Improved RF Fingerprint-based Identity Verification in the Presence of an SEI Mimicking Adversary

Donald R. Reising[1,*], Joshua H. Tyler[1], Mohamed K. M. Fadul[1], Matthew R. Hilling[1] and T. Daniel Loveless[2]

[1]*The University of Tennessee at Chattanooga, Chattanooga, TN 37403 USA*
[2]*Indiana University, Bloomington, IN 47405 USA*
E-mail: donald-reising@utc.edu; joshua-tyler@mocs.utc.edu; mohammed-fadul@utc.edu; matthew-hilling@mocs.utc.edu; dlovele@iu.edu
[*]*Corresponding Author*

Received 19 October 2023; Accepted 11 March 2024

Abstract

Specific Emitter Identification (SEI) is advantageous for its ability to passively identify emitters by exploiting distinct, unique, and organic features unintentionally imparted upon every signal during formation and transmission. These features are attributed to the slight variations and imperfections in the Radio Frequency (RF) front end; thus, SEI is being proposed as a physical layer security technique. Most SEI work assumes the targeted emitter is a passive source with immutable and difficult-to-mimic signal features. However, Software-Defined Radio (SDR) proliferation and Deep Learning (DL) advancements require a reassessment of these assumptions because DL can learn SEI features directly from an emitter's signals, and SDR enables signal manipulation. This paper investigates a strong adversary that uses SDR and DL to mimic an authorized emitter's signal features to circumvent SEI-based identity verification. The investigation considers three SEI mimicry approaches, two different SDR platforms, the application of matched filtering

Journal of Cyber Security and Mobility, Vol. 13_5, 887–916.
doi: 10.13052/jcsm2245-1439.1354

before SEI feature extraction, and selecting the most informative portions of the signals' time-frequency representation using entropy. The results show that "off-the-shelf" DL achieves effective SEI mimicry. Additionally, SDR constraints impact SEI mimicry effectiveness and suggest an adversary's minimum requirements. Our results show matched filtering results in the identity of all authorized emitters being correctly verified at a rate of 90% or higher, the rejection of all other authorized emitters–whose IDs are *not* being verified–at a rate of 97% or higher, and rejection of forty-five out of forty-eight SEI mimicry attacks. Based on the results presented herein, future SEI research must consider adversaries capable of mimicking another emitter's SEI features or manipulating their own.

Keywords: Specific emitter identification (SEI), ID verification, security, SEI mimicry, adversary, RF fingerprint.

1 Introduction

Specific Emitter Identification (SEI) was introduced nearly thirty years ago to enable electronic warfare systems to detect, characterize, and identify radar systems of the same type by exploiting intra-pulse modulation features [16]. These features are a byproduct of "unintentional modulation on pulse" and attributed to the systems, sub-systems, and components (e.g., power amplifier) comprising a radar's Radio Frequency (RF) front-end. SEI is appealing because (i) it is a passive technique, which means that the targeted emitter maintains its intended mission and generates signals without external stimulation; (ii) it exploits the distinct, unique, and organic features unintentionally imparted to the transmitted signal by the emitter's RF front-end; (iii) it can quantitatively measure those exploited features present within the signal; and (iv) it can exploit persistent features across time, location, and environments.

The success of radar SEI led to its extension to the identification of wireless communications emitters in an attempt to augment digital security measures such as encryption and MAC address filtering. Current literature shows that SEI can achieve serial number discrimination (i.e., emitters of the same manufacturer and model), which is the most challenging SEI case. Of particular interest are Deep Learning (DL) based SEI works because they show learning and exploitation directly from the signal's baseband, discrete-time In-Phase, and Quadrature (IQ) samples [5, 13, 17–19, 23, 25, 29, 30]. Direct feature learning opens the door for an emitter to manipulate those features before transmission to prevent or inhibit SEI.

Traditionally, SEI assumes the emitters are passive sources and exploited SEI features are immutable and difficult-to-mimic [11, 32]. These assumptions imply that an emitter is unwilling or unable to develop and implement effective SEI countermeasures. However, recent DL advances and Software-Defined Radio (SDR) flexibility necessitate investigation into whether or not and the degree to which these assumptions hold. We are interested in the ease and degree to which an emitter can inhibit or thwart SEI by masquerading as another emitter through SEI mimicry. Early SEI mimicry investigations are presented in [3, 14, 24, 26]. The authors of [3] investigate defeating transient-based SEI techniques using signal replay; however, they replay the signals using a high-end arbitrary waveform generator and at an unspecified high Signal-to-Noise Ratio (SNR). The authors of [24] present *FIRNet*, a Deep Neural Network (DNN) based system purpose-built to mimic the signal features in another emitter's signals to inhibit or defeat a DL-based SEI process. In [24], all emitters are USRP N210 SDRs, even the authorized emitters, which do not reflect typical communications systems whose user equipment is unlikely to be SDRs. This also makes it easier for the adversary to mimic the authorized emitters' SEI features because emitters of the same manufacturer and model (i.e., only different serial numbers) exhibit the greatest similarity in SEI features. In [26], the authors investigate replay and Generative Adversarial Network (GAN) based signal feature spoofing versus a Deep Learning (DL) based SEI process; however, the approach suffers from two key drawbacks. First, the work is conducted in simulation only, and second, the adversary's receiver must be located near the emitter whose signals are being collected to ensure the channel conditions are the same as those experienced by the SEI process's receiver. Lastly, the authors of [14] use online learning to manipulate the adversary's IQ samples to mimic the signal features of the authorized emitters and achieve a high success rate (i.e., 90% or better) at SNRs of 15 dB and above. Similar to the authors of [24], the authors of [14] use Analog Devices Active Learning Module (ADALM) Pluto SDRs for the adversary and authorized emitters. Thus, their work suffers the same concern mentioned for [24]. Our work in [22] and that presented herein does not suffer any of these highlighted drawbacks. Additionally, our work in [22] differed from these works and extended SEI mimicry understanding in the following ways.

- Assesses the effectiveness of SEI mimicry in degrading or defeating an identity (ID) verification-based SEI process while previous works target classification-based processes.

- Uses "off-the-shelf" algorithms to perform SEI mimicry while previous efforts use purpose or tailored algorithms specifically designed for SEI mimicry. We aim to determine the ease at which SEI mimicry can be implemented.
- Addresses an adversary capable of mimicking SEI features using one of two SDR platforms of differing Size, Weight, Power, and Cost (a.k.a., SWaP-C) to determine how SDR capabilities impact SEI mimicry performance.
- Assesses SEI mimicry performance using an adversary that employs signal replay, a Multi-Layer Perceptron–AutoEncoder (MLP-AE), or Convolutional–Generative Adversarial Network (C-GAN) to mimic the SEI features present within the signals of another emitter.
- Considers the presence or lack of signal energy, which is motivated by our published work in [30].
- Considers the presence or lack of "decoy" emitter signals to aid the ID verification process in discerning the adversary from the authorized emitter being mimicked. A decoy emitter is an SDR of the same manufacturer and model as that used by the adversary.

Despite these contributions, the work in [22] showed that SDR SWaP-C does impact SEI mimicry effectiveness. When the adversary employs a higher SWaP-C SDR, the SEI process is not able to verify the identities of three authorized emitters without also incorrectly identifying the adversary as an authorized emitter when a decoy emitter's RF fingerprints are part of the ID verification process' training set. Even when the adversary uses an SDR of lower SWaP-C, the SEI process still requires using a decoy emitter's RF fingerprints to correctly reject the adversary's mimicked SEI features. To address the ID verification process' susceptibility to SEI mimicry, we extend our previous work in [22] in the following ways.

- Investigates matched filtering of the signals to maximize the SNR before RF fingerprint generation.
- Investigates using entropy to identify the most informative sub-regions to generate RF fingerprints from. These sub-regions are extracted from the two-dimensional (2D) Time-Frequency (TF) representations of the authorized emitters' signals.

Our results show "off-the-shelf" DL algorithms, and SDR enables SEI mimicry; however, adversary success is impacted by (i) the use of decoy emitter signals, (ii) energy normalization, and (iii) SDR SWaP-C constraints. The integration of matched filtering before Gabor-based RF fingerprint generation

allows (i) the IDs of all authorized emitters to be verified at a rate of 90% or higher, (ii) the rejection of all other authorized emitters–whose IDs are *not* being verified–at a rate of 97% or higher, and (iii) rejection of forty-five out of forty-eight SEI mimicry attacks at a rate of 93% or higher.

This paper is organized as follows: Section 2 describes the threat model; Section 3 explains ID verification versus classification, the Signal of Interest, RF fingerprint generation, Relief-F feature selection, and Support Vector Machines (SVM); Section 4 describes the signal collection, detection and post-processing steps; Section 5 presents the results and the paper is concluded in Section 6.

2 Threat Model

The threat model extends our model in [20] whose adversary (a.k.a., Eve) uses simple software tools to falsify its digital credentials (e.g., MAC address) to gain unauthorized network access by digitally posing as an authorized (a.k.a., Alice) device and being incorrectly authenticated by the network monitor (a.k.a., Bob). As in [20], Eve is not an authorized network device and does not have inherent access to the network or the network's devices. Lastly, the network's communication links and hardware are not initially compromised.

This work aims to determine the ease at which Eve can implement SEI mimicry using "off-the-shelf" technologies and gauge attack success. Therefore, Eve is implemented using commercial SDRs, compute resources, and open-source DL algorithms. Two SDR platforms are used to determine SWaP-C impacts on SEI mimicry effectiveness. The SDRs are: Ettus Research's Universal Software Radio Peripheral (USRP) B210 (∼$4,000 per unit) and Great Scott Gadget's HackRF One (∼$470 per unit) [4, 10]. The B210 SDR gives Eve a distinct advantage over the HackRF One SDR because it can receive its own transmitted signals. The HackRF One SDR is a half-duplex system, so a second SDR is needed to receive the signals transmitted by the first, which increases Eve's complexity. Eve uses one of three SEI mimicry attacks.

- *Replay-based SEI Mimicry*: Eve collects, saves, and re-transmits the signals of an authorized emitter. Eve may adjust its transmit power and CFO behavior to modify the replayed signals, but this work does not assess them.
- *AE-based SEI Mimicry:* Eve collects a set of signals transmitted by the targeted, authorized emitter. The MLP-AE is implemented using a

hidden layer of size one hundred, and a targeted Mean Squared Error (MSE) of 1×10^{-6} [29].

- *GAN-based SEI Mimicry:* Eve collects its signals and those of a targeted, authorized emitter. The GAN is trained by assigning the targeted emitter's collected preambles to Class #1 of the GAN's discriminator D and its preambles to Class #2. Class #2 is the GAN's generator G input during training. The G's weights are updated to learn the mapping needed to modify the SEI features–present in Eve's signals–to match those of the targeted emitter (a.k.a., Class #1) at the G's output. How often the D assigns the G's output to Class #1 determines the mapping's success. The D is a basic, three-layer CNN trained to discriminate Class #1's signals from those of Class #2 modified by the G. The G and D are iteratively updated until the D can no longer discern a Class #1 signal from those output by the G. Ideally, the G learns to remove Class #2's inherent SEI features and insert Class #1's.

When training the AE and GAN networks, the energy of each collected preamble is normalized to unity to ensure the corresponding DL network learns the correct feature distributions without signal energy biasing the process. Only the collected preambles' raw IQ samples are used during training. AE- and GAN-based SEI mimicry is implemented by passing the Wi-Fi preamble's IQ samples through the trained AE and G. The resulting outputs are stored for later transmission by Eve.

Eve uses SEI mimicry with a falsified digital credential to increase the chances that Bob incorrectly authenticates it. Bob uses the digital credential and SEI features to verify the ID of the to-be-authenticated device. Lastly, Eve conducts the attack at an SNR of 9 dB because this is an SNR at which SEI processes typically struggle to differentiate one from another, and Wi-Fi is still capable of supporting communications.

3 Background

3.1 ID Verification versus Classification

Most SEI processes use classification to discriminate between emitters [16]. A decision is made in classification using a *one-to-many* comparison. For instance, assume a new RF fingerprint (a.k.a., a collection of SEI features extracted from an emitter's signal)–whose originating emitter is unknown–is collected and input into a classification-based SEI process trained to discriminate between six known emitters. The classifier compares the new

RF fingerprint with each of the six learned classes and assigns the new RF fingerprint to the emitter whose class model results in the "best" match. This decision uses a predetermined criterion and is made no matter how poor the match is, thus representing a vulnerability that an adversary can exploit to circumvent classification-based SEI security mechanisms. This is due to the classifier's forced assignment of every incoming RF fingerprint to one of the six classes, even if those RF fingerprints are from a seventh, unknown emitter.

The work in [2] introduced ID verification to eliminate this vulnerability. ID verification performs a *one-to-one* comparison between a previously unseen RF fingerprint and the stored, learned model of the emitter whose digital credential is claimed by the unseen RF fingerprint's originator. Since the comparison is *one-to-one*, the SEI process is not forced to make a class (a.k.a., originating emitter) assignment when the RF fingerprint poorly matches the stored, learned model. In addition to the work in [2], the works in [20,21] also successfully demonstrated the effectiveness of ID verification in defeating adversaries that falsify their digital IDs in an attempt to circumvent traditional network security approaches (e.g., MAC address filtering); however, none of them considered an adversary capable of employing SEI mimicry.

3.2 Signal of Interest

All results are generated using IEEE 802.11a Wireless-Fidelity (Wi-Fi) preambles. IEEE 802.11a Wi-Fi signals consist of sixty-four Orthogonal Frequency Division Multiplexing (OFDM) sub-carriers transmitted within the 5 GHz Industrial, Scientific, and Medical (ISM) band. The IEEE 802.11a Wi-Fi frame's first 16μs are occupied by a preamble, which is used for carrier frequency and phase correction, equalization, and synchronization [12]. The 802.11a Wi-Fi preamble consists of ten Short Training Symbols (STS), designated s_1 through s_{10}, a Guard Interval (GI), designated s_G, and two Long Training Symbols (LTS), designated s_{L1} and s_{L2}, as shown in Figure 1. IEEE 802.11a Wi-Fi is used because (i) it is an IoT communications protocol [31], (ii) OFDM is used in the 802.11ac, 802.11ad, 802.11ax, 802.11p, and Long Term Evolution (LTE) standards, (iii) it is used in prior SEI work [5, 22, 23, 29, 30], and (iv) access to a set of Commercial-Off-The-Shelf (COTS) IEEE 802.11a Wi-Fi emitters.

3.3 RF Fingerprint Generation

RF fingerprints are generated from the TF representations of each emitter's preambles. The TF representation is the normalized, magnitude-squared

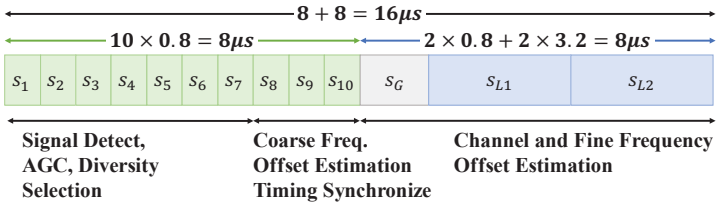

Figure 1 The structure of the IEEE 802.11a Wi-Fi preamble that occupies the first 16 μs of every transmitted signal [12].

coefficients calculated by the Gabor Transform (GT) [8],

$$G_{mk} = \sum_{n=1}^{MN_\Delta} s(n)W^*(n - mN_\Delta)exp^{-j2\pi kn/k_G}, \tag{1}$$

where G_{mk} are the complex-valued Gabor coefficients, $s(n) = s(n+lMN_\Delta)$ is the periodic input signal, $W(n) = W(n-lmN_\Delta)$ is the Gaussian window function, l is the period of the signal or window, the window is shifted N_Δ samples between calculations, $m = 1, 2, \ldots, M$ for M total shifts, $k = 1, 2, \ldots, K_G - 1$ for $K_G \geq N_\Delta$, and $\mod(MN_\Delta, K_G) = 0$ is satisfied [1]. In this work, K_G is greater than N_Δ because it results in an *oversampled* GT, which is advantageous when processing signals collected at low SNR values.

