ADL-ID: Adversarial Disentanglement Learning for Wireless Device Fingerprinting Temporal Domain Adaptation

Abdurrahman Elmaghbub, Bechir Hamdaoui and Weng-Keen Wong
School of Electrical Engineering and Computer Science, Oregon State University
{elmaghba, hamdaoui, wongwe}@oregonstate.edu

Abstract—As the journey of 5G standardization is coming to an end, academia and industry have already begun to consider the sixth-generation (6G) wireless networks, with an aim to meet the service demands for the next decade. Deep learning-based RF fingerprinting (DL-RFFP) has recently been recognized as a potential solution for enabling key wireless network applications and services, such as spectrum policy enforcement and network access control. The state-of-the-art DL-RFFP frameworks suffer from a significant performance drop when tested with data drawn from a domain that is different from that used for training data. In this paper, we propose ADL-ID, an unsupervised domain adaption framework that is based on adversarial disentanglement representation to address the temporal domain adaptation for the RFFP task. Our framework has been evaluated on real LoRa and WiFi datasets and showed about 24% improvement in accuracy when compared to the baseline CNN network on short-term temporal adaptation. It also improves the classification accuracy by up to 9% on long-term temporal adaptation. Furthermore, we release a 5-day, 2.1**TB, large-scale WiFi 802.11b dataset collected from 50 Pycom devices to support the research community efforts in developing and validating robust RFFP methods.**

Index Terms—RF Fingerprinting, Domain Adaptation, Disentanglement Learning, WiFi Dataset, IoT Testbed.

I. INTRODUCTION

An important highlight of evolving next-generation wireless networks is security and privacy, which have been somehow overlooked, to some extent, in previous generations [1]. While traditional security relies heavily on the upper-layer mechanisms, it is essential to complement these mechanisms with unclonable, physical-layer security technologies to increase the security robustness of such systems. As an efficient physical-layer security mechanism, RF fingerprinting (RFFP) technology allows for device identification based on unique fingerprints inherited in the RF signal. Due to the imperfection of the manufacturing process, circuit components in the signal path introduce a variety of random hardware impairments such as I/Q mismatch, phase noise, and non-linearity [2] that collectively impact the transmitted signal, constituting a unique RF Fingerprint (RFF). Deep Learning (DL) has reduced the need for preprocessing and expert-defined feature extraction techniques, building state-of-the-art RFFP frameworks [3], [4]. However, most of these rely on the assumption that the training and testing sets are drawn from the same distribution, which falls short of the conditions of realistic RF scenarios.

We define domain portability as the ability of a model to maintain its training domain performance when applied to new domains. In the RFFP context, the domain refers to the settings in which data has been collected. This includes time, channel, receiver, protocol, and environment. Hence, any considerable change in the training settings yields a different domain. In [5], [6], we conducted an experimental study on LoRa devices, which disclosed the sensitivity of DL-RFFP to domain changes. In that work, the models (i) fail to maintain their high accuracy when channel conditions change and (ii) completely lose their classification ability when the protocol configuration or receiver changes. This lack of portability represents a substantial impediment to the implementation of DL-based RFFP methods in real-world applications.

The complexity of the RFFP problem comes from the fact that we are dealing with two layers of obscurity: the fingerprint interactions in both circuit and channel levels and the feature vector learned by the deep learning network. The lack of behavioral comprehension of RFFs and deep learning networks opens up the floor to several hypotheses to explain the exposed sensitivity of deep learning models when the domain changes. It could be that the fingerprints themselves change as the domain changes; hence, identifying devices based on their old/learned fingerprints yields poor accuracy in cross-domain validation. This assumption disqualifies these features as fingerprints since authentication schemes require them to be robust and permanent. More reasonably, the feature vector learned by the network includes both RFFs and a domain-specific component which leads to confusion when the latter component changes as the domain changes. Most of the proposed solutions adapted the latter assumption in their endeavors. Some works have focused on eliminating the background effects from the raw signal to improve its domain generalization. One such example is the use of channel equalization to improve the performance by eliminating channel influences [7], but this is a partial solution since they eliminate many beneficial spurious emissions in the passband, which reduces the user capacity of the RFFs [8]. On the other hand, DeepLoRa [9], a data augmentation technique based on the ITU-R channel models, was proposed to counter the degradation introduced by the wireless channel. The best performance reported by DeepLoRa was 36% with only 5 devices being classified, achieved when the model is trained on one day and tested on another day. This accuracy drops to 8% when the number of devices goes to 50.

More recently, the portability problem of the RFFP has been

formulated as a domain adaptation problem [10], [11], allowing us to leverage the advancements of this branch of transfer learning in the RF context. This formulation assumes that the source and target domains have slightly different distributions. In [12], a multi-discrepancy domain adaptation method has been proposed to minimize the intraclass discrepancy across the source and target domains while maximizing the inter-class discrepancy, leading to 15% improvement over the ResNet network. Adversarial domain adaptation [13] has also been proposed to reduce the influence of domain features on the RFFs by maximizing the domain discrimination loss, resulting in an enhancement in the portability of the network. However, the target domains on the two mentioned domain adaptation methods have been created via simulation by introducing white Gaussian noise to the source domain dataset to create several target domains with different signal-to-noise ratio conditions. These simulation settings are far from realistic RF scenarios where multiple settings change at the same time. Addressing multiple varying settings, Tweak [11] combined the metric learning with a lightweight calibration method to address the portability problem in hardware, channel, and configuration dimensions, separately. Tweak shows a considerable performance improvement in each of these dimensions, but needs to be recalibrated each time the target domain changes.

Another exciting domain adaptation category is disentangled representation learning, an unsupervised learning technique that separates (disentangles) features into disjoint parts of the underlying representation [14]. Motivated by the intuition that the RFFs are entangled with domain-specific features and inspired by the disentanglement concept [15], we propose ADL-ID, a domain adaptation framework that integrates disentanglement representation learning with adversarial learning to address the temporal portability problem in RFFP. Our framework factors the feature vector into a device-specific component (RFFs) and a domain-specific component. Then, it uses only the device-specific component, which contains domain-invariant features, as input to the classifier. Our proposed framework, tested using real WiFi dataset (detailed in Section II) and LoRa datasets (described in [16]), provides a high-performance gain, $20\% - 24\%$, in short-term-varied domain adaptation. We are also excited to release our large-scale WiFi 802.11b dataset of 50 Pycom devices to the research community to support the efforts toward enabling RFFP portability. This is the first public WiFi 802.11b RFFP dataset, which is common among low-rate IoT devices. Our testbed includes the same devices used before to build and release the LoRa dataset [16], which extends the opportunity to analyze other domain dimensions, such as protocol configurations and hardware receivers. Our main contributions are as follows:

- Proposed an adversarial-based disentanglement learning network to address the RFFP temporal portability issue.
- Released the first large-scale WiFi 802.11b RFFP dataset of 50 Pycom devices over 5 consecutive days for both indoor and outdoor scenarios.
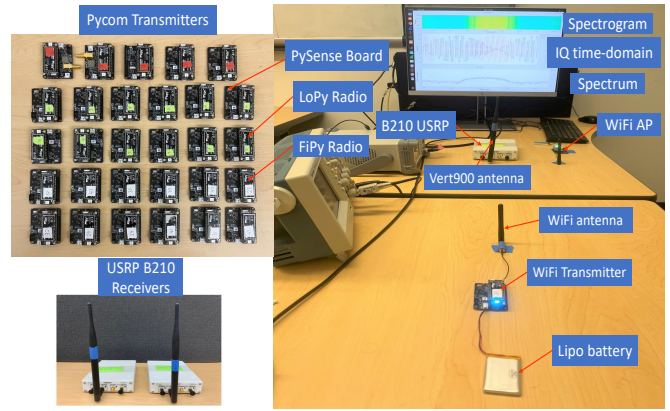- Provided novel insights into the performance of RFFP

Fig. 1: 50-device Pycom WiFi testbed

methods due to short-term and long-term temporal domain variations using both WiFi and LoRa datasets.

The rest of the paper is organized as follows. Section II describes the device testbed and WiFi dataset. Following that, Section III reveals the investigation on short-term variations performance. Then, our framework is explained in Section IV while our results are presented and discussed in Section V. The paper is concluded in Section VI.

II. TESTBED AND WiFi DATASET DESCRIPTION

In this section, we begin by describing the hardware, software, and protocol components of the tested used for collecting our dataset. Then, we describe the collected WiFi dataset. The description and download information of the LoRa dataset that we also used in our evaluation can be found in [5], [16].

