

Physical-Layer Fingerprinting of LoRa devices using Supervised and Zero-Shot Learning

Pieter Robyns*
Hasselt University - tUL - imec
Martelarenlaan 42
Hasselt 3500, Belgium
pieter.robyns@uhasselt.be

Eduard Marin*
ESAT-COSIC and imec
Kasteelpark Arenberg 10, box 2452
Heverlee 3001, Belgium
eduard.marin@esat.kuleuven.be

Wim Lamotte
Hasselt University - tUL - imec
Martelarenlaan 42
Hasselt 3500, Belgium
wim.lamotte@uhasselt.be

Peter Quax
Hasselt University - tUL - imec
Martelarenlaan 42
Hasselt 3500, Belgium
peter.quax@uhasselt.be

Dave Singelée
ESAT-COSIC and imec
Kasteelpark Arenberg 10, box 2452
Heverlee 3001, Belgium
dave.singelee@esat.kuleuven.be

Bart Preneel
ESAT-COSIC and imec
Kasteelpark Arenberg 10, box 2452
Heverlee 3001, Belgium
bart.preneel@esat.kuleuven.be

ABSTRACT

Physical-layer fingerprinting investigates how features extracted from radio signals can be used to uniquely identify devices. This paper proposes and analyses a novel methodology to fingerprint LoRa devices, which is inspired by recent advances in supervised machine learning and zero-shot image classification. Contrary to previous works, our methodology does not rely on localized and low-dimensional features, such as those extracted from the signal transient or preamble, but uses the entire signal. We have performed our experiments using 22 LoRa devices with 3 different chipsets. Our results show that identical chipsets can be distinguished with 59% to 99% accuracy per symbol, whereas chipsets from different vendors can be fingerprinted with 99% to 100% accuracy per symbol. The fingerprinting can be performed using only inexpensive commercial off-the-shelf software defined radios, and a low sample rate of 1 Msps. Finally, we release all datasets and code pertaining to these experiments to the public domain.

CCS CONCEPTS

•Security and privacy → Mobile and wireless security; •Networks → Network privacy and anonymity;

KEYWORDS

LoRa; PHY layer; fingerprinting

ACM Reference format:

Pieter Robyns, Eduard Marin, Wim Lamotte, Peter Quax, Dave Singelée and Bart Preneel. 2017. Physical-Layer Fingerprinting of LoRa devices using Supervised and Zero-Shot Learning. In *Proceedings of WiSec '17, Boston, MA, USA, July 18-20, 2017*, 6 pages.
DOI: 10.1145/3098243.3098267

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WiSec '17, Boston, MA, USA

© 2017 Copyright held by the owner/author(s). Publication rights licensed to ACM.
978-1-4503-5084-6/17/07...\$15.00
DOI: 10.1145/3098243.3098267

1 INTRODUCTION

Physical (PHY)-layer device identification is a technique through which it is possible to uniquely identify devices by looking at small differences in their analog Radio Frequency (RF) signals. These differences are caused by imperfections of their analog hardware components, which allow to create a unique fingerprint of the device [2, 10, 22]. PHY-layer device identification has been proposed for various purposes such as access control and the detection of cloning and wormholes [21]. Several articles have demonstrated the feasibility of PHY-layer device fingerprinting for a variety of wireless technologies such as Radio Frequency Identification (RFID) or High Frequency (HF) (e.g. [5, 6]).

Over the past few years, the rise of Internet of Things (IoT) appliances has introduced numerous new PHY-layer protocols such as SigFox, Wi-Fi HaLow, LTE for IoT, Weightless, and LoRa. Among those, LoRa is currently one of the most promising wireless technologies, as it allows long-range, low-power and low-cost communication. This makes LoRa suitable for devices which need to send small amounts of data a few times per day over long distances. In terms of security and privacy, the current LoRaWAN Medium Access Control (MAC) specification provides built-in data confidentiality, integrity, and device authentication. However, it does not offer privacy. More specifically, all LoRaWAN messages contain a unique MAC address that can identify the sender device [26]. Nonetheless, one could envision a future, privacy-preserving version of the LoRaWAN protocol to support applications where a certain degree of privacy is needed. For example, the MAC address could be periodically randomized similarly to current Wi-Fi implementations. However, even if these changes are applied to the current LoRaWAN protocol, it remains unclear whether adversaries could still identify LoRa devices based on their analog RF signals alone. In this paper we tackle this problem and investigate whether an adversary can identify, locate or track any LoRa device regardless of the cryptographic mechanisms being used in the higher layers.

Contributions: Our first contribution is a fully functional, open-source software decoder of the LoRa modulation scheme. A second contribution is a novel fingerprinting methodology that applies *supervised machine learning* techniques to radio signals in order

* Both authors contributed equally to this work and share first authorship.

to distinguish between multiple known LoRa transmitters. Our methodology is inspired by recent advances in image and speech recognition, where state-of-the-art performance was achieved using raw data [16, 24, 25]. Our classifier achieves 59%–99% accuracy, even when devices with identical chipsets are up to 100 meters away from the fingerprinter. Additionally, our methodology can be applied to any part of the frame in contrast to previous works, where only part of the frame (e.g. the preamble [10, 19, 20]) is used for fingerprinting. Our technique is fully automated, passive, does not rely on underlying properties of the modulation scheme, and can be performed with sample rates as low as 1 Msps. Moreover, we show how our classifier can be extended to recognize *previously unseen* transmitters using zero-shot learning techniques. Finally, we performed a number of experiments that reveal the effects of (i) the distance between the fingerprinter and the LoRa device, (ii) time, and (iii) the sample rate on the classifier’s performance.

2 FINGERPRINTING LORA

PHY-layer fingerprinting leverages on small differences in the analog RF signals sent by wireless devices to uniquely identify them. These differences are caused by imperfections introduced in the analog hardware components during the manufacturing process [2].

2.1 Classification

In traditional approaches for classifying devices based on PHY-layer features, N acquisitions of several features are reduced in dimensionality and then averaged into one final, low-dimensional feature vector. This feature vector is subsequently learned by a classifier in order to construct a template for a specific device [4, 8, 10, 22]. Finally, the template is matched with the device by means of a similarity metric, such as Euclidian distance [22], entropic distance [10], or KL-divergence [28]. Although such approaches allow to distinguish between devices with high accuracy, they have several shortcomings. First, N is determined empirically, which makes these approaches hard to generalize for different channel conditions, hardware, or modulation schemes. Second, the selection of which features to use depends on the expertise of the researcher and can influence the result on multiple levels.

We propose a novel per-symbol classification methodology that aims to overcome some of these shortcomings by using high-dimensional features learned in an automated fashion. Our methodology is inspired by recent advances in Computer Vision (CV), more specifically in image and speech recognition. Here, state-of-the-art classification results are achieved by learning on raw data, such as the image pixels or time-domain waveform samples rather than manually selected features [12, 24, 27]. To limit the dimensionality and to ensure payload independence of the classification, we apply our methodology to the information contained in each i -th LoRa symbol $s(t)^{(i)} \dots s(t)^{(n)}$ separately for a frame consisting of n symbols.

2.1.1 Supervised classification. In our supervised classification approach, the fingerprinter is given a reference set of F -dimensional input features $x^{(1)} \dots x^{(n)} \in X^{n \times F}$ extracted from $s(t)^{(1)} \dots s(t)^{(n)}$, and corresponding C -dimensional class label tensor $y^{(1)} \dots y^{(n)} \in Y^{n \times C}$. Here, the term “class” refers to a single radio chip of a device, where the corresponding label can be “LoRa 1–22”. Given a model parameterized by the learned variables θ , the following

loss function $L(\theta)$ is minimized during an initial training phase: $L(\theta) = -\frac{1}{n} \sum_{i=1}^n \sum_{f=1}^F y_f^{(i)} \log(h_\theta(x_f^{(i)}))$. In this equation, the hypothesis function h_θ outputs the predicted class for each of the different models given the learning model parameters θ , and the loss function minimizes the cross entropy between the predicted and true classes.

After the training phase, the classifier extracts features and evaluates the model for each symbol in a LoRa frame in order to predict the most likely class. The class of the transmitter can be determined by performing majority voting on the symbols of the frame.

A requirement for accurate results under this classification approach is that a sufficiently large reference set of training samples for each of the device classes must be available. Furthermore, the accuracy depends on the quality of the training samples. A device should ideally be fingerprinted under different conditions. For example, by acquiring samples over a long period of time in order to prevent overfitting of the model. In Sect. 3.2, we will evaluate the effect of different channel conditions on the accuracy of our classifier.

2.1.2 Zero-shot classification. When fingerprinting a random observed radio signal, one typically would not possess a reference database of training samples for the associated (unknown) transmitters. In this case, the difficulty of the classification task is increased. Techniques that deal with the absence of training data for a set of unknown classes have been given several names over different domains: zero-shot classification [15, 25], semi-supervised anomaly detection [17], or open set recognition. We will use the term zero-shot classification henceforth in this work.