The magnitude of the complex-valued coefficients, G_{mk}, is calculated, the resulting magnitude surface squared, and normalized such that all values are in the range of zero and one. The resulting TF representation is,

$$\overline{|G_{mk}|^2} = \frac{|G_{mk}|^2 - \min\{|G_{mk}|^2\}}{\max\{|G_{mk}|^2\} - \min\{|G_{mk}|^2\}}. \tag{2}$$

Generation of an RF fingerprint starts by dividing the TF representation into N_R two-dimensional sub-regions (a.k.a., patches) that are each comprised of $N_T \times N_F$ values where N_T and N_F are the length of the sub-region along the time and frequency dimension, respectively. A given sub-region is selected, reshaped into a $1 \times N_T \cdot N_F$ vector, and variance, skewness, and kurtosis statistics are calculated. The statistics of subsequent sub-regions are calculated and appended to those from all preceding sub-regions. After all sub-regions statistics are calculated, the statistics are calculated over the entire TF representation and added to the end of the RF fingerprint. An illustration of the RF fingerprint generation process is shown in Figure 2. For the results in Section 5, the RF fingerprints are generated using $N_T = 53$

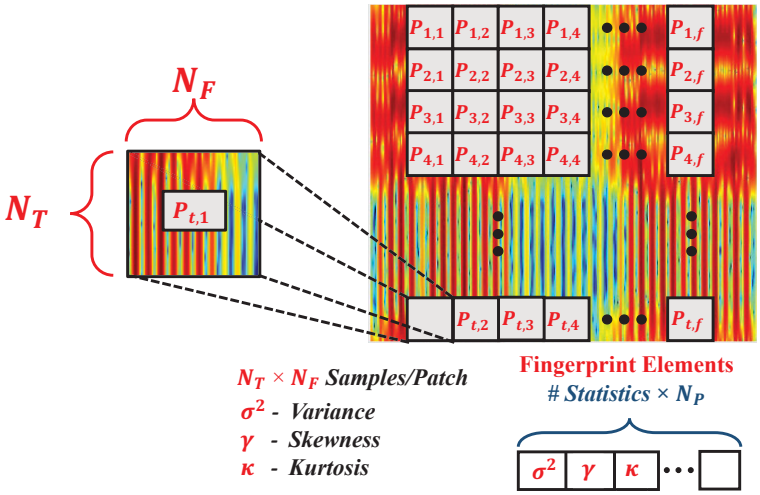

Figure 2 Gabor-based RF fingerprint generation is detailed in [21] and adopted for time-frequency-based SEI.

and $N_F = 4$ for a total of $N_R = 212$ sub-regions per TF representation and $N_f = 639$ features per RF fingerprint. The RF fingerprint generation process is detailed in [6, 20, 21].

3.4 Relief-F Feature Selection

Relief-F is an improved Relief algorithm that accounts for noise, missing values, and more than two classes [15]. Each RF fingerprint feature's quality is iteratively determined by calculating the within-feature dimension distances between a randomly chosen RF fingerprint and each of its N_k nearest in-class and out-of-class neighbors. The value of N_k is empirically chosen for each authorized emitter without knowledge of Eve's signals or RF fingerprints. For a given iteration, each feature's weight value is updated using the calculated distances in conjunction with the weight's previous value, the prior probability of the emitter from which the chosen RF fingerprint is drawn, and the prior probability of the emitter's class from which the nearest out-of-class neighbor is drawn.

3.5 Support Vector Machines

All results are generated using SVMs, which are non-probabilistic, linear two-class classifiers. An SVM maps the RF fingerprints into a space that

maximizes the margin between the classes. The trained SVM maps previously unseen RF fingerprints into the same space, and their predicted class is selected based on the side of the margin on which they fell. SVMs perform non-linear classification using a kernel that maps the RF fingerprints into a higher-dimensional feature space. As in [20], ID verification is performed by a non-linear SVM using a Radial Basis Function kernel.

4 Methodology

4.1 Signal Collection, Detection, and Post-Processing

The SEI process and Eve perform signal collection, detection, and post-processing.

4.1.1 The SEI Process

All signals are transmitted at 5.805 GHz and collected using a Tektronix Real-time Spectrum Analyzer (RSA) 5126B and ultra wide-band antenna. The authorized devices are eight TP-Link Archer T3U USB 802.11a Wi-Fi-compliant emitters. The TP-Links transmit a 2GB binary file to an SFTP server. This ensures a sufficient transfer time with many transmissions from the TP-Link under test. All signals are sampled at 200 MHz and collected at an antenna incident SNR of 9 dB–average value calculated across all collected signals–by placing all emitters eight feet away from the RSA's antenna. The entire collection record is filtered using a fourth-order elliptic filter with a passband ripple of 0.5 dB, a stopband attenuation of 20 dB, and an 8.865 MHz cutoff frequency. Individual IEEE 802.11a Wi-Fi frames are detected and removed from the collection record using an empirically selected amplitude-based threshold. Each IEEE 802.11a Wi-Fi frame's preamble is detected using cross-correlation of the entire frame with an ideal preamble, Carrier Frequency Offset (CFO) correction performed, downsampled to 20 MHz and stored for later use. One thousand preambles are collected for each emitter.

4.1.2 The Adversary

Eve continuously observes the 5.805 GHz Wi-Fi channel and collects all signals using a USRP B210 or HackRF One SDR. A sampling rate of 40 MHz and 20 MHz is used by the USRP B210 and HackRF One, respectively. For each targeted emitter, Eve records four seconds worth of communication using the OSMOCOM source block in GNU Radio. When collecting its signals, Eve transmits an ideally generated Wi-Fi frame. The B210 SDR is

capable of full-duplex communications; thus, Eve can receive its transmitted signals. The HackRF One is a half-duplex SDR; thus, a second HackRF One SDR is used to collect the signals transmitted by a HackRF One-based Eve. Each recording is saved as a complex Float32 binary file (a.k.a., the I and Q channels are saved using 16-bit floating point precision) because the maximum resolution of the Digital-to-Analog Converter (DAC) is 12 bits; there is no need to record the data at single or double precision. Eve loads the binary file(s) into MATLAB® to perform preamble extraction, CFO correction, and energy normalization (i.e., the preambles are of unit energy) as described in Section 4.1.1. Before constructing a specific mimicry attack, all B210 SDR collected preambles are downsampled to 20 Mhz from 40 Mhz. The DL architectures are trained using $1,000$ collected preambles for the AE- and GAN-based SEI mimicry attacks. Once Eve has applied the prepared SEI mimicry approach to one of its preambles, the OFDM payload–constructed using a random sequence of bits–is appended to the preamble, and the resulting frame is saved as a complex Float32 binary file. The saved frame is transmitted via GNU Radio using the OSMOCOM sink block.

4.2 SEI Feature Enhancement Techniques

Two SEI feature enhancement techniques are implemented to improve RF fingerprint-based ID verification performance in the presence of an SEI-mimicking adversary. These techniques are explained in the following two sections.

4.2.1 Matched Filter Processing

The author of [27] defines a matched filter as a conjugated, time-reversed copy of a known signal sequence often used for symbol detection in correlation-based receivers. Matched filters are linear filters that optimize the SNR for additive, stochastic noise channels. A received symbol or symbols are defined as,

$$r(t) = s(t) + n(t), \tag{3}$$

where $s(t)$ is the original, known transmitted symbol or symbols, and $n(t)$ is the additive, stochastic noise. The matched filter's impulse response is,

$$h(t) = \begin{cases} s^*(T-t), & 0 \leq t \leq T \\ 0, & \text{otherwise} \end{cases}, \tag{4}$$

where T is the symbol or symbols duration [27].

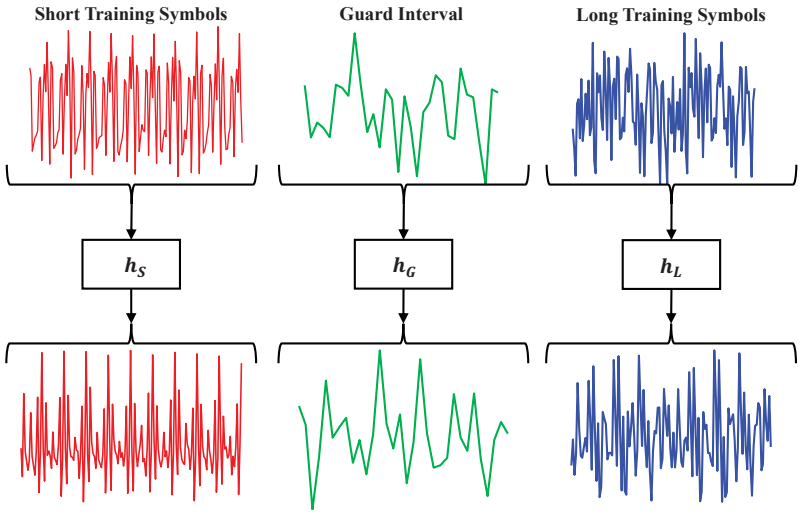

Figure 3 Representative illustration showing, from left to right, matched filtering of a collected preamble's ten short training symbols using h_S, guard interval using h_G, and two long training symbols using h_L. The output of these matched filters is sequentially concatenated to form a matched filtered preamble whose RF fingerprint is then generated using the process described in Section 3.3.