A. Testbed Description

1) Hardware: Our testbed, shown in Fig. 1, consists of 50 Pycom devices (25 FiPy & 25 LoPy boards on top of 50 PySense extension boards). These boards mainly include ESP32-D0WDQ, Semtech SX1276, and Sequans Monarch chips to support WiFi, Bluetooth, LoRa, Sigfox, and LTE networks, respectively. On the reception side, we used an Ettus USRP (Universal Software Radio Peripheral) B210 with a VERT900 antenna for the data acquisition. The RF-front of the USRP was configured with a center frequency of 2.412GHz, a 20MHz Bandwidth, and a 25MSps sample rate. We used lipo batteries to power the Pycom devices. Each Pycom device was connected to a dedicated WiFi antenna and configured to transmit IEEE802.11b packets in the 2GHz spectrum.

2) Software: Pycom devices are based on FREERTOS operating system. To perform transmission, we programmed and configured our Pycom boards using MicroPython, an efficient implementation of Python3 composed of a subset of standard Python libraries and optimized to run on microcontrollers and constrained environments. Also, we used Pymakr plugin as a REPL console that connects to Pycom boards to run codes or upload files. To perform reception, we used the GNURadio software, a real-time signal processing graphical tool, along with USRP Hardware Driver (UHD) to set up and configure

the USRP receiver to capture WiFi transmissions, plot their time and frequency domains, and store the collected samples in their files.

3) Protocol: We used the WiFi 802.11b protocol with the high-rate direct-sequence spread-spectrum (HR/DSSS) physical layer. Transmitters were configured to transmit at 1Mbps, with a carrier frequency of 2.412GHz and 20MHz bandwidth.

B. Dataset Description

Our WiFi dataset contains 2.1TB of WiFi transmissions of 50 Pycom devices captured over 5 consecutive days in both indoor (laboratory) and outdoor environments. For each day, we capture 5 transmissions from each device in a round-robin fashion, where each capture consists of 50M complex-valued I/Q samples. The time-gap between two consecutive captures of the same device is 5mins. Transmitters were located about 5 meters away from the access point and the receiver, with a clear line of sight. The collected signals are down-converted by the USRP to the baseband and then stored in the form of I/Q data in both time and frequency domain representations. For each capture, associated metadata compliant with the latest SigMF specifications and additional field extensions for usability are recorded.

More details and use cases of the WiFi dataset can be found in [5]. The datasets can be downloaded from NetSTAR Lab at http://research.engr.oregonstate.edu/hamdaoui/datasets.

III. THE TEMPORAL PORTABILITY PROBLEM

Several works [5], [9], [17] consider the days as the change unit in the temporal dimension of the domain portability problem. Hence, they analyze the performance of the models when the training and testing data are from different days. However, there is nothing special about the 24 hours as far as the RF fingerprints are concerned. We postulate that the data distribution of fingerprints changes dramatically even within the same day. To validate this hypothesis, we used our LoRa and WiFi datasets in which we have several captures of the same device within the same day. Starting with the LoRa dataset [16], the time gap between two consecutive captures is 5 minutes with 5m distance between transmitters and the receiver. The captures were collected in indoor and outdoor scenarios, and the devices have not been rebooted between the captures. We trained a 6-Layer CNN network [5] using capture 1 and tested it on data from captures 1, 2, 3, and 4 of the 4 days. Surprisingly, the model's performance suffers from a huge degradation even when the time gap between the training and testing sets is as short as 5 minutes. Furthermore, as we can see in Fig 2—depicting accuracy results when using the LoRa data, the performance continues to degrade as we move further in time from the training domain. This trend manifests itself in the four days and both scenarios.

To confirm our observation, we carried out the same investigation in our 50-device WiFi dataset, which also includes 5 captures of the same devices within the same day for 5 consecutive days, where the time gap is similar to that in the LoRa dataset. The difference is that the devices here have

(a) LoRa Indoor Scenario.
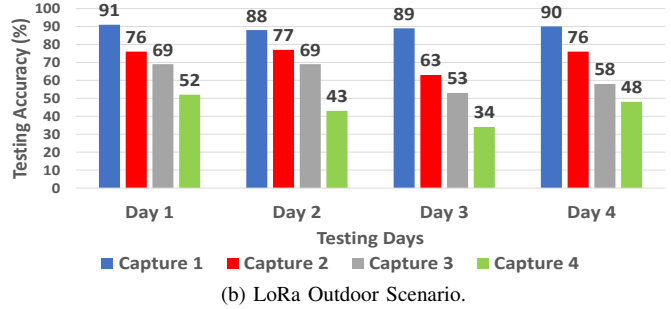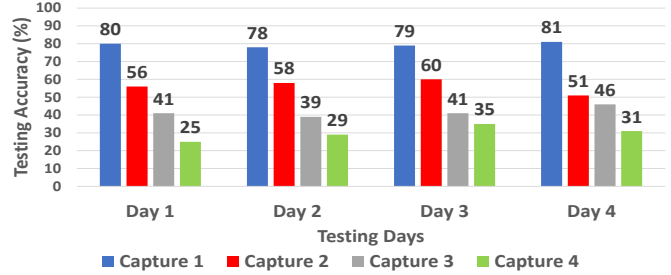

(b) LoRa Outdoor Scenario.

Fig. 2: Performance in short-term variations of LoRa captures where the time gap between two consecutive captures is 5mins. The CNN model was trained on captures 1 of each day and tested on captures 1, 2, 3, and 4.
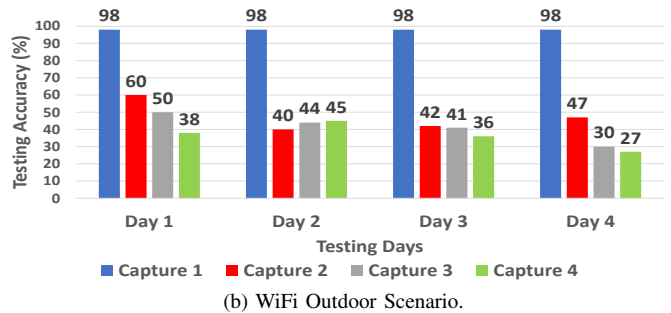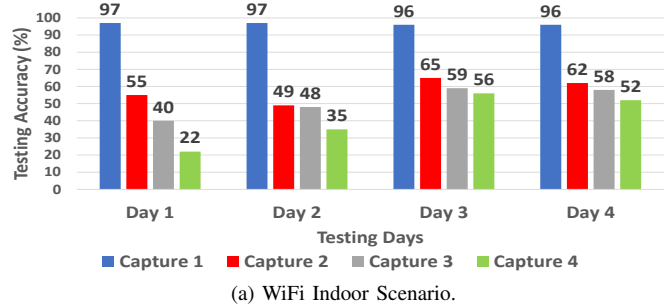

(a) WiFi Indoor Scenario.

(b) WiFi Outdoor Scenario.

Fig. 3: Performance in short-term variations of WiFi captures where the time gap between two consecutive captures is 5mins. The CNN model was trained on captures 1 of each day and tested on captures 1, 2, 3, and 4.
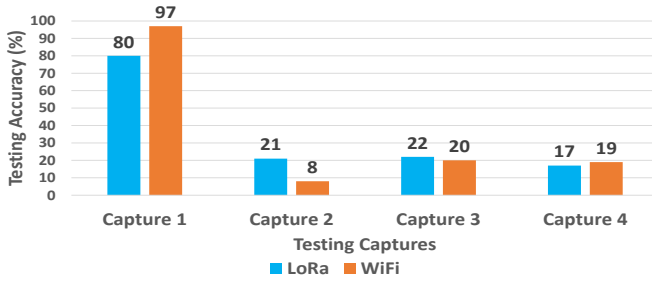
Fig. 4: Performance in long-term variations on LoRa (1st column) and WiFi (2nd column) captures where the captures are collected on 4 different days. The CNN model was trained on capture 1 and tested on captures 1, 2, 3, and 4.

been rebooted between the captures. Results in Fig. 3 show that the WiFi dataset is no different, and the model also suffers from short-term variations on the WiFi dataset. These results consolidate our hypothesis that even a slight departure from the network's training domain can introduce confusion to the classifier and hurt its performance severely.

We also notice that on most of the days, the most significant drop rate in performance occurs in the transition from capture 1 to capture 2, and that rate decreases as we move away from the training domain. We claim that the drop rate decreases until it reaches a point when the time gap does not matter anymore; i.e., aging has no to little effect. Hence, with respect to domain adaptation performance, 2-day or 10-day time separation between the source and target domains has a similar effect. Some of that can be seen in Fig. 4, which indicates that the effect on the different days is similar. We refer to this effect as the effect of long-term variations. As expected, Fig. 4 shows that the drop in the network's performance due to the long-term variations is more severe than the previous scenario as the accuracy drops from around 90% to 20%, when the model is trained on capture 1 and tested on captures 2, 3, and 4, which have been collected on day 2, 3, and 4, respectively.

IV. THE PROPOSED FRAMEWORK: ADL-ID

The proposed framework, shown in Fig. 5, aims to train an encoder, called fingerprint encoder, E_F, to extract only the device-specific feature vector (RFFs) and feed it to the cross-domain device classification task. The intuition here is that the RFFs are shared across domains while the background and domain artifacts are different from one domain to another. This is achieved by having three encoders: fingerprinting encoder, E_F, source encoder, E_S, and target encoder, E_T that are trained to model different components of the domain representation. Using labeled source domain data, $X_S = \{(x_s^i, y_s^i)\}_{i=0}^{N_s}$, and unlabeled target domain data, $X_T = \{(x_t^i)\}_{i=0}^{N_t}$, the fingerprint encoder learns to extract discriminative and shared features across the source and target domains while the source and target encoders learn the features that are specific to the source and target domains, respectively.

At each training iteration, the output of the fingerprinting encoder will be fed to the domain discriminator, $D(x)$, to eliminate domain-indicative features to make the RFFs similar to each others regardless of their original domains. We minimize the domain classification loss, L_{Domain}, while having the gradient reversal layer (GRL) [18] in the data path. GRL acts as an identity transformation in the forward path while reversing the gradient's sign in the backpropagation. Hence, it maximizes the confusion of the domain discriminator as the updated weights of E_F result in RFFs with less domain-indicative information. To enhance the separation between the RFF and the domain-specific representations of an input frame as the output of E_F and E_S or E_T encoders, we apply a difference loss, $L_{difference}$ that encourages the orthogonality of the two representations and pushes them further away from each other by minimizing the squared Frobenius norm of the multiplication of two matrices H_f and H_s or H_t whose rows are h_f and h_s or h_t, respectively [15]. To avoid trivial solutions from the previous losses, we use a signal decoder, $G(x)$, that decodes the combination of the RFF and domain-specific features to get back the original input frame. The reconstruction process is governed by the reconstruction loss, $L_{Reconstruct}$, which is a scale–invariant mean squared error (SI-MSE) loss that experimentally proves to be more effective for disentanglement representations [19]. Finally, h_f^s, the RFF representation of X_s is passed to the main-task classifier with a cross-entropy loss, L_{Class}, that encourages the RFF representation to not just be domain-invariant and separable from the domain-specific representation, but also to be discriminative with regard to the device classification task. The total loss of our network can thus be written as

$$L_{total} = L_{Class} + \alpha L_{Reconstruct} + \beta L_{Difference} + \gamma L_{Domain}$$

where the hyperparameters, α, β, and γ, control the strength of these losses in the overall training process.

Each encoder in our framework consists of 6 convolution blocks followed by a fully-connected layer. Each convolution block contains 2D convolution, batch normalization, leakyReLu, and max pooling layers. The signal decoder has also the same components in an opposite order. On the other hand, the device classifier and domain discriminator has three and two fully-connected layers, respectively, followed by a Softmax layer. We start the training with a warm up phase of 10000 steps where only the supervised task and the reconstruction losses are used before switching to the other losses. This can help in speeding up the training process as well as leading to a better generalization [20].

V. PERFORMANCE EVALUATION AND DISCUSSION

In this section, we evaluate our proposed framework for the indoor setup of our LoRa and WiFi datasets. For each device, we used 20,000 I/Q frames each of size 2x1024, partitioned into training, validation, and testing sets. Our network has been trained for 50 epochs on an Nvidia DGX2 node. We use two testing methodologies to run our evaluation: (i) testing on same-day captures to evaluate the short-term temporal adaptation; and (ii) testing on different-days captures to evaluate the long-term temporal adaptation.
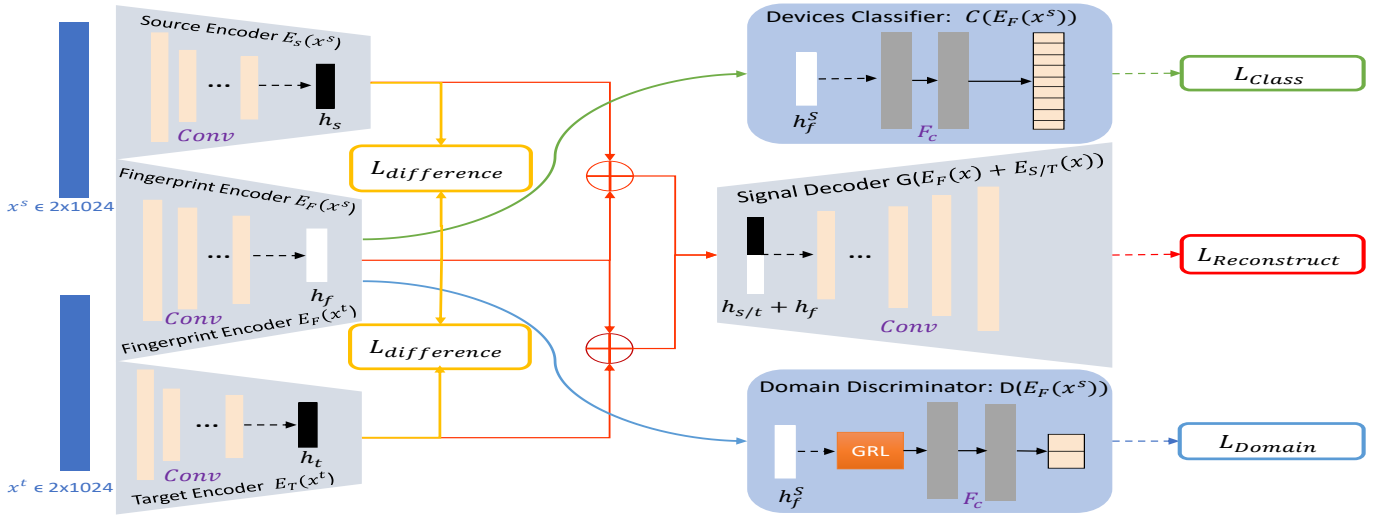
Fig. 5: Architecture of the proposed unsupervised adversarial-based disentanglement learning network for the RFFP task.
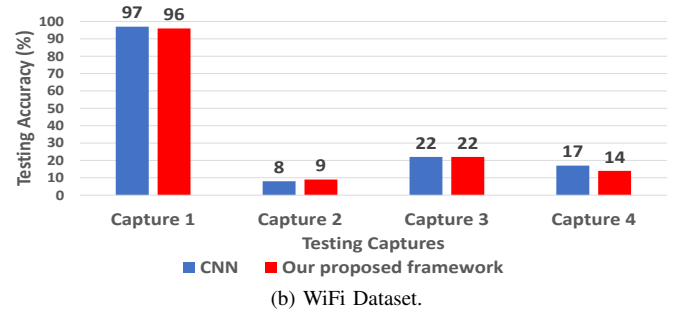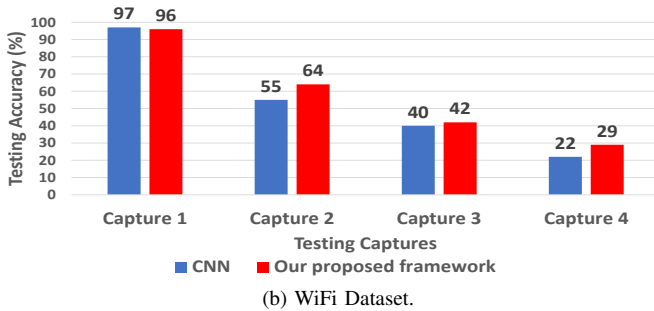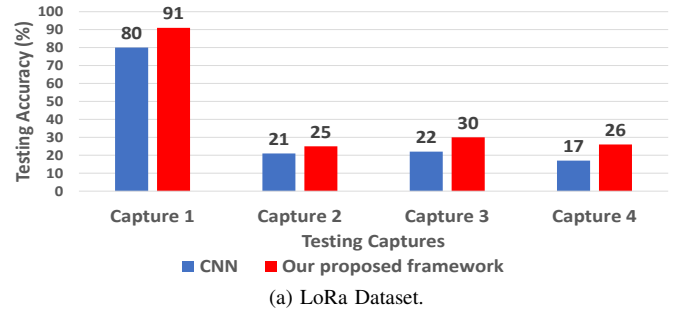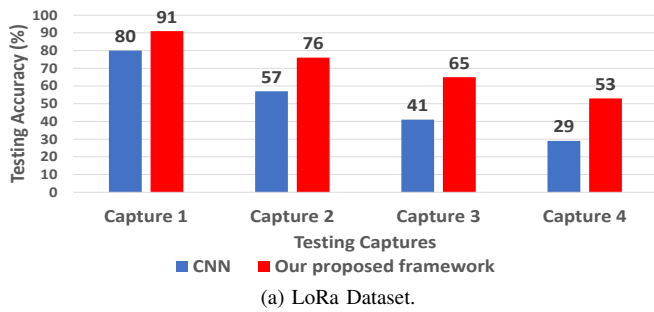

(a) LoRa Dataset.