Despite not having training data for unseen classes, the fingerprinter can still learn discriminative *attributes* for a given set of known classes [14]. Such attributes can be interpreted as high-level, semantically meaningful properties that are used to describe a new class [13]. For fingerprinting LoRa devices, we were inspired by the zero-shot image classification approaches proposed by Socher et al. [25] and Lu [15].

Before applying zero-shot classification to PHY-layer fingerprinting, we first need to determine whether an observed symbol belongs to a known class or to a previously unknown class. To accomplish this goal, we have modeled the output values of the supervised classifier under a mixture of K multivariate Gaussian distributions, similarly to the work of Socher et al. [25]. Here, K is the number of known classes, which is also equal to the number of output neurons. The parameters μ_k and σ_k of each Gaussian $\mathcal{N}_k(\mu_k, \sigma_k)$ are determined by respectively taking the mean and standard deviation of the output values after feeding the input features of a known class k to the neural network. Then, we perform outlier detection by evaluating the indicator function $\mathbb{1}\{\sum_{k=1}^K h_\theta(x) \mathcal{N}_k(\mu_k, \sigma_k) < T\}$, where $h_\theta(x)$ is the output of the hypothesis function parameterized by θ given x , and T is an outlier tolerance threshold.

Next, the actual classification can be performed. If a symbol is not an outlier, it should belong to a known class. Therefore, it can be classified using the supervised classification approach from Section 2.1.1. Otherwise, we used the unsupervised DBSCAN [9] algorithm to cluster symbols transmitted by the same unknown class together. The ϵ parameter of the DBSCAN algorithm, which indicates the maximum distance between two points for them to be considered as in the same neighborhood, was set to the mean of the minimum Euclidean distance between all combinations of centroid

Table 1: Overview of all LoRa devices involved in the experiments and their chipset, identifiers, and quantity.

Device	Chipset	Identifiers	Quantity
Custom board	RN2483	LoRa 1-3	3
Pycom LoPy	SX1272	LoRa 4	1
Dragino LoRa/GPS HAT	RF96	LoRa 5	1
Adafruit Feather 32u4	RF96	LoRa 6	1
RN2483 breakout board	RN2483	LoRa 7-22	16

pairs in $\{\mu_k \dots \mu_K\}$. This ensures that symbols transmitted by different devices are appropriately mapped to different clusters, while symbols transmitted by the same device are mapped to the same cluster.

2.2 Learning models

To model the observed features for our supervised and zero-shot classifiers, several approaches could be considered. In this work, we examined Multilayer Perceptrons (MLPs) and Convolutional Neural Networks (CNNs), due to their success in similar classification tasks for other domains such as facial and speech recognition. Additionally, we briefly discuss Support Vector Machine (SVM)-based models due to their popularity in previous PHY-layer fingerprinting works.

Multilayer Perceptron: Our MLP model consists of one fully connected hidden layer with ReLU activation functions, and one fully connected output layer. Hence, the input features are mapped to the output device classes using the hypothesis function $h_\theta(x) = \sigma(\text{ReLU}(xW_1 + b_1)W_2 + b_2)$, where σ denotes the softmax function, and W and b respectively denote the weights and biases of the neurons. The softmax function scales each output from the classification of $x^{(i)}$ to form a discrete probability distribution for each $y \in Y$. Thus, the model learns the estimated probability that the symbol was transmitted by the device with label y .

Convolutional Neural Network: CNNs learn parameters to cross-correlation filter layers, which allows them to identify both low-level details at lower layers and high-level features at higher layers. Our CNN fingerprinting architecture consists of two hidden 1D convolution layers with kernel width 8 and ReLU activation functions, followed by a fully connected layer and softmax function for performing the classification.

C-Support Vector Classification: SVMs are trained to find an optimal hyperplane, in which the margin of separation between two classes is maximized [11]. This model is described in detail by Chang et al. [3]. In our experiments, we have used the SVM implementation of `sklearn` [18], which uses a one-vs-one scheme to perform multiclass classification.

3 IMPLEMENTATION AND RESULTS

Our laboratory setup comprises an Ettus Research B210 Universal Serial Radio Peripheral (USRP), antennas and a standard desktop computer. Table 1 gives an overview of all LoRa classes including their chipset, identifiers and quantity. We designed a custom board, which was always fixed in the same position, where we plugged in the LoRa transceivers before starting each of the fingerprinting experiments. This ensures that our results are not influenced by the distance between devices.

Table 2: Overview of the datasets used in this paper.

