Crowdsourced wireless spectrum anomaly detection

Sreeraj Rajendran*, Vincent Lenders[†], Wannes Meert[‡], and Sofie Pollin*

Email: {sreeraj.rