

Radio Frequency Fingerprinting via Deep Learning: Challenges and Opportunities

Saeif Al-Hazbi , Ahmed Hussain , Savio Sciancalepore , Gabriele Oligeri , Panos Papadimitratos 

*College of Science and Engineering (CSE), Hamad Bin Khalifa University (HBKU) – Doha, Qatar
 {salhazbi, goligeri}@hbku.edu.qa

†Networked Systems Security group, KTH Royal Institute of Technology – Stockholm, Sweden
 ahmed.hussain@ieee.org, papadim@kth.se

‡Eindhoven University of Technology – Eindhoven, Netherlands. s.sciancalepore@tue.nl

Abstract—Radio Frequency Fingerprinting (RFF) techniques promise to authenticate wireless devices at the physical layer based on inherent hardware imperfections introduced during manufacturing. Such RF transmitter imperfections are reflected into over-the-air signals, allowing receivers to accurately identify the RF transmitting source. Recent advances in Machine Learning, particularly in Deep Learning (DL), have improved the ability of RFF systems to extract and learn complex features that make up the device-specific fingerprint. However, integrating DL techniques with RFF and operating the system in real-world scenarios presents numerous challenges, originating from the embedded systems and the DL research domains. This paper systematically identifies and analyzes the essential considerations and challenges encountered in the creation of DL-based RFF systems across their typical development life-cycle, which include (i) data collection and preprocessing, (ii) training, and finally, (iii) deployment. Our investigation provides a comprehensive overview of the current open problems that prevent real deployment of DL-based RFF systems while also discussing promising research opportunities to enhance the overall accuracy, robustness, and privacy of these systems.

Index Terms—Physical Layer Security, Specific Emitter Identification, Deep Learning, Internet of Things, Wireless Security.

I. INTRODUCTION

Radio Frequency Fingerprinting (RFF) techniques have recently gained attention in the scientific community as a way to authenticate Radio Frequency (RF) devices based on samples of their Physical (PHY) layer signals [1]. RFF grounds on the assumption that two RF devices are unlikely to transmit signals with the same PHY layer features, even when they convey the same data. Therefore, a rogue transmitter that tries to mimic a legitimate Internet of Things (IoT) device by replaying its messages would inevitably leave its own fingerprint, distinct at the PHY layer from that of the legitimate device. This unique property can serve as either an alternative or a direct complement to existing cryptography-based authentication protocols. In particular, being orthogonal to cryptographic protocols and requiring no modifications to the transmitter’s hardware or software, RFF is well-suited for situations where updates are

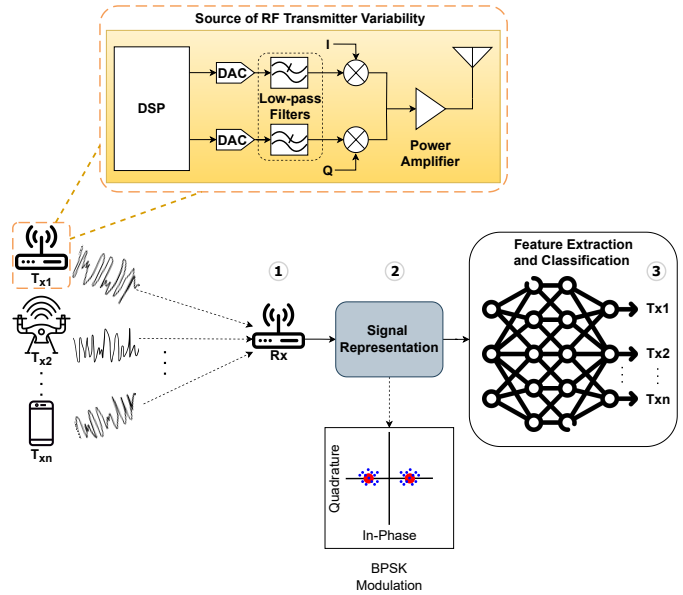


Fig. 1. RFF building blocks: The receiver identifies the transmitting source by analyzing the received signal at the PHY layer.

difficult (e.g., satellite networks) or devices are characterized by strict hardware constraints (e.g., IoT) [2].

The identification and extraction of PHY layer features required for RFF involve integrating advanced data mining techniques as part of the system. Specifically, the use of Deep Learning (DL) techniques in RFF systems has shown significant potential to improve performance, due to the remarkable ability of DL models to identify complex patterns and extract relevant features from raw data [3]. The effectiveness of DL-based RFF systems in accurately identifying RF devices has been reported in the literature, but these outstanding results were obtained primarily under controlled laboratory conditions. At the time of this writing, RFF deployment remains challenging in real-world applications.

Factors such as the calibration and configuration of the DL methods, especially within the context of RFF systems, the unpredictability of the radio channel, and the deployment conditions contribute to the difficulties of adapting DL solu-

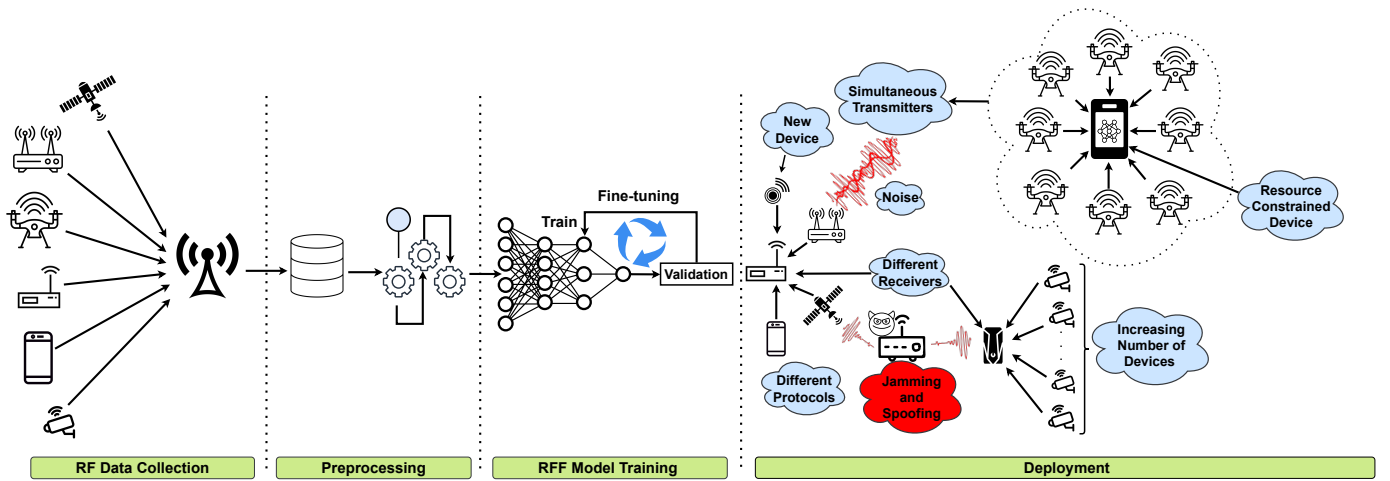


Fig. 2. Typical workflow of DL-based RFF.

tions for RFF systems. Challenges associated with DL-based RFF systems are discussed in several recent studies [4]–[8]; however, such considerations are disconnected from the real-world deployment of such systems.

Thus, we currently miss a systematic classification of the main challenges that prevent the adoption of DL-based RFF systems on a large scale in real-world scenarios.

Contribution. In this paper, we systematically identify and classify the building blocks necessary for assembling DL-based RFF systems, as well as the challenges that currently prevent their real-world deployment.

We do so with reference to the standard development pipeline of DL-based RFF systems, which includes: *RF Data Collection and Preprocessing*, *Training*, and *Deployment*.

For all identified challenges and each phase, we provide explanations and examples grounded in the literature. Addressing such challenges can be a springboard for DL-based RFF deployment in real-world scenarios and further future research. Essentially, we bridge the gap between research and practice towards DL-based RFF deployment.

Our contribution is not yet another survey on RFF. This paper addresses researchers in Industry and Academia approaching the RFF domain, providing a tutorial-style discussion of the aspects connected to both the DL and the embedded systems areas and their intersection.

The remainder of this paper is organized as follows. Section II introduces DL-based RFF systems. Section III highlights key considerations required for both the stage of data collection and preprocessing as well as the training, and Section IV discusses deployment challenges. We highlight future research opportunities in Section V, before we conclude in Section VI.

II. PRELIMINARIES ON DL-BASED RFF

RFF techniques identify wireless devices based on their unique signal characteristics. As shown in Fig. 1, RF devices emit unique waveforms in the radio spectrum due to manu-

facturing imperfections in various components, e.g., digital-to-analog converters, mixers, power amplifiers, antennas, etc. Once captured by a receiver, these signals can be represented in various formats and then processed by a DL model to identify and extract subtle variations per device. These imperceptible variations inherent to each signal constitute the signature of the device and serve as the foundation for authentication. The process of designing and implementing DL-based RFF systems typically involves three essential phases, known as the DL pipeline: (i) *data collection and pre-processing*, (ii) *training*, and finally (iii) *deployment* (see Fig. 2).

Data collection and Pre-processing. RF signals emitted by wireless devices within a monitored environment are captured and stored as In-phase (I) and Quadrature (Q) samples. This format preserves information about the amplitude and phase distortions. After acquisition, the acquired data are prepared for model training through a series of preprocessing techniques, including, e.g., feature selection, data augmentation, normalization, and noise removal. In the context of the RFF building pipeline, data preprocessing is particularly critical given that data are collected from the wireless channel, where RF signals might experience significant unpredictable and time-varying distortion due to phenomena such as multipath propagation, interference, shadowing, and fading.

Training. This phase involves designing and training a DL architecture capable of effectively capturing the underlying patterns corresponding to the unique RF fingerprints of wireless devices. Factors such as architecture selection, hyper-parameter tuning, processing variable input lengths, and designing proper validation and testing metrics are crucial for building robust and accurate RFF systems. Throughout this phase, the model undergoes iterative fine-tuning and rigorous evaluation to ensure accuracy, generalization, and adaptability.

Deployment. The trained model is integrated into a production system to perform the authentication of wireless devices. In this context, authentication is achieved by comparing incoming RF fingerprints with those learned during training.

The RFF system should efficiently and reliably authenticate wireless devices under various environmental conditions and deployment scenarios. However, various issues, such as system scalability, adaptability to new conditions, and robustness to adversarial attacks, affect the system’s performance.

III. DATA COLLECTION, PRE-PROCESSING, AND TRAINING

Training Dataset Collection Settings. Existing techniques to collect RF fingerprints typically involve generating short, intermittent snapshots of device wireless data over several days [7], [9]. The aim is to capture the RF fingerprint across varying wireless channel conditions and ensure a manageable dataset size for efficient training. However, these techniques often overlook additional factors that could impact the devices’ actual RF fingerprint, including effects from power cycles, temperature variations in internal circuits, and device aging [10]. In real-world scenarios, RF devices operate continuously, posing a challenge in developing data collection methods that can effectively capture the dynamic nature of the device fingerprint.

Input Data Selection. Each segment of the transmitted wireless signal possesses unique characteristics, potentially useful for RFF. However, selecting the optimal data segment as input to the RFF model can be challenging. Ideally, the chosen data segment should show a consistent and repetitive pattern. In fact, this can guarantee that the DL model learns from the unique patterns of the RF signal rather than becoming biased towards the specific content of the wireless segment. The preamble of the wireless packet, containing synchronization-related information, is particularly appealing as it is consistent nature across different devices and packets (for the same communication technology). Overall, selecting the optimal wireless signal segment for RFF is still an open challenge that requires further research and investigation.

Data Augmentation. Refers to a set of established transformations applied to the original data to diversify it while preserving its statistical properties. The main objective is to increase the size of the data set and improve the generalization and robustness of the model. For example, in computer vision, augmentation techniques such as flipping, rotation, scaling, and cropping are used successfully to expand the data and mitigate overfitting. However, in the context of RFF, identifying universally effective techniques to enhance wireless data remains challenging.

This issue arises due to the complexity and dynamic nature of wireless signals, which are subject to channel conditions and hardware-based effects.

Signal Representation. Wireless signals can be represented in various forms using different signal processing methods, each capturing different characteristics of the signal. The features extracted from these analyses significantly impact the performance of the DL model. Representations such as raw IQ data, spectral characteristics, and the combination of time- and frequency-related metrics are often considered [4]. Raw IQ data capture amplitude and phase variations caused by

hardware imperfections. Spectral analysis extracts frequency-domain characteristics of the RF signal. Furthermore, the time-frequency representations capture the combined time-varying frequency behavior of the RF signal. Since each representation has its own strengths and weaknesses, selecting the most suitable one for training RFF models is not straightforward and requires further research.

Feature Normalization. Wireless data is inherently dynamic and subject to time-varying channel conditions, hardware imperfections, and noise, all of which can introduce power fluctuations and scaling issues in the received data. Given that DL models are scale-sensitive, these inconsistencies, if not addressed, can negatively affect the models’ learning and generalization. For example, when training with raw IQ data, significant scale differences between the I and Q components may lead the model to assign greater weight to the component with a larger scale, resulting in a scale-dependent solution for device identification. By equalizing all wireless data feature ranges (normalization), DL models can correctly capture the underlying pattern of the data distribution and reduce potential biases towards signal power or feature scale.

Neural Network Structure & Hyper-parameter Tuning. The DL architecture is crucial for the overall accuracy and robustness of RFF systems. Many available solutions reuse or modify existing DL networks, reporting outstanding performance in other domains, such as computer vision and Natural Language Processing (NLP) [4]. However, these models, initially designed to capture domain-specific features, cannot be expected to perform well with wireless data. Therefore, it becomes imperative to develop specialized DL models to deal with the complexity of wireless data and adapt to its inherently dynamic nature. Fig. 3a compares neural network architectures trained on preamble raw IQ data with a defined sequence. To assess model robustness against temporal changes, the sequence of the preamble data is randomized during testing. While all models achieved high training accuracy, only the Transformer architecture (*without positional encoder*) demonstrated strong generalization on both training and testing data, regardless of the raw IQ data sequence. Furthermore, the configuration of DL models with different hyper-parameter values can affect their performance during training and deployment. Suboptimal hyper-parameter selection might cause overfitting or underfitting, resulting in degraded performance. For example, analyzing the impact of batch size and epochs, Fig. 3b illustrates how varying the batch size affects the model learning. Smaller batch sizes achieve higher accuracy and converge faster but are more computationally demanding, whereas larger batches converge more slowly but demonstrate steady progress over more epochs.

Design of a Proper Validation/Testing Methodology. While conventional validation techniques like *Holdout* or *K-fold cross-validation* can be utilized to evaluate DL models with wireless data, the dynamic nature of wireless data requires special considerations. Wireless channels frequently experience time-varying conditions, e.g., fading, multipath