ID	# Symbols	Sampling rate	Date
I	495,216	1 Msps	January 27, 2017
II	124,740	1 Msps	January 30, 2017
III	497,595	2 Msps	January 17, 2017
IV	127,476	2 Msps	January 27, 2017
V	221,622	5 Msps	February 2, 2017
VI	55,908	5 Msps	February 3, 2017
VII	219,718	10 Msps	January 31, 2017
VIII	56,528	10 Msps	February 3, 2017

Each of the 22 LoRa devices used in the experiments was configured to continuously transmit frames with a 4-byte payload, using coding rate $4/8$ and SF 7 at 868.1 MHz. This configuration resulted in 36 symbols per frame. The payload bytes were randomized to ensure that the resulting symbol values are random as well, thus removing any bias due to the payload data. Both the training and test samples from the LoRa transmissions were acquired with the USRP tuned to a 868 MHz instead of 868.1 MHz carrier, in order to mitigate the effect of the USRP's Direct Current (DC) bias filter. To capture raw PHY frames, we have built a custom LoRa decoder using GNU Radio named `gr-lora`¹. The proprietary coding and whitening algorithms in the LoRa modulation scheme were reverse engineered to achieve this goal.

After extracting all synchronized symbols from the frame, the classifier needs to be trained on their features in order to distinguish between different LoRa transmitters. To reduce the number of possible symbol values from 2^{SF} to 1, we first calculate the ideal cyclic shift k of the symbol (i.e. its demodulated value) and modulate the symbol with value $\bar{k} = -k$. As a result of this operation, each modulated symbol $\hat{s}(t)$ is transformed back to the unmodulated chirp $s(t)$, and the errors introduced by the hardware are preserved.

An overview of all collected datasets is given in Table 2. We will refer to these datasets in future sections of this work using their respective Roman numeral label.

3.1 Classifier training

For implementing and training the models described in Sect. 2.2, we use the `tensorflow` machine learning library presented by Abadi et al. [1]². Before the training process, the entire dataset of collected features and labels is uniformly randomized. Then, a training set of 10,000 symbols and a cross validation set of 10,000 symbols are randomly fetched from the dataset for evaluation during training. A test set of 1,500 symbols is used for evaluation after training. The randomization process ensures that the training set is not biased by any particular device.

The LoRa symbols from each dataset are converted to a feature tensor. Considering the LoRa configuration parameters and receiver sample rate, the matrix of feature tensors for n symbols is $X \in \mathbb{R}^{n \times m}$, with $m = 2^{SF} \frac{f_s}{BW}$ where SF is the spreading factor, f_s is the sampling rate of the receiver, and BW is the bandwidth.

¹The decoder is available at <https://github.com/rpp0/gr-lora>.

²The code can be found at <https://github.com/rpp0/lora-phy-fingerprinting>.

Next, each used feature tensor $x^{(i)}$ was z-score normalized to prevent extreme gradient values from occurring during the training phase: $\hat{x}^{(i)} = \frac{x^{(i)} - \mu_{x^{(i)}}}{\sigma_{x^{(i)}}}$. This normalization also helps to reduce the effect of the absolute amplitude on classification, which is undesired since we do not distinguish devices based on their transmission power or physical location.

Finally, in each training step we feed a mini-batch of $b < n$ tensors to the classifier, and periodically log training set accuracy, test set accuracy, and cost function.

3.2 Fingerprinting experiments

Using the models defined in Sect. 2.2, we trained and evaluated our classifiers to distinguish between vendor models as well as individual LoRa devices. Furthermore, we investigated the effect of the sample rate, distance, and time on the subset accuracy, macro-averaged precision, and macro-averaged recall of the classifier.

3.2.1 Supervised classification experiment. In a first experiment, we trained our models with labeled instances to distinguish between different device vendors and different devices of the same type. Intuitively, the former should be easier since the analog hardware layout and design between various vendors may differ significantly. On the other hand, devices of the same vendor model only differ as a result of manufacturing variations.

Based on our findings, the crystal oscillator of the radio chip is especially susceptible to fingerprinting, since small differences in the oscillation frequency of the crystal will introduce a measurable Carrier Frequency Offset (CFO) error [28]. Contrary to previous works, where the CFO is explicitly measured as a scalar based on (averaged) samples of the signal [2, 7, 28], our approach uses the raw signal directly. As such, the CFO error manifests itself as a constant drift of the phase. However, since the phase difference is small for each of the many sample points, it is difficult for our classifier to learn this feature. We mitigate this issue by transforming the signal to the frequency domain using the Fast Fourier Transform (FFT). As a result of this transformation, the phase drift will shift the entire frequency spectrum, resulting in a large frequency difference for a few sample points of the FFT, which is easier for a classifier to learn.