As described in Section 3.2 and shown in Figure 1, the IEEE 802.11a Wi-Fi preamble is comprised of thirteen symbols: ten STS $[s_1(t), s_2(t), \ldots, s_{10}(t)]$, one GI $s_G(t)$, and two LTS $[s_{L1}(t), s_{L2}(t)]$. Ideally, the ten STS are all the same, and the two LTS are copies of one another; thus, three matched filters are constructed. The first matched filter is constructed by extracting a single STS from an ideal IEEE 802.11a Wi-Fi preamble (i.e., one that is free from SEI features, channel, and other impairments or distortions) and its corresponding matched filter, h_S, generated using Equation (4). The GI and LTS matched filters are generated following the same process and designated h_G and h_L, respectively. For a sampling rate of 20 MHz, h_S, h_G, and h_L are comprised of 16, 32, and 64 complex-valued coefficients. Matched filtering is performed by convolving each of a received preamble's thirteen symbols with their corresponding matched filter. Figure 3 provides a representative illustration of a received preamble's ten STS, GI, and two LTS being matched filtered using h_S, h_G, and h_L, respectively. The thirteen matched-filtered symbols are sequentially concatenated in the order corresponding to their unfiltered counterparts. Lastly, each matched filtered preamble is normalized to be of unit energy.

4.2.2 Entropy-Informed Sub-Region Selection

Inspired by the work in [7, 28], which uses entropy to identify the most informative regions of a painting or Gabor-based representation of an IEEE 802.11a Wi-Fi preamble, we investigate using entropy to select the most informative N_R sub-regions from the GT-based images generated from the preambles' TF representation as described in Section 3.3. The TF representation $|\bar{G_{mk}}|^2$ is converted to a grayscale image $\hat{I}_G$. The entropy-informed sub-region selection differs from the approach in [28] in that the selected sub-regions can overlap, are chosen using an exhaustive search in which consecutive sub-regions only differ by one row or column of pixels, and the total number of sub-regions selected is set to 212. The latter point is important because it facilitates direct comparison with the results generated using the RF fingerprint generation process described in Section 3.3.

A GT image's most informative $N_R = 212$ sub-regions are selected using Claude Shannon's definition of entropy [9],

$$\epsilon(\iota) = -\sum_{j=0}^{255} f_\iota[j] \log(f_\iota[j]), \tag{5}$$

where the intensity of a pixel is j, ι is the intensity random variable, and $f_\iota[j]$ is the probability that a given pixel intensity level is within the GT image.

Entropy-informed sub-region selection uses only the GT images generated from the authorized emitters' preambles. In other words, the decoy emitters' and adversary's preambles or their TF representations are never used to select the sub-regions. Sub-region selection begins by randomly selecting 900 preambles from the 1,000 collected for each authorized emitter following the process in Section 4.1.1. The TF representation of each randomly chosen preamble is generated per Section 3.3. However, instead of dividing the TR representation into N_R non-overlapping, two-dimensional sub-regions, the sub-regions are selected from the TF representation using a two-dimensional sliding window that permits overlap between sub-regions. Algorithm 1 describes the process followed in selecting the top N_R entropy-ranked sub-regions from a preamble's GT image $\hat{I}_G$. This algorithm is followed for a given authorized emitter's 900 randomly selected preambles. The entropy values returned across all 900 images are then organized from highest to lowest. The sub-regions associated with the top N_R entropy values across all 900 images–without duplication–are the sub-regions from which RF fingerprints are generated for that authorized emitter and all other emitters when the ID of the selected authorized emitter is being verified. This ensures

Algorithm 1 The algorithm followed to select the N_R top entropy-ranked sub-regions for each authorized emitter.

Input: $\hat{I}_G, N_T, N_F, M, K_G, N_R$
Output: I_T, I_F
Calculate $N_\tau = \left\lfloor \frac{M}{N_T} \right\rfloor$
Calculate $N_\kappa = \left\lfloor \frac{K_G}{N_F} \right\rfloor$
Define $\epsilon_{\text{min}} = -\infty$
Define $E(j) \leftarrow -\infty, j = [1, 2, \ldots, N_R]$
Define $I_T(i, j) \leftarrow 0, i = [1, 2, \ldots, N_\tau]$
Define $I_F(l, j) \leftarrow 0, l = [1, 2, \ldots, N_\kappa]$
for $m = 1$ to N_τ **do**
  Sub-region time indices $\tau_p = \{(m-1)+1 : (m-1)+N_T\}$
  for $k = 1$ to N_κ **do**
    Sub-region frequency indices $\kappa_p = \{(n-1)+1 : (n-1)+N_F\}$
    $S_p = \hat{I}_G(\tau_p, \kappa_p)$
    Calculate entropy $\epsilon(S_p)$ using Eq. (5)
    if $\epsilon(S_p) > \epsilon_{\text{min}}$ **then**
      $[\hat{\epsilon}, \hat{j}] \leftarrow \underset{j}{\text{argmin}}\, E(j)$
      $E(\hat{j}) = \epsilon(S_p)$
      $\epsilon_{\text{min}} \leftarrow \underset{j}{\text{argmin}}\, E(j)$
      $I_T(i, \hat{j} = \tau_p$
      $I_F(l, \hat{j} = \kappa_p$
    end if
  end for
end for

the ID verification process is authorized emitter-centric, thus preserving its one-to-one nature. This process is repeated for the remaining authorized emitters.

ID verification results are generated considering two cases: (i) removing the Relief-F feature selection step and performing ID verification using the RF fingerprints generated using the entropy-selected $N_R = 212$ sub-regions, and (ii) performing Relief-F feature selection on the RF fingerprints generated using the entropy-selected $N_R = 212$ sub-regions before performing ID verification.

5 Results

All results are generated by randomly partitioning the RF fingerprints into a training and a "blind" test data set. The training set consists of the eight

Table 1 Identity Verification Outcomes (adopted from [21])

	System Declaration	
Actual ID	Authorized	Adversary
Authorized	True Verification (TVR)	False Reject (FRR)
Adversary	False Verification (FVR)	True Reject (TRR)

authorized emitters' RF fingerprints; thus, Eve's signals or RF fingerprints are never used in the feature selection and SVM training processes. The number of retained features ranges from four to six-hundred thirty-eight in steps of two between consecutive values. The training set consists of nine hundred RF fingerprints per authorized emitter, and the remaining one hundred is assigned to the "blind" test set. An SVM is trained for each authorized emitter using ten-fold cross-validation. The model associated with the smallest error across all folds is stored and used for blind testing and assessing each SEI mimicry attack's effectiveness. Figure 4 provides the process for verifying an emitter's identity. The four outcomes in Table 1 are adopted from [21]. From the SEI process' perspective, it is desirable to ID authorized emitters at a True Verification Rate (TVR)$\geq$90% while achieving a False Verification Rate (FVR)$\leq$10% when processing RF fingerprints extracted from the signals of the remaining authorized emitters–whose ID's are not being verified–and each of Eve's SEI mimicry attacks.

An SVM model is trained for each number of retained features and authorized emitter, but only one is used to represent each authorized emitter. The "best" SVM model is chosen using our approach in [20]. The approach in [20] is advantageous because SVM model selection uses only knowledge of the authorized emitters' RF fingerprints to achieve a TVR$\geq$90% and FVR$\leq$10% while simultaneously achieving effective adversary rejection performance *without* any knowledge of the adversary's RF fingerprints.

Additionally, this work also considers signal energy and whether or not a "decoy" emitter's RF fingerprints are included with the training set. We consider the "organic" and "unity" energy cases of [30] for signal energy. Organic energy means that the energy of each preamble is unchanged from that at its point of collection. For the unity energy case, the energy of every preamble is normalized to 1 J. In this work, a "decoy" emitter is of the same manufacture and model as the emitter employed by Eve. So, in this case, the decoy emitter is either a USRP B210 or a HackRF One SDR that differs in serial number from Eve's. When a decoy emitter is employed, its RF fingerprints are included with those of the remaining authorized emitters

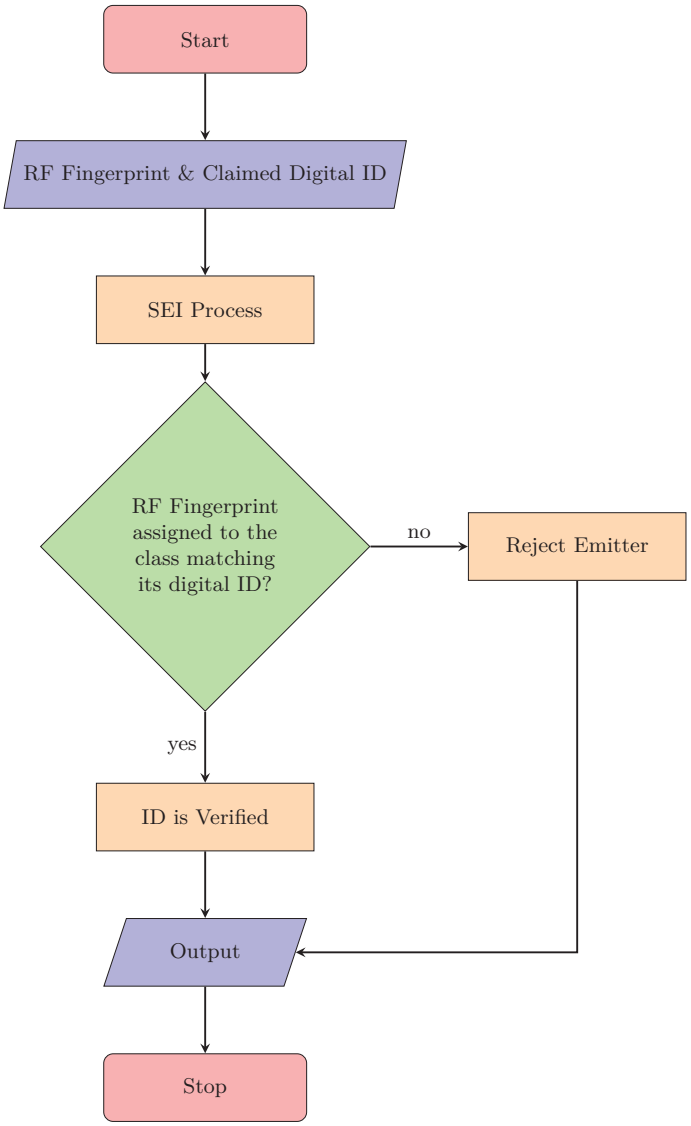

Figure 4 Flowchart of the process used to verify an emitter's identity (ID) or reject that emitter due to the SVM-based ID verification decision not matching the emitter's provided digital ID. It is important to note that when the emitter under test is an authorized emitter (a.k.a., Alice), a "ID is Verified" decision increases the TVR. In contrast, a "Reject Emitter" increases Alice's FRR. For the case of Eve, an "ID is Verified" decision increases the FVR while a "Reject Emitter" increases Eve's TRR.

whose IDs are not verified and designated as "Others" in Figure 5. and Figure 6.

5.1 Results: SEI Mimicry Using a B210 SDR

All B210-based SEI mimicry attacks and authorized emitter ID verification results are shown in Figure 5. Each sub-figure row corresponds to a specific signal energy case, while the columns correspond to the decoy status. It is important to note that a verification rate of 50% represents a guess by the SEI process because it employs a two-class SVM model.

When considering *organic* signal energy *without decoy* emitter RF fingerprints present in the training set, see Figure 5(a), Eve is capable of diminishing the SEI process' ability to distinguish Eve from the targeted, authorized emitter. Eve can mimic the SEI features of four authorized emitters at an FVR of at least 16% and as high as 46% using SEI mimicry. GAN-based proves to be the most effective in terms of the highest achieved FVR values. Interestingly, replay-based mimicry attacks are the second most effective regarding the number of FVR values above 10%.

For *organic* signal energy *with decoy* emitter RF fingerprints present in the training set, Figure 5(b), the inclusion of decoy emitter RF fingerprints does diminish SEI mimicry effectiveness when compared with the corresponding organic signal energy without decoy emitter cases, Figure 5(a).

Unit signal energy *without decoy* emitter RF fingerprints, and a B210-based Eve ID verification results are shown in Figure 5(c). All authorized emitter IDs are verified at a TVR$\geq$90%–the lowest TVR is 96% for Emitter #8–while also correctly rejecting all "Others" and nineteen of Eve's twenty-four SEI mimicry attacks at an FVR$\leq$10%. Replay- and GAN-based SEI mimicry attacks are the most effective regarding the number of attacks with an FVR$>$10% and the highest FVR values. The number of successful mimicry attacks is consistent with the organic signal energy without decoy emitter case, Figure 5(a). Unlike organic energy, ID verification performance changes very little when signal energy is unity.

Unit signal energy *with decoy* emitter ID verification results are shown in Figure 5(d). The results in Figure 5(d) are poorer for Emitter #1, Emitter #2, and Emitter #3 when compared with their results in Figure 5(b). Suggesting that organic signal energy may aid the latter SEI process's ability to discern Eve from these three emitters. Overall, Eve's GAN-based approach is the most successful SEI mimicry attack. This is unsurprising as it is the most

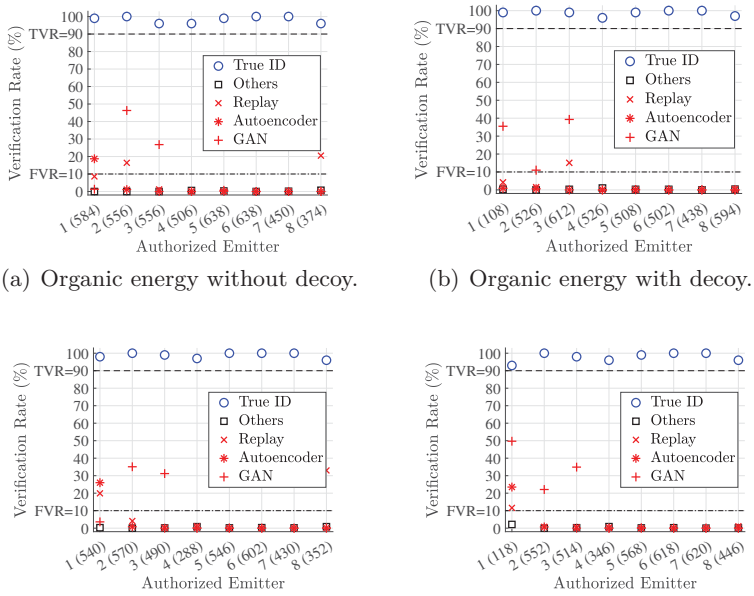

(a) Organic energy without decoy. (b) Organic energy with decoy.

(c) Unit energy without decoy. (d) Unit energy with decoy.

Figure 5 *B210-based SEI Mimcry:* Authorized emitter whose ID is being verified ($\circ$), the remaining authorized emitters ($\square$), replay- ($\times$), AE- ($*$), and GAN-based ($+$) SEI mimicry attacks. The numbers in parentheses along the x-axis are the number of Relief-F retained features associated with the authorized emitter whose ID is being verified. "Others" are authorized emitter whose ID is *not* being verified.

sophisticated and ensures optimal SEI feature manipulation through the D's and G's combative relationship.

5.2 Results: SEI mimicry using a HackRF One SDR

All HackRF-based SEI mimicry attacks and corresponding authorized emitter ID verification results are shown in Figure 6 and whose organization is consistent with Figure 5.

Figure 6(a) shows ID verification and SEI mimicry attack rejection results for the *organic* signal energy *without decoy* emitter case. All but two SEI mimicry attacks are detected without violating the FVR=10% requirement. Compared to the results in Figure 5(a), these results suggest that a lower SWaP-C SDR reduces SEI mimicry attack effectiveness. The HackRF SDR results in less successful mimicry of Emitter #2. Replay- and GAN-based mimicry dropped from 16.4% and 46.3% to 11.4% and 18.6%, respectively.

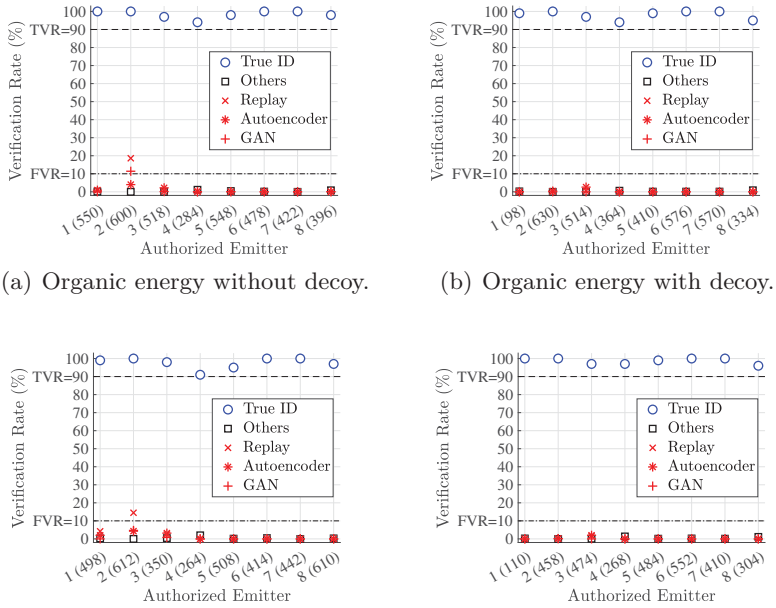

(a) Organic energy without decoy. (b) Organic energy with decoy.

(c) Unit energy without decoy. (d) Unit energy with decoy.

Figure 6 *HackRF One-based SEI Mimcry:* Authorized emitter whose ID is being verified ($\circ$), the remaining authorized emitters *not* being verified ($\square$), replay- ($\times$), AE- ($*$), and GAN-based ($+$) SEI mimicry attacks.

Figure 6(b) presents the *organic* energy *with decoy* emitter ID verification results. Integration of decoy emitter RF fingerprints within the SVM training process allows all–replay, AE, and GAN–of Eve's SEI mimicry attacks to be detected at an FVR of 10% or better while achieving a TVR>90%. None of Eve's attacks exceed an FVR of 5%. These results show that ID verification is robust against a HackRF One-based Eve regardless of signal energy status so long as a decoy emitter's RF fingerprints are used during SVM training.

Unit signal energy *without decoy* emitter RF fingerprints, and HackRF One-based Eve ID verification results are shown in Figure 6(c). All authorized emitter IDs are verified at a TVR$\geq$90%–lowest TVR is 91% for Emitter #4–while also correctly rejecting all "Others" and all but one SEI mimicry attacks with an FVR$\leq$10%. The one SEI mimicry attack with an FVR>10% is when Eve replays Emitter #2's signals, which has an FVR of 14.5%. However, ID verification and adversary rejection are improved over the B210-based results shown in Figure 5(c) that wrongly verified Eve as an authorized emitter five separate times and at FVR$\in\{20, 35\}$%.

Compared with the organic case results in Figure 6(a), the results show that removing signal energy makes ID verification more resilient to SEI mimicry and reinforces our previous findings in [30].

Unit energy *with decoy* ID verification results are shown in Figure 6(d). Eve's SEI mimicry attack is defeated (i.e., FVR$\leq$10%) regardless of the authorized emitter being mimicked. These results are consistent with the organic signal energy results shown in Figure 6(b), thus showing the benefit of using decoy emitter RF fingerprints. Although signal energy does not appear to be a factor in defeating Eve when decoy emitter RF fingerprints are or are not used, it is worth noting that Eve's signal energy was not altered during the attacks.

5.3 Results: SEI Feature Enhancement Techniques

Based upon the results presented and analyzed in Section 5.1 and Section 5.2, the results presented and interpreted in this section are generated using unit energy preambles and without a decoy emitter represented within the SVM training set. This is because signal energy is an adversary exploitable feature [30] while removing the ID verification process' dependency on decoy emitters. The latter is important because using a decoy emitter seems impractical for real-world deployments.

Figure 7 shows the ID verification and rogue rejection results associated with the matched filtering, Figure 7(c) and Figure 7(d), and entropy-informed patch selection, Figure 7(e) and Figure 7(f), techniques for enhancing SEI features. Figure 7(a) and Figure 7(b) are the same as those shown in Figure 5(c) and Figure 6(c), respectively to facilitate direct comparison between them and the remaining results in Figure 7.

Matched filtered preambles ID verification and rogue rejection results are shown in Figure 7(c) when Eve uses a B210 SDR to mimic the SEI features of each of the authorized emitters. All authorized emitters have their IDs verified at a TVR$\geq$90% with the lowest TVR of 94% for Emitter #8 while simultaneously rejecting all "Others" and all SEI mimicry attacks with a FVR$\leq$10%. These results are significant because they not only outperform those shown in Figure 5(c) but also show that the SEI-based ID verification process can correctly identify each of the authorized emitters while rejecting all twenty-four SEI mimicry attacks launched by Eve without the need for a decoy emitter's RF fingerprints being part of the SVM's training set. Figure 7(d) shows the ID verification and rogue rejection results when using matched filtered preambles, and Eve performs SEI mimicry using a HackRF

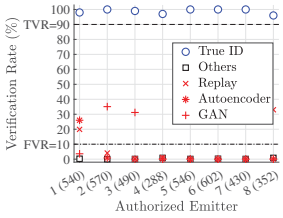

(a) B210 adversary and SEI using the preambles [22].

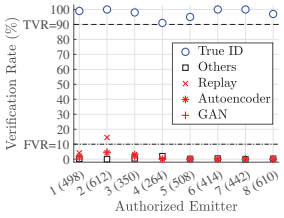

(b) HackRF adversary and SEI using the preambles [22].

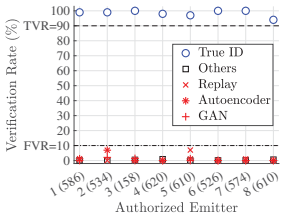

(c) B210 adversary and SEI using matched filtered preambles.

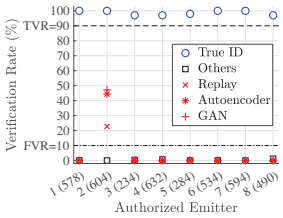

(d) HackRF adversary and SEI using matched filtered preambles.

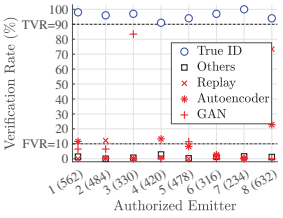

(e) B210 adversary and SEI using entropy-informed patches.

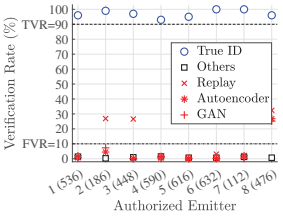

(f) HackRF adversary and SEI using entropy-informed patches.

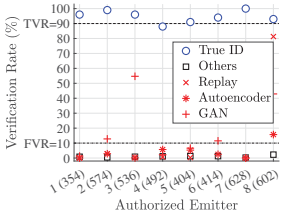

(g) B210 adversary and SEI using matched filtered preambles and entropy-informed patches.

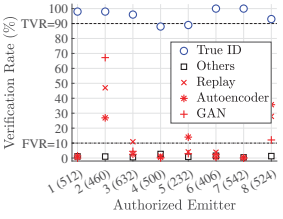

(h) HackRF adversary and SEI using matched filtered preambles and entropy-informed patches.

Figure 7 *Enhanced Feature SEI:* Authorized emitter whose ID is being verified ($\circ$), the remaining authorized emitters ($\square$), replay- ($\times$), AE- ($*$), and GAN-based ($+$) SEI mimicry attacks. The numbers in parentheses along the x-axis are the number of Relief-F retained features associated with the authorized emitter whose ID is being verified. "Others" are authorized emitter whose ID is *not* being verified.

One SDR. As with the B210-based Eve results in Figure 7(c), all authorized emitters have their ID verified at a TVR$\geq$90% with the lowest TVR of 97% for Emitter #3 and Emitter #8 while simultaneously rejecting all "Others" and twenty-one of the SEI mimicry attacks with an FVR$\leq$10%. However, when Eve mimics Emitter #2 using signal replay, an AE, or a trained GAN's G, the SEI process incorrectly verifies Eve as Emitter #2 at a false verification rate of 22.7%, 44.4%, and 47.2%, respectively. One possibility for this poor adversary rejection performance could be that the matched filtering process increases the similarity between the SEI features in Emitter #2's and those in Eve's mimicked signals. Another consideration is that the patch size and location were selected to optimize classification and not ID verification performance [6], so one possible solution is to optimize these values for ID verification. Patch location is considered in the entropy-informed patch selection approach. It is also worth noting that when considering Emitter #2 ID verification and rogue rejection performance in Figure 7(d) versus those in Figure 7(c), the results contradict the findings in [22] that concluded that the HackRF One's lower SWaP-C made it easier for the SEI-based ID verification process to discern Eve from the authorized emitter versus when Eve employs the higher SWaP-C B210. This observation further justifies using multiple SWaP-C SDRs in this and future SEI works.

Entropy-informed selected patches ID verification and rogue rejection results are shown in Figure 7(e) when Eve employs a B210 SDR to perform SEI mimicry of the eight authorized emitters. All eight authorized emitters have their IDs verified at a TVR$\geq$90% with the lowest TVR of 91% corresponding to Emitter #4. All of the "Others" are rejected as rogues at a FVR$\leq$10%; however, Eve is falsely verified (i.e., the FVR is greater than 10%) as an authorized emitter seven times out of the twenty-four SEI mimicry attacks launched by Eve. Eve's greatest success occurs when using a GAN to mimic the SEI features in Emitter #3's preambles, resulting in an FVR of 83.4%. Eve also achieves an FVR of 73.3% when replaying Emitter #8's preambles. When Eve employs AE-based SEI mimicry, FVRs of 11.6%, 13.4%, and 22.8% result when mimicking the SEI features of Emitter #1, Emitter #4, and Emitter #8. The ID verification and rogue rejection results for entropy-informed selected patches and a HackRF One-based Eve are shown in Figure 7(f). The IDs of all authorized emitters are verified at a TVR$\geq$90%, with Emitter #4 achieving a TVR of 93%. All "Others" emitters are correctly rejected at FVRs below 10%. Despite this success, Eve can still achieve FVRs greater than 10% for five of its twenty-four SEI mimicry attacks. Unlike the B210-based Eve, the greatest success occurs when replaying the preambles

of an authorized emitter, specifically when Eve replays the preambles of Emitter #2, Emitter #3, and Emitter #8 to achieve FVRs of 26.9%, 26.6%, and 32.4%, respectively. Eve's next most successful SEI mimicry attack occurs when using an AE to mimic the SEI features present in Emitter #8's preambles, resulting in an FVR of 26.8%. When mimicking the SEI features of Emitter #8, all three of Eve's mimicry attacks are falsely verified above the 10% threshold; however, none exceed 32.4%. Compared to the B210-based Eve results in Figure 7(e), Eve's SEI mimicry attacks are less successful regarding the highest FVR achieved. This reaffirms SWaP-C's impact on the success of Eve's SEI mimicry attacks.

Matched filtered preambles and Entropy-Informed Patch Selection The final set of ID verification and rogue rejection results are generated by selecting the highest entropy-ranked sub-regions drawn from the GT images of each emitter's matched filtered preambles. ID verification and rogue rejection performance results for a B210-based Eve are presented in Figure 7(g). Seven of the eight authorized emitters are verified at a TVR$\geq$90%. Emitter #4 has its ID verified at a TVR of 88%. Despite Emitter #4's TVR, Eve's SEI mimicry attacks are successfully rejected at FVRs lower than 10%. All 'Others' are correctly rejected as rogue emitters at an FVR$\leq$10%, and seventeen of Eve's SEI mimicry attacks are successfully rejected with FVRs lower than the desired 10% rate. Eve's most successful SEI mimicry attack occurs when Eve replays Emitter #8's preambles, resulting in an FVR of 81.3%. Eve's GAN-based SEI mimicry attack achieves the most occurrences of FVRs above 10%. These FVRs occur when mimicking the SEI features of Emitter #2, Emitter #3, Emitter #6, and Emitter #8 with FVRs of 12.8%, 54.7%, 11.5%, and 42.9%, respectively. Eve's AE-based mimicry attack is successfully rejected in seven of the eight attempts. An FVR of 15.7% is achieved when Eve uses the AE-based approach to mimic the SEI features in Emitter #8's preambles. Eve's attacks achieve an FVR above 10% when mimicking the SEI features of Emitter #8, with the AE-based approach being the least successful. Figure 7(h) shows the ID verification and rogue rejection performance associated with a HackRF One-based Eve. Six of the eight authorized emitters have their IDs verified at TVRs greater than 90%. Emitter #4 and Emitter #5 are verified at TVRs of 88% and 89%. Again, Eve's attacks are successfully rejected at FVRs lower than 10% when mimicking Emitter #4's SEI features. The fifty-six 'Others' cases are successfully rejected at FVRs below the 10% threshold. Eve's most successful replay, AE, and GAN-based attacks occur when mimicking the SEI features of Emitter #2, Emitter #8, and Emitter #2 with FVRs of 47%, 35.7%, and 67.2%, respectively. The

GAN-based mimicry of Emitter #2 is Eve's most successful attack using a HackRF One SDR. None of Eve's Emitter #2 mimicking attacks are rejected at FVRs lower than 26.9%, but none exceed the 67.2% achieved using GAN-based mimicry. Eve achieves FVRs of 26.9%, 14%, and 35.7% when mimicking the SEI features of Emitter #2, Emitter #5, and Emitter #8 using an AE.

When considering the B210-based Eve results in Figure 7(e) and Figure 7(g) versus those in Figure 7(c), ID verification and rogue rejection performance is poorer when using entropy-informed sub-region selection. When considering the number of incidents in which the FVR exceeds the 10% threshold, this is also the case when comparing the HackRF One-based Eve results in Figure 7(f) and Figure 7(h) with those in Figure 7(d). For the B210 and HackRF One-based Eve, this could be because the SEI features in the authorized emitters' high entropy sub-regions are more similar to those in Eve's high entropy sub-regions, thus making it more difficult for the corresponding SVM models to discern Eve from the authorized emitters. These results suggest that only matched filtered preambles (i.e., without entropy-informed selection) are needed to improve ID verification and rogue rejection performance for the case when a decoy emitter's RF fingerprints are not present in the SVM training set.

6 Conclusion

SEI mainly treats emitters as passive devices unwilling or incapable of developing and implementing countermeasures to inhibit or thwart SEI. Additionally, SEI work has primarily assumed the exploited features are immutable and challenging to mimic. However, the presented results show that SDR's flexibility and DL's ability to learn SEI features directly from discrete-time signals call into question the extent to which these assumptions hold. This work investigates the adversarial use of "off-the-shelf" algorithms and SDRs to inhibit or thwart ID verification by mimicking authorized emitter SEI features. Our results show that effective SEI mimicry can be implemented using "off-the-shelf" DL algorithms; however, adversary success can be diminished using decoy emitter RF fingerprints and accounting for signal energy. Matched filtering of the signals before SEI feature generation results in the IDs of all authorized emitters being verified at a rate of 90% or higher, the rejection of all other authorized emitters–whose IDs are *not* being verified–at a rate of 97% or higher, and rejection of forty-five out of forty-eight SEI mimicry attacks. This is an improvement over the thirty-nine out

of forty SEI mimicry attacks that are correctly rejected 90% of the time when using non-matched filtered, unit energy signals and without decoy emitter signals or RF fingerprints present in the SVM's training set. Based on the results presented, the viability of SEI–as an effective, operational security mechanism–rests on using strong adversaries within the research and development process. Additionally, this work exposes the technological and algorithmic entry levels required to implement an effective SEI-mimicking adversary. Future research will investigate methods to improve SEI-based ID verification in the presence of a strong adversary. This work will investigate alternate machine learning algorithms that include but are not limited to deep learning architectures. Future research will also consider alternate operating conditions, including channel models that coincide with real-world situations such as multipath and mobility. Lastly, the presented results are limited to eight authorized emitters and two adversary SDRs (i.e., one B210 and one HackRF One); thus, further research is needed when the number of authorized emitters is more in line with a typical IoT deployment (e.g., 30 to 50 emitters).

References

[1] Martin Bastiaans and Marc Geilen. On the discrete gabor transform and the discrete zak transform. *Signal processing*, 49(3), 1996.

[2] William E Cobb, Eric D Laspe, Rusty O Baldwin, Michael A Temple, and Yong C Kim. Intrinsic physical-layer authentication of integrated circuits. *IEEE Transactions on Information Forensics and Security*, 7(1):14–24, 2011.

[3] Boris Danev, Heinrich Luecken, Srdjan Capkun, and Karim El Defrawy. Attacks on physical-layer identification. In *Proceedings of the third ACM conference on Wireless network security*, pages 89–98, 2010.

[4] Ettus Research. USRP Hardware Driver and USRP Manual. https://www.ettus.com/all-products/ub210-kit/, May 26, 2020.

[5] Mohamed KM Fadul, Donald R Reising, and Mina Sartipi. Identification of ofdm-based radios under rayleigh fading using rf-dna and deep learning. *IEEE Access*, 9:17100–17113, 2021.

[6] Mohamed KM Fadul, Jordan T Willis, Donald R Reising, and T Daniel Loveless. An analysis of process parameters for the optimization of specific emitter identification under rayleigh fading. In *Global IoT Summit*, pages 277–291. Springer, 2022.

[7] Steven Frank. This AI can Spot an Art Forgery: With millions at stake, deep learning enters the art world. *IEE Spectrum*, Aug 2021.

[8] Dennis Gabor. Theory of communication. part 1: The analysis of information. *Journal of the Institution of Electrical Engineers-part III: radio and communication engineering*, 93(26):429–441, 1946.

[9] Rafael C Gonzalez, Steven L Eddins, Richard Eugene Woods, et al. *Digital image processing using MATLAB*. Upper Saddle River, NJ: Pearson/Prentice Hall,, 2004.

[10] Great Scott Gadgets. HackRF One. https://greatscottgadgets.com/hackrf/one/, May 26, 2020.

[11] Hao Han, Li Cui, Wen Li, Luying Huang, Yuan Cai, Jihao Cai, and Yuli Zhang. Radio frequency fingerprint based wireless transmitter identification against malicious attacker: An adversarial learning approach. In *2020 International Conference on Wireless Communications and Signal Processing (WCSP)*, pages 310–315. IEEE, 2020.