(b) WiFi Dataset.

Fig. 6: ADL-ID vs. CNN on captures from Day 1, Indoor scenario for both (a) LoRa and (b) WiFi. ADL-ID was trained on labeled capture 1 data (source) and unlabeled capture 2 data (target) and tested on captures 1, 2, 3, and 4.

(a) LoRa Dataset.

(b) WiFi Dataset.

Fig. 7: ADL-ID vs. CNN on captures from 4 different days, Indoor scenario for both (a) LoRa and (b) WiFi. ADL-ID was trained on labeled capture 1 data (source) and unlabeled capture 2 data (target) and tested on captures 1, 2, 3, and 4.

First, we start with the short-term temporal adaption scenario. In the LoRa dataset, the time gap between the source domain and target domain is 5mins, and the transmitters have not been rebooted between the captures. Specifically, we train our framework with labeled data from capture 1 (source domain) and unlabeled data from capture 2 (target domain) and then test it on captures 1, 2, 3, and 4, all captured on the same day with 5mins in between. Compared to the CNN baseline, Fig. 6a shows that ADL-ID (in red), when

tested on LoRa data, provides a $20\% - 24\%$ improvement in the classification accuracy, pushing the performance from 57% to 76% on capture 2, from 41% to 65% on capture 3, and from 29% to 53% on capture 4. This improvement can be linked to the efforts of our network to minimize the presence of domain-specific features from the source and target domains in the RFF feature vector, which enhances the network's generalization. On the WiFi dataset (Fig. 6b), our proposed framework continues to show an improvement in

accuracy with highest gain of 9%, increasing the accuracy from 55% to 64% on capture 2, and from 40% to 42% on capture 3, and from 22% to 29% on capture 4. It can be noticed that the improvement percentage is higher in the LoRa dataset. We speculate that one reason could be that, in contrast to the LoRa dataset, the transmitters in the WiFi scenario were rebooted between the different captures. Consequently, a partial RF calibration has taken place, which affects the RFF distribution and stretches the distance between the source and target domains.

Second, we evaluate our network on long-term temporal adaptation. Specifically, we consider day 1 data as a source domain and day 2 data as the target domain. Hence, we train the network with labeled day 1 data and unlabeled day 2 data and then test it on data from the 4 days. While Fig. 7 shows that our network provides an improvement over the baseline CNN network in the long-term temporal adaptation as well, it must be marked that it did not succeed in bringing the performance on long-term varying domains closer enough to the source domain's performance. A reason for this is that our framework struggles to adapt to new domains when the difference between the source and target domains is too severe, which is the case in the long-term variation scenario. This results also conveys that performance of the framework is dependent on the similarity between the source and target domains. It is worth mentioning that similar behavior has been reported when domain adaptation techniques are applied in vision applications [20]. The variations in target domains in real-world RF applications are exceptionally more complicated compared to the vision domain. In the RFFP context, the variations in the target domain might include changes in the temporal dimension, channel characteristics, protocol configurations, and receiver hardware variations. Furthermore, these variations have made the interaction between the RFFs and domain-specific features more subtle, which results in a more difficult disentanglement process.

VI. Conclusion

We presented a large-scale WiFi 802.11b dataset, collected from 50 Pycom devices, and analyzed the performance of CNN-based RFFP on short- and long-term variations using LoRa and WiFi data. We also proposed ADL-ID, which provides up to 24% improvement in short-term temporal adaptations and up to 10% improvement in long-term variations over the baseline CNN. These results prove the potential of disentanglement representation techniques in enabling robust RFFP. However, more work is still needed to fully understand the RFF behaviors when the domain changes and to develop more robust RFFP methods.

VII. Acknowledgment

References

[1] S. Dang, O. Amin, B. Shihada, and M.-S. Alouini, "What should 6g be?" *Nature Electronics*, vol. 3, no. 1, pp. 20–29, 2020.

[2] A. Elmaghbub, B. Hamdaoui, and A. Natarajan, "Widescan: Exploiting out-of-band distortion for device classification using deep learning," in *GLOBECOM 2020-2020 IEEE Global Communications Conference*. IEEE, 2020, pp. 1–6.