For the test sets under identical channel conditions (I, III, V, VII), a disjoint and randomized subset from the same dataset was selected. The test sets (II, IV, VI, VIII) were sampled from data sets which were recorded on a different day from the training set. Each model was trained for 10,000 epochs, which depending on the used model corresponds to several hours of training on a Dell Precision T3610 with an Intel® Xeon® E5-1620 v2 CPU (3.70GHz). Table 3 summarizes the results of feeding these datasets to the MLP, CNN, and SVM classifiers when fingerprinting individual chipsets. When fingerprinting the 3 chipset vendors, the accuracy is 99% to 100% for all datasets and classifiers.

Figure 1 shows a t-SNE visualization of the output weights learned by the MLP model after training on dataset III. We can observe several clusters corresponding to the different LoRa devices. Due to space limitations, we will only discuss the results of the remaining experiments for dataset III and the MLP model. We believe these results are the most interesting, since 2 Msps is the maximum sample rate of low-cost Software Defined Radio (SDR) devices, such as

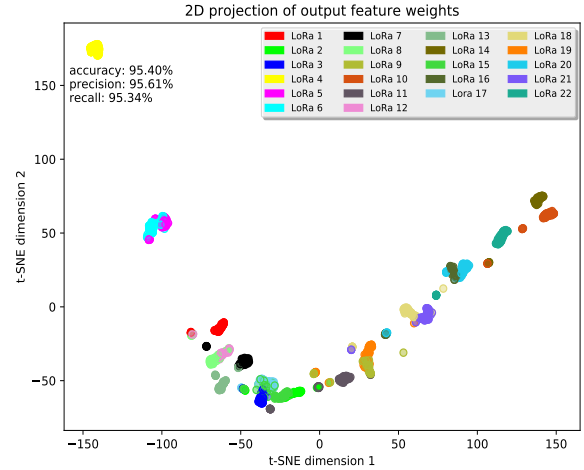


Figure 1: 2D t-SNE visualization of the output feature weights learned by the MLP model given dataset III. Each point represents a LoRa symbol, where the fill color indicates the true value and the outline color represents the predicted value.

the RTL-SDR, and the MLP model is faster to train while achieving similar or better accuracy compared to the CNN and SVM models.

3.2.2 Zero-shot classification experiment. A second experiment evaluates our zero-shot classification approach from Sect. 2.1.2 after training on 10,000 random symbols from dataset III. Here, the randomization and training procedures were identical to the previous experiment, except that we excluded symbols belonging to certain classes from the training set. Subsequently, we observed whether the classifier was able to cluster these unknown classes together. The results of this experiment are shown in Table 4.

We observed that the accuracy of the zero-shot classification largely depends on which devices are excluded from the training set. For example, in experiment ZS1, LoRa 4, 5 and 6 were excluded from the training set. Hence, the model was trained only on devices that have a RN2483 chipset. As a result, the classifier was not able to distinguish LoRa 5 and 6, i.e. both were grouped in the same cluster. This problem can be mitigated by including LoRa devices with similar fingerprints in the training set (see ZS6 in Table 4).

3.2.3 Effect of sample rate and time. Ideally, our fingerprinter should be able to classify devices at low sample rates and consistently over time. Therefore, we investigated the effect of these aspects on the classifier accuracy.

Ramsey et al. found that the fingerprinting accuracy increases with the sampling rate, but does not further improve above the Nyquist frequency [19]. On the contrary, in our experiments we observed that, under identical channel conditions, a sampling rate above the Nyquist frequency (250 KHz) increases the accuracy when devices have similar fingerprints. A higher sampling rate results in a higher granularity of frequency bins of the FFT spectrum, which allows the fingerprinter to detect more fine-grained frequency errors.

Table 3: Accuracy, precision, and recall when fingerprinting individual chipsets using supervised learning with test sets of 1,500 symbols.