[12] IEEE. *IEEE Std 802.11-2007, Local and Metropolitan Area Networks, Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, Jun 2007.

[13] Hossein Jafari, Oluwaseyi Omotere, Damilola Adesina, Hsiang-Huang Wu, and Lijun Qian. Iot devices fingerprinting using deep learning. In *Military Communications Conference (MILCOM)*, pages 1–9. IEEE, 2018.

[14] Samurdhi Karunaratne, Enes Krijestorac, and Danijela Cabric. Penetrating rf fingerprinting-based authentication with a generative adversarial attack. In *International Conference on Communications (ICC)*, pages 1–6. IEEE, 2021.

[15] Igor Kononenko. Estimating attributes: Analysis and extensions of relief. In *European conference on machine learning*, pages 171–182. Springer, 1994.

[16] Lawrence E Langley. Specific emitter identification (sei) and classical parameter fusion technology. In *Proceedings of WESCON'93*, pages 377–381. IEEE, 1993.

[17] Guyue Li, Jiabao Yu, Yuexiu Xing, and Aiqun Hu. Location-invariant physical layer identification approach for wifi devices. *Ieee Access*, 7:106974–106986, 2019.

[18] Gihan Janith Mendis, Jin Wei-Kocsis, and Arjuna Madanayake. Deep learning based radio-signal identification with hardware design. *IEEE Transactions on Aerospace and Electronic Systems*, 55(5):2516–2531, 2019.

[19] Yiwei Pan, Sihan Yang, Hua Peng, Tianyun Li, and Wenya Wang. Specific emitter identification based on deep residual networks. *IEEE Access*, 7:54425–54434, 2019.

[20] Donald Reising, Joseph Cancelleri, T Daniel Loveless, Farah Kandah, and Anthony Skjellum. Radio identity verification-based iot security using rf-dna fingerprints and svm. *IEEE Internet of Things Journal*, 8(10):8356–8371, 2020.

[21] Donald R Reising, Michael A Temple, and Julie A Jackson. Authorized and rogue device discrimination using dimensionally reduced rf-dna fingerprints. *IEEE Transactions on Information Forensics and Security*, 10(6):1180–1192, 2015.

[22] Donald R Reising, Joshua H Tyler, Mohamed KM Fadul, Matthew R Hilling, and T Daniel Loveless. Rf fingerprint-based identity verification in the presence of an sei mimicking adversary. In *2023 19th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*, pages 438–444. IEEE, 2023.

[23] Francesco Restuccia, Salvatore D'Oro, Amani Al-Shawabka, Mauro Belgiovine, Luca Angioloni, Stratis Ioannidis, Kaushik Chowdhury, and Tommaso Melodia. Deepradioid: Real-time channel-resilient optimization of deep learning-based radio fingerprinting algorithms. In *Proceedings of the Twentieth ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pages 51–60, 2019.

[24] Francesco Restuccia, Salvatore D'Oro, Amani Al-Shawabka, Bruno Costa Rendon, Kaushik Chowdhury, Stratis Ioannidis, and Tommaso Melodia. Hacking the waveform: Generalized wireless adversarial deep learning. *arXiv preprint arXiv:2005.02270*, 2020.

[25] Kejin Sa, Dapeng Lang, Chenggang Wang, and Yu Bai. Specific emitter identification techniques for the internet of things. *IEEE Access*, 8:1644–1652, 2019.

[26] Yi Shi, Kemal Davaslioglu, and Yalin E Sagduyu. Generative adversarial network in the air: Deep adversarial learning for wireless signal spoofing. *IEEE Transactions on Cognitive Communications and Networking*, 7(1):294–303, 2020.

[27] Bernard Sklar. *Digital communications: fundamentals and applications*. Pearson, 2021.

[28] Mohamed A Taha, Mohamed KM Fadul, Joshua H Tyler, Donald R Reising, and T Daniel Loveless. An assessment of entropy-based data reduction for sei within iot applications. In *Military Communications Conference (MILCOM)*, pages 385–392. IEEE, 2022.

[29] Joshua H Tyler, Mohamed KM Fadul, Donald R Reising, and Erkan Kaplanoglu. Simplified denoising for robust specific emitter identification of preamble-based waveforms. In *2021 IEEE Global Communications Conference (GLOBECOM)*, pages 1–7. IEEE, 2021.

[30] Joshua H Tyler, Mohamed MK Fadul, Donald R Reising, and Farah I Kandah. An analysis of signal energy impacts and threats to deep learning based sei. In *ICC 2022-IEEE International Conference on Communications*, pages 2077–2083. IEEE, 2022.

[31] WiSilica. Top 6 IoT Communication Protocols. https://wisilica.com/company/top-6-iot-communication-protocols/, Aug. 2020.

[32] Qiang Xu, Rong Zheng, Walid Saad, and Zhu Han. Device fingerprinting in wireless networks: Challenges and opportunities. *IEEE Communications Surveys & Tutorials*, 18(1):94–104, 2015.

Biographies

Donald R. Reising received a B.S. degree in electrical engineering from the University of Cincinnati, Cincinnati, OH, in 2006 and an M.S. and Ph.D. in electrical engineering from the Air Force Institute of Technology, Dayton, OH, in 2009 and 2012, respectively. From 2008 to 2014, he was an electronics engineer with the U.S. Air Force Research Laboratory at Wright-Patterson Air Force Base, Dayton, OH. He is an Alexander and Charlotte Guerry and University of Chattanooga (UC) Foundation Associate Professor of Electrical Engineering with the University of Tennessee at Chattanooga, Chattanooga, TN, USA. His research interests include wireless device discrimination using RF distinct native attribute fingerprints, deep learning, next-generation communications systems, dynamic spectrum access, and critical infrastructure protection. He is a senior member of the Institute of Electrical and Electronics Engineers (IEEE) and a member of Eta Kappa Nu and Tau Beta Pi.

Joshua H. Tyler is a doctoral student and researcher at the University of Tennessee at Chattanooga. His research interests include specific emitter identification, power quality analysis, digital communications, deep learning, and edge computing. He graduated Bachelor of Science in Electrical Engineering from the University of Tennessee at Chattanooga in 2020, with a Master of Science in Electrical Engineering in 2022. He is currently employed as a research associate in the Electrical Engineering department at UTC. He is a student member of the Institute of Electrical and Electronics Engineers (IEEE).

Mohamed K. M. Fadul received a B.S. degree in electrical and electronics engineering from the University of Khartoum, Khartoum, Sudan, in 2012 and the M.S. and Ph.D. degrees in electrical engineering and computational engineering from The University of Tennessee at Chattanooga, Chattanooga, TN, USA, in 2018 and 2022, respectively. He is a postdoctoral researcher at The University of Tennessee at Chattanooga, Chattanooga, TN, USA. His research interests include software-defined radios, wireless device discrimination using RF distinct native attribute fingerprints, and deep learning.

Matthew R. Hilling received a B.S. in electrical engineering from The University of Tennessee at Chattanooga, Chattanooga, TN, USA, in 2021

and was a graduate research assistant within Dr. Reising's Wireless Sensing Group (WSG) during the 2021–2022 academic year.

T. Daniel Loveless is an Associate Professor of Intelligent Systems Engineering at Indiana University. They received a B.S. degree in electrical engineering from Georgia Institute of Technology, Atlanta, Georgia, in 2004 and M.S. and Ph.D. in electrical engineering from Vanderbilt University, Nashville, Tennessee, in 2007 and 2009, respectively. Dr. Loveless was a Guerry Professor of Electrical Engineering at the University of Tennessee at Chattanooga from 2014 to 2023. Before joining UTC in 2014, Dr. Loveless was a senior engineer and Research Assistant Professor at the Institute for Space and Defense Electronics (ISDE) at Vanderbilt University. Their research interests include radiation effects and reliability in electronic and photonic integrated circuits; high-performance and radiation-hardened digital, mixed-signal, and analog integrated circuit design; embedded systems; field-programmable gate arrays (FPGAs); microprocessors and microcontrollers; systems-on-chip; and CubeSat design. Dr. Loveless has published over 110 articles in peer-reviewed journals, is a Senior Member of IEEE, and is an Associate Editor of the IEEE Transactions on Nuclear Science. Dr. Loveless' honors include the inaugural 2019 Nuclear and Plasma Sciences Society (NPSS) Radiation Effects Early Achievement Award, five best conference paper awards, and the Institute of Electrical and Electronics Engineers (IEEE) NPSS Graduate Scholarship Award for recognition of contributions to the fields of nuclear and plasma sciences. He is a member of the American Society for Engineering Education (ASEE) and a senior member of IEEE.