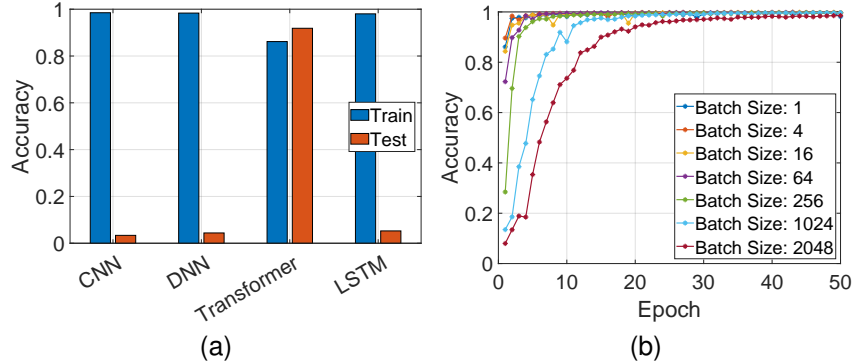


Fig. 3. Impact of DL architecture and hyper-parameter tuning on model performance based on the analysis of the dataset released in [5]. (a) Comparison of different DL architectures, including Convolutional Neural Networks (CNNs), Dense Neural Networks (DNNs), Transformers, and Long Short-Term Memory Networks (LSTMs), trained on fixed, consistent raw IQ sequences. The tests were performed on randomly shuffled raw IQ sequences. (b) Validation accuracy as a function of the number of epochs for various batch sizes.




Fig. 4. Visualization of feature embedding for a reference DL-based RFF model trained to identify 143 transmitters, based on our own analysis of the dataset released in [5]. (a) Cosine similarity analysis highlights the degree of similarity among features. (b) Projection using Principal Component Analysis (PCA) shows the distribution variation and clustering of similar transmitters. (c) t-Distributed Stochastic Neighbor Embedding (t-SNE) visualization of the transmitters' features, providing a dimensionality-reduced view of the crowded feature space and its overlapping regions. The colors in (b) and (c) represent individual transmitters.

propagation, and Doppler effects. These elements change the data distribution over time, making it difficult to guarantee that the validation and testing sets are representative of future data [9]. Moreover, spatial variability introduces another level of complexity, as specific physical environments can greatly influence wireless data. Consequently, it is imperative to validate and test the model on diverse data captured from different locations and periods, as this will provide an accurate assessment of its performance under real-world conditions.

IV. DEPLOYMENT CHALLENGES

Model Interpretability Issues. Although DL techniques often show remarkable performance compared to other Machine Learning techniques, their *blackbox* nature and lack of interpretability are significant drawbacks. In particular, it is difficult to explain how these models classify the data and which specific underlying features contribute to the final decision produced by the model. This issue becomes even more relevant when considering that DL-based RFF systems are deployed for security applications, such as device authentication and intrusion detection. Given the security implications, a transparent

understanding of the rationale of the decision-making process of these models is even more critical to mitigate the risk of false predictions while maintaining robustness and reliability.

Lack of Benchmark Dataset. Standardized benchmark datasets for computer vision and NLP provide a common ground for comparing the performance of DL models in those domains. On the contrary, for RFF, developing a widely accepted benchmark dataset remains a challenge. To overcome this, researchers often resort to locally generated datasets for model validation. Although such datasets could be practical for initial model testing, they may only represent a subset of real deployment conditions.

This bias may affect the robustness and generalizability of the model when deployed in real-world scenarios. In general, creating a comprehensive and standardized RFF benchmark dataset is crucial to evaluate deployed DL-based RFF models across various modulation schemes, transmission rates, communication technologies, and channel conditions.

Receiver Hardware Bias. Most existing RFF solutions collect data using a single receiver to identify a specific wireless transmitter within a pool of devices. Using this

technique, it is commonly assumed that the receiver hardware does not introduce its own variability to the captured RF fingerprint. Additionally, the receiver hardware used during training is typically expected to remain unchanged during deployment. In reality, the RF fingerprint is affected by the entire communication chain, including the transmitter, the channel, and finally, the receiver [9], [11]. Therefore, the resulting trained model is significantly biased towards the physical characteristics of the receiver used during training. With this limitation, any change at the receiver—where the model has been deployed—would require retraining the model. This is paramount for any generalized deployment, with a trained model for one transmitting device, \mathcal{V} , distributed to many (or any) devices that seek to identify/authenticate \mathcal{V} .

Scalability Issues. DL-based RFF systems are designed to be deployed in large-scale network environments with numerous devices. However, as the number of devices increases, the performance of the DL-based RFF might degrade [5]. This degradation is mainly due to the difficulty in distinguishing devices with similar characteristics. For instance, in a network employing BPSK modulation, using raw IQ data for device authentication could introduce scalability issues. Specifically, the presence of limited features (strong I and nearly negligible Q) combined with a densely populated feature space can lead to overlapping features from similar devices, resulting in performance degradation. Fig. 4 presents a visualization of the feature embedding and cosine similarity for a model designed to identify 143 wireless transmitters when analyzing the dataset in [5]. As shown in Fig. 4a, cosine similarity reveals areas of high similarity between the transmitter features. Moreover, Fig. 4b and 4c reveal a dense feature space and many overlapping regions, indicating the potential challenge for the model to distinguish transmitters. Overall, improving the scalability of DL based RFF systems requires developing novel solutions, which may involve improved feature engineering and model architecture design.

Protocol Dependency. Current DL-based RFF systems are inherently protocol-dependent and sensitive to protocol configuration, which in turn affects their device identification capabilities [8]. For example, WiFi and LoRa define a range of parameters, including modulation schemes, bit rates, spreading sequences (in the case of LoRa), allowed frequencies, bandwidth, and transmit power, which can be fine-tuned based on operational conditions. Changes in these parameters affect the acquired RF fingerprint, leading to a decrease in overall system performance. This performance drop is especially noticeable when training and testing data are collected under different protocol configurations.

Detection of Unauthorized Devices. Most existing DL-based RFF systems operate under closed-set classification conditions, i.e., they only recognize devices that were included in their training set. While this approach works well in controlled environments, it presents a significant challenge when new, previously unseen devices transmit data. In such a scenario, closed-set DL-based RFF systems cannot reliably recognize new devices, primarily because they have not encountered

their unique RF fingerprints during the training phase. As a result, the system compares the new unseen devices based on the one(s) that closely match the RF fingerprint from its training set. To overcome this challenge, it is crucial to develop open-set classification techniques (e.g., zero-shot learning) that can identify, handle, and reject new unauthorized devices during the deployment phase, thus improving the security and adaptability of the system.

Simultaneous Transmissions. During training, RF profiles of transmitting devices are usually acquired by isolating a specific transmitter device that operates on a single frequency channel and collecting data with a receiver [12]. However, in real-world RFF environments, multiple devices may transmit simultaneously, either on the same frequency or across different channels. Such concurrent transmissions can degrade the performance of DL-based RFF systems due to signal overlap at the receiver. Furthermore, any change in the training channel profile can further affect performance, as the system may encounter unfamiliar channel characteristics.

Deployment on Constrained Devices. Most current DL-based RFF studies focus on developing RFF systems intended to be deployed in computing environments with ample processing power, memory, and energy. However, in many real-world scenarios, deploying these computationally intensive models on lightweight devices is often impractical. This is particularly relevant for scenarios such as edge computing and resource-constrained IoT devices. Therefore, more research is required to explore optimization techniques that allow RFF models to run efficiently on lightweight devices without compromising performance.

Variable Input Length. Wireless signals can vary in size and duration, depending on the specific communication technology and modulation scheme [13]. This variability is particularly evident in adaptive communication protocols such as LoRaWAN, where the length of transmitted symbols varies in response to network conditions and proximity. This is challenging for conventional DL models, designed to handle fixed-size input data. These models require the input data to be truncated or padded to fit the input size, which can potentially lead to information loss and increase the computational resources to compensate for missing data. However, the most critical implications lie in the negative impact on overall model performance.

System Adaptiveness. Existing DL-based RFF systems typically operate under the assumption that the set of authorized devices present during the training phase remains the same at the time of the deployment phase. However, real-world wireless networks are often dynamic, and devices frequently join and leave the network. The current solution to this challenge involves retraining the system model whenever there is a change in the network—a computationally expensive and time-consuming process. This issue becomes particularly critical in large critical infrastructures, where updates must be executed seamlessly without disrupting regular operations.

Inference Time. In real-world settings, the time to perform the tests, namely the inference time, is a key factor

in determining the suitability of DL-based RFF systems to achieve their objective. Short inference times are particularly important in time-sensitive applications, where delays cannot be tolerated. To ensure the practicality of DL-based RFF systems, it is essential to evaluate and optimize the model inference time across a range of devices that could potentially be deployed within the specific network. This evaluation and optimization is necessary, given that the inference time of DL models varies based on the hardware components in which they are installed.

Robustness to Channel Variations. DL-based RFF systems rely on data collected from the physical layer of the wireless spectrum. However, these data are subject to strong channel fluctuations, noise, multipath, and shadowing effects caused by environmental changes. These channel variations can significantly affect the resulting profile of RF signals, making the channel conditions experienced during training often different from those experienced during the deployment phase. Eventually, this variability leads to performance degradation [9].

Jamming. Jamming attacks occur when a malicious entity intentionally transmits a signal on the same frequency band as the target communication link, disrupting ongoing communication. In our context, jamming affects the reception of the clean RF signal, making the extraction of RFF features challenging. As mentioned above, this challenge is exacerbated as the RF signal is already affected by other noise sources present in the wireless channel. As a result, jamming attacks can significantly compromise the reliability of the system.

Spoofing. The openness of wireless communications allows malicious entities to eavesdrop on radio signals and inject new signals to gain unauthorized access. In the RFF context, an attacker can capture the unique RF fingerprint of a target device, emulate or distort the signal to create a deceptive one, and then transmit it to try to impersonate the legitimate device. More research is needed to validate the possibility of this adversarial scenario and test the resilience of DL-based RFF systems against this attack.

Adversarial Attacks. DL models are vulnerable to adversarial attacks [14]. For example, in image classification tasks, malicious entities can trick the model by manipulating the input imperceptibly to the human eye, but can fool the model into making wrong predictions. In the context of RFF, adversarial attacks can significantly impact the performance and reliability of DL-based RFF systems. By exploiting the sensitivity of the model to input perturbations, an attacker can carefully craft deceptive adversarial samples to compromise the integrity of the system. At the time of this writing, only a limited number of studies have investigated the robustness of DL-based RFF systems to such attacks, and only under ideal conditions, without considering the requirements necessary for executing successful attacks.

V. RESEARCH OPPORTUNITIES

Realistic Channel Simulation. One promising approach to mitigate channel variability involves simulating RF fingerprint data under various realistic wireless channel conditions, with

data taken from real data wireless channels [15]. Although this research direction is still in its early stages, it holds significant potential, as it exposes the signal to real-world deployment conditions and introduces variability and augmentation to the data before training.

Continual Learning. Another effective strategy might involve integrating continuous adaptive learning into their deployment. By adopting this approach, these systems can continuously update in real-time without being re-trained, allowing them to maintain accuracy and adapt to the highly dynamic nature of wireless channels while being deployed.

Collaborative Learning and Multimodal Approach. To improve the reliability and accuracy of DL-based RFF systems, researchers can investigate the concept of multimodal and collaborative learning where different types of input data or model predictions are fused together to strengthen the overall system's performance.

RFF Anonymization. Once the RF device fingerprint is leaked, it can be used against the targeted device. Specifically, malicious entities can exploit it for activities such as device tracking, profiling, and spoofing. Consequently, developing privacy-preserving RFF techniques is crucial to protect individual privacy and data while ensuring the usability of these DL-based RFF systems.