Dataset	Model	Accuracy	Precision	Recall	Ident. chan. cond.
I	SVM	68.80%	63.08%	69.37%	Yes
I	CNN	89.40%	90.32%	89.23%	Yes
I	MLP	93.33%	93.32%	93.04%	Yes
II	SVM	53.80%	48.17%	52.09%	No
II	CNN	58.60%	55.94%	58.15%	No
II	MLP	58.67%	52.87%	58.19%	No
III	SVM	76.27%	76.01%	76.05%	Yes
III	CNN	94.27%	95.16%	94.88%	Yes
III	MLP	95.40%	95.61%	95.34%	Yes
IV	SVM	59.53%	62.66%	59.99%	No
IV	CNN	67.60%	73.64%	68.17%	No
IV	MLP	71.47%	75.04%	72.36%	No
V	SVM	83.00%	83.07%	82.77%	Yes
V	CNN	96.53%	97.03%	96.78%	Yes
V	MLP	99.00%	99.03%	98.98%	Yes
VI	SVM	69.33%	67.07%	70.23%	No
VI	CNN	76.80%	82.56%	76.72%	No
VI	MLP	75.07%	74.89%	75.37%	No
VII	SVM	81.27%	80.62%	81.13%	Yes
VII	CNN	98.00%	98.11%	98.01%	Yes
VII	MLP	98.67%	98.77%	98.63%	Yes
VIII	SVM	56.53%	53.02%	57.94%	No
VIII	CNN	60.33%	62.00%	62.22%	No
VIII	MLP	60.80%	58.75%	63.07%	No

Table 4: Accuracy, precision, and recall for the unknown “outlier” classes from the zero-shot classification experiments. The evaluation was performed on 1,500 symbols from dataset III.

Experiment	Excluded	Accuracy	Precision	Recall
ZS1	4,5,6	70.98%	75.36%	75.00%
ZS2	4	100.0%	100.0%	100.0%
ZS3	2,3,9,10,11	66.67%	41.45%	38.78%
ZS4	8,12,14,16,21	65.22%	48.55%	53.33%
ZS5	15,17,20	75.00%	63.44%	71.43%
ZS6	7,13,14	88.35%	67.82%	65.00%

Table 3 shows the classifier accuracy, precision, and recall when using the learning models with datasets of different sample rates. The MLP classifier for example achieves 93% per-symbol accuracy for 22 LoRa devices under identical channel conditions, with sample rates as low as 1 Msps. However, the accuracy drops for each of the sample rates over time. This was expected, since the crystal oscillator may undergo temperature changes over time, which causes its frequency to change and subsequently overlap with training data from a different radio chip. When examining the confusion matrix, we observed that the misclassifications indeed mostly occur with neighboring clusters (see Fig. 1). This issue could be mitigated by providing more training data gathered over an extended period of time or by periodically updating the model (i.e. “adaptive learning”)

to reflect changes in the channel conditions. Such adaptive learning models could be considered in future work.

3.2.4 Effect of distance. We evaluated how increasing the distance between the LoRa devices and the fingerprinter affects the accuracy of our classifier. For this purpose, we performed a series of experiments within a building in which the fingerprinter was always kept in the same location, whereas the LoRa devices were placed in three different locations. In the first experiment, the LoRa devices were in an adjacent room which is approximately 20 meters away from the fingerprinter (D1). In the second and third experiment, the LoRa devices were placed in a room that is 50 meters (D2) and 100 meters away (D3) from the fingerprinter, respectively.³ In future work, we plan to further increase the distance between the LoRa devices and the fingerprinter.⁴

For the signal test sets collected from D1, D2, and D3, our classifier respectively achieves an accuracy of 94.33% 98.40% and 96.40% after training on signals from the respective location. However, we found that the classifier achieves only 22.00% – 26.53% accuracy when a test set from one location is evaluated on a model that was previously trained on signals from a different location. From this observation we can conclude that the channel conditions significantly impact the accuracy of our classifier. In Sect. 4, we will briefly describe two possible ways to overcome this problem.

4 DISCUSSION AND IMPLICATIONS

Training with artificial noise: Our experiments reveal that the accuracy of our classifier degrades when the LoRa signals in the training and testing phases are captured under different channel conditions. Intuitively, one way to overcome this problem would be to use adaptive learning, which allows the classifier to continuously learn and dynamically adapt the models. However, this approach is susceptible to attacks where adversaries try to maliciously influence the way the classifier learns in their favor such that their signals are eventually considered as valid ones. Another possibility to mitigate this problem would be to add artificial noise to the training signals of the classifier, which could be used to simulate varying channel conditions in practice. Similar techniques are applied in the domain of image recognition to increase the robustness of the classification.

Resistance to attacks: We acknowledge that all (including ours) existing PHY-layer identification systems are susceptible to impersonation attacks. In an extended version of this paper, we will discuss the feasibility of performing attacks against our system.

Defensive fingerprinting: Despite achieving high accuracy when fingerprinting LoRa devices both from the same and different chipsets, PHY-layer fingerprinting should never be used for access control or authentication purposes alone. We emphasize that PHY-layer fingerprinting should only be implemented in combination with other security mechanisms, e.g. as a second factor authentication. Based on our results, we can also conclude that LoRa should not be used in applications where strong anonymity guarantees are needed, as fingerprinting will allow to de-anonymize the traffic.