[3] K. Merchant, S. Revay, G. Stantchev, and B. Nousain, "Deep learning for rf device fingerprinting in cognitive communication networks," *IEEE Journal of Selected Topics in Signal Processing*, vol. 12, no. 1, pp. 160–167, 2018.

[4] B. Hamdaoui, A. Elmaghbub, and S. Mejri, "Deep neural network feature designs for RF data-driven wireless device classification," *IEEE Network*, vol. 35, no. 3, pp. 191–197, 2020.

[5] A. Elmaghbub and B. Hamdaoui, "LoRa device fingerprinting in the wild: Disclosing RF data-driven fingerprint sensitivity to deployment variability," *IEEE Access*, 2021.

[6] B. Hamdaoui and A. Elmaghbub, "Deep-learning-based device fingerprinting for increased lora-iot security: Sensitivity to network deployment changes," *IEEE Network*, vol. 36, no. 3, pp. 204–210, 2022.

[7] F. Restuccia, S. D'Oro, A. Al-Shawabka, M. Belgiovine, L. Angioloni, S. Ioannidis, K. Chowdhury, and T. Melodia, "Deepradioid: Real-time channel-resilient optimization of deep learning-based radio fingerprinting algorithms," in *Proc. of the 20th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, 2019, pp. 51–60.

[8] S. Rajendran and Z. Sun, "Rf impairment model-based physical-layer identification for enhanced domain generalization," *IEEE Transactions on Information Forensics and Security*, vol. 17, pp. 1285–1299, 2022.

[9] A. Al-Shawabka, P. Pietraski, S. B. Pattar, F. Restuccia, and T. Melodia, "Deeplora: Fingerprinting lora devices at scale through deep learning and data augmentation," in *Proceedings of the Twenty-second International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing*, 2021, pp. 251–260.

[10] I. Redko, E. Morvant, A. Habrard, M. Sebban, and Y. Bennani, "A survey on domain adaptation theory: learning bounds and theoretical guarantees," *arXiv preprint arXiv:2004.11829*, 2020.

[11] J. Gaskin, B. Hamdaoui, and W.-K. Wong, "Tweak: Towards portable deep learning models for domain-agnostic LoRa device authentication," in *2022 IEEE conference on communications and network security (CNS)*. IEEE, 2022.

[12] T. Wang, L. Yu, W. Wang, and J. Wei, "Specific emitter identification based on the multi-discrepancy deep adaptation network," *IET Radar, Sonar & Navigation*, 2022.

[13] J. Wang, B. Zhang, J. Zhang, N. Yang, G. Wei, and D. Guo, "Specific emitter identification based on deep adversarial domain adaptation," in *2021 4th International Conference on Information Communication and Signal Processing (ICICSP)*. IEEE, 2021, pp. 104–109.

[14] I. Higgins, D. Amos, D. Pfau, S. Racaniere, L. Matthey, D. Rezende, and A. Lerchner, "Towards a definition of disentangled representations," *arXiv preprint arXiv:1812.02230*, 2018.

[15] K. Bousmalis, G. Trigeorgis, N. Silberman, D. Krishnan, and D. Erhan, "Domain separation networks," *Advances in neural information processing systems*, vol. 29, 2016.

[16] A. Elmaghbub and B. Hamdaoui, "Comprehensive rf dataset collection and release: A deep learning-based device fingerprinting use case," in *2021 IEEE Globecom Workshops (GC Wkshps)*. IEEE, 2021, pp. 1–7.

[17] A. Al-Shawabka, F. Restuccia, S. D'Oro, T. Jian, B. C. Rendon, N. Soltani, J. Dy, S. Ioannidis, K. Chowdhury, and T. Melodia, "Exposing the fingerprint: Dissecting the impact of the wireless channel on radio fingerprinting," in *IEEE INFOCOM 2020-IEEE Conference on Computer Communications*. IEEE, 2020, pp. 646–655.

[18] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky, "Domain-adversarial training of neural networks," *The journal of machine learning research*, vol. 17, no. 1, pp. 2096–2030, 2016.

[19] D. Eigen, C. Puhrsch, and R. Fergus, "Depth map prediction from a single image using a multi-scale deep network," *Advances in neural information processing systems*, vol. 27, 2014.

[20] M. Wulfmeier, A. Bewley, and I. Posner, "Addressing appearance change in outdoor robotics with adversarial domain adaptation," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 1551–1558.