VI. CONCLUSION

Research on reliable and robust DL-based RFF systems has progressed significantly in recent years, yielding effective solutions for various scenarios and applications. However, deploying RFF systems in real-world settings still faces numerous challenges inherent to DL techniques and the highly dynamic nature of wireless scenarios. This paper has systematically identified and analyzed the most significant considerations and challenges towards deploying real-world DL-based RFF systems, categorized per phase in the development pipeline. We also outlined promising future research opportunities. The aim is to provide researchers and industry experts with succinct guidance on the main challenges and assist efforts toward real-world RFF deployment.

ACKNOWLEDGEMENT

This research was made possible by the award GSRA7-1-0510-20045 and NPRP12C-0814-190012-SP165 from Qatar National Research Fund (a member of Qatar Foundation) and by the INTERSECT project, Grant ID NWA.1162.18.301, funded by Netherlands Organisation for Scientific Research (NWO). The work of P. Papadimitratos and A. Hussain was supported in part by the Swedish Science Foundation (VR). The contents herein are solely the responsibility of the author(s).

REFERENCES

- [1] Soltanieh, N. et al., "A Review of Radio Frequency Fingerprinting Techniques," *IEEE Journal of Radio Frequency Identification*, vol. 4, no. 3, pp. 222–233, 2020.
- [2] Oligieri, G. et al., "PAST-AI: Physical-Layer Authentication of Satellite Transmitters via Deep Learning," *IEEE Transactions on Information Forensics and Security*, vol. 18, pp. 274–289, 2023.

- [3] LeCun, Y. et al., "Deep Learning," *nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [4] Jagannath, A. et al., "A Comprehensive Survey on Radio Frequency (RF) Fingerprinting: Traditional Approaches, Deep Learning, and Open Challenges," *Computer Networks*, p. 109455, 2022.
- [5] Hanna, S. et al., "WiSig: A Large-Scale WiFi Signal Dataset for Receiver and Channel Agnostic RF Fingerprinting," *IEEE Access*, vol. 10, pp. 22 808–22 818, 2022.
- [6] Rehman, S. et al., "Radio Frequency Fingerprinting and its Challenges," in *IEEE Conf. on Commun. and Network Security*, 2014, pp. 496–497.
- [7] Elmaghub, A. et al., "Comprehensive RF Dataset Collection and Release: A Deep Learning-Based Device Fingerprinting Use Case," in *IEEE Globecom Wkshps.*, 2021, pp. 1–7.
- [8] —, "LoRa Device Fingerprinting in the Wild: Disclosing RF Data-Driven Fingerprint Sensitivity to Deployment Variability," *IEEE Access*, vol. 9, pp. 142 893–142 909, 2021.
- [9] Al-Shawabka, A. et al., "Exposing the Fingerprint: Dissecting the Impact of the Wireless Channel on Radio Fingerprinting," in *IEEE INFOCOM* 2020, 2020, pp. 646–655.
- [10] Alhazbi, S. et al., "The day-after-tomorrow: On the performance of radio fingerprinting over time," in *Proceedings of the 39th Annual Computer Security Applications Conference*, 2023, pp. 439–450.
- [11] Merchant, K. et al., "Toward Receiver-agnostic RF Fingerprint Verification," in *2019 IEEE Globecom Workshops (GC Wkshps)*. IEEE, 2019, pp. 1–6.
- [12] J. A. Gutierrez del Arroyo, B. J. Borghetti, and M. A. Temple, "Considerations for radio frequency fingerprinting across multiple frequency channels," *Sensors*, vol. 22, no. 6, p. 2111, 2022.
- [13] Shen, G. et al., "Toward Length-Versatile and Noise-Robust Radio Frequency Fingerprint Identification," *IEEE Transactions on Information Forensics and Security*, vol. 18, pp. 2355–2367, 2023.
- [14] Goodfellow, I. et al., "Explaining and Harnessing Adversarial Examples," *arXiv preprint arXiv:1412.6572*, 2014.
- [15] Yan, W. et al., "RRF: A Robust Radiometric Fingerprint System that Embraces Wireless Channel Diversity," in *Proc. of ACM WiSec*, 2022, pp. 85–97.