³Note that for the second and third experiments we used another antenna for LoRa 4, since the received signal was too weak.

⁴When using SF=7 in an indoor environment, 100 meters was the maximum distance from which the fingerprinter was able to receive signals.

Choice of learning models: Besides the MLP, SVM, and CNN learning models discussed in this paper, other models could have been considered for fingerprinting. Similarly, different hyperparameters, e.g. number of hidden layers, dropout probability, number of neurons, etc. could have been selected. We considered the examination of optimal architectural choices for the models out of the scope of this paper. Nevertheless, we believe this would be an interesting and useful subject for future work, since this could allow to further increase the accuracy of PHY-layer fingerprinting systems.

5 RELATED WORK

Remley et al. analysed the feasibility of fingerprinting 802.11 devices by extracting time- and frequency-domain features from 6 devices that belong to 3 different vendors [23]. Brik et al. proposed PARADIS, a system to fingerprint 802.11 devices based on modulation-specific errors in the frame with an accuracy of 99% [2]. The main limitations of this system are that it uses sophisticated equipment with a high sampling rate for capturing the signals. Han et al. proposed a technique called Geneprint, which identifies Ultra High Frequency (UHF) RFID devices with an accuracy of 99.68%. Geneprint uses features extracted from the signal's preamble using a USRP and a sampling rate of 10 Msps. Ramsey et al. introduced a technique to fingerprint IEEE 802.15.4 devices based on a combination of features extracted from the signal's preamble. This includes the variance, skewness, kurtosis of the instantaneous phase, frequency and amplitude [19, 20]. In [19], they also demonstrated how the fingerprinting can be done with a USRP and PXIe-1085 with a relatively low sampling rate that varies from 5 to 20 Msps. The accuracy was reported as 100% in high Signal-to-Noise Ratio (SNR) conditions. However, their experiments involved at most 6 devices. The previous two works are the closest to ours in terms of the selection of the sampling rate. However, they extract features only from the preamble, which may facilitate impersonation attacks. In this paper we are the first to fingerprint LoRa devices, and show that fingerprinting is possible even when the LoRa devices are up to 100 meters away from the fingerprinter.

6 CONCLUSIONS

This paper demonstrates an automated supervised classification approach that can distinguish LoRa devices by analyzing their RF signals. Our classifier achieves 59%–99% accuracy when fingerprinting identical chipsets, and 99%–100% accuracy when fingerprinting chipset models. We extended the classifier with zero-shot learning methods to recognize previously unseen classes and achieve 65%–88% accuracy for those classes under similar channel conditions. Our results show that an adversary can identify a transmitter independently of the used modulation scheme or cryptographic mechanisms being used in the higher layers.

7 ACKNOWLEDGEMENTS

The authors would like to thank the anonymous reviewers for their helpful comments, and Enrique Argones, Bram Bonn , Rafael Galvez and Balazs Nemeth for their support. This work was supported in part by a Ph.D. grant of the Research Foundation Flanders (FWO), the Research Council KU Leuven C16/15/058, the Flemish Government through the imec Distributed Trust program, in particular the Netsec project, and through ICON project Diskman.

