C. Xu, H. Haas, and L. Hanzo, "Holographic metasurface-based beamforming for multi-altitude LEO satellite networks," *IEEE Trans. Wireless Commun.*, vol. 24, no. 4, pp. 3103–3116, Apr. 2025.

[11] A. Bhattacharjee, V. Kohli, S. Shailendra, and B. Sikdar, "Dataset of physical layer features in inter-satellite communication for LEO non-terrestrial networks," Tech. Rep., 2025, doi: 10.21227/y55b-6137.

[12] Y. Zhang and Q. Yang, "A survey on multi-task learning," *IEEE Trans. Knowl. Data Eng.*, vol. 34, no. 12, pp. 5586–5609, Dec. 2022.