REFERENCES

- [1] Martin Abadi, Ashish Agarwal, and others. 2016. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467* (2016).
- [2] Vladimir Brik, Suman Banerjee, Marco Gruteser, and Sangho Oh. 2008. Wireless device identification with radiometric signatures. In *Proc. 14th Int. Conf. Mobile computing and networking*. ACM, 116–127.
- [3] Chih-Chung Chang and Chih-Jen Lin. 2011. LIBSVM: a library for support vector machines. *Trans. on Intelligent Systems and Technology (TIST)* 2, 3 (2011), 27.
- [4] Boris Danev and Srdjan Capkun. 2009. Transient-based identification of wireless sensor nodes. In *Proc. 2009 Int. Conf. on Information Processing in Sensor Networks*. IEEE Computer Society, 25–36.
- [5] Boris Danev, Srdjan Capkun, Ramya Jayaram Masti, and Thomas S. Benjamin. 2012. Towards practical identification of HF RFID devices. *Trans. on Information and System Security (TISSEC)* 15, 2, Article 7 (July 2012), 24 pages.
- [6] Boris Danev, Thomas S. Heydt-Benjamin, and Srdjan Capkun. 2009. Physical-layer Identification of RFID Devices. In *Proc. 18th Conf. on USENIX Security Symposium (SSYM'09)*. USENIX Association, Berkeley, CA, USA, 199–214.
- [7] Boris Danev, Heinrich Luecken, Srdjan Capkun, and Karim El Defrawy. 2010. Attacks on physical-layer identification. In *Proc. 3rd Conf. on Wireless network security*. ACM, 89–98.
- [8] Boris Danev, Davide Zanetti, and Srdjan Capkun. 2012. On physical-layer identification of wireless devices. *Computing Surveys (CSUR)* (2012).
- [9] Martin Ester, Hans-Peter Kriegel, J rg Sander, Xiaowei Xu, and others. 1996. A density-based algorithm for discovering clusters in large spatial databases with noise.. In *Kdd*, Vol. 96. 226–231.
- [10] Jinsong Han, Chen Qian, Panlong Yang, Dan Ma, Zhiping Jiang, Wei Xi, and Jizhong Zhao. 2016. Geneprint: generic and accurate physical-layer identification for UHF RFID tags. *IEEE/ACM Trans. on Networking* 24, 2 (2016), 846–858.
- [11] Marti A. Hearst, Susan T Dumais, Edgar Osuna, John Platt, and Bernhard Scholkopf. 1998. Support vector machines. *IEEE Intelligent Systems and their Applications* 13, 4 (1998), 18–28.
- [12] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*. 1097–1105.
- [13] Christoph H Lampert, Hannes Nickisch, and Stefan Harmeling. 2014. Attribute-based classification for zero-shot visual object categorization. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 36, 3 (2014), 453–465.
- [14] Hugo Larochelle, Dumitru Erhan, and Yoshua Bengio. 2008. Zero-data Learning of New Tasks. In *Proc. 23rd Nat. Conf. on Artificial Intelligence*. 646–651.
- [15] Yao Lu. 2015. Unsupervised Learning on Neural Network Outputs. *CoRR* abs/1506.00990 (2015).
- [16] Aaron Van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. 2016. Pixel Recurrent Neural Networks. In *Proc. 33rd Int. Conf. on Machine Learning*.
- [17] C David Page and Sriraam Natarajan. 2014. Encyclopedia of Machine Learning and Data Mining. 2 (2014), 1–24. DOI: <https://doi.org/10.1007/978-1-4899-7502-7>
- [18] F. Pedregosa, G. Varoquaux, and others. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [19] Benjamin W. Ramsey, Stubbs, and others. 2015. Wireless infrastructure protection using low-cost radio frequency fingerprinting receivers. *Int. Journal of Critical Infrastructure Protection* 8 (2015), 27–39.
- [20] Benjamin W Ramsey, Michael A Temple, and Barry E Mullins. 2012. PHY foundation for multi-factor ZigBee node authentication. In *Global Communications Conference (GLOBECOM)*. IEEE, 795–800.
- [21] K. Bonne Rasmussen and S. Capkun. 2007. Implications of radio fingerprinting on the security of sensor networks. In *Proc. 3th. Int. Conf. Security and Privacy in Communications Networks and the Workshops - SecureComm*. 331–340.
- [22] Saeed Ur Rehman, Kevin Sowerby, and Colin Coghill. 2012. RF fingerprint extraction from the energy envelope of an instantaneous transient signal. In *Communications Theory Workshop (AusCTW), 2012 Australian*. IEEE, 90–95.
- [23] KA Remley, CA Grosvenor, RT Johnk, DR Novotny, PD Hale, MD McKinley, A Karygiannis, and E Antonakakis. 2005. Electromagnetic signatures of WLAN cards and network security. In *Proc. 5th IEEE Int. Symposium on Signal Processing and Information Technology*. IEEE, 484–488.
- [24] Karen Simonyan, Sander Dieleman, Andrew Senior, and Alex Graves. 2016. WaveNet: A Generative Model for Raw Audio. (2016), 1–15.
- [25] Richard Socher, Milind Ganjoo, Christopher D Manning, and Andrew Ng. 2013. Zero-shot learning through cross-modal transfer. In *Advances in Neural Information Processing Systems*. 935–943.
- [26] N Sornin, M Luis, T Eirich, T Kramp, and O Hersent. 2015. LoRaWAN™ Specifications. *LoRa™ Alliance* (2015).
- [27] Christian Szegedy, Wei Liu, and others. 2015. Going deeper with convolutions. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*. 1–9.
- [28] Tien Dang Vo-Huu, Triet Dang Vo-Huu, and Guevara Noubir. 2016. Fingerprinting Wi-Fi Devices Using Software Defined Radios. In *Proc. 9th Conf. Security & Privacy in Wireless and Mobile Networks*. ACM, 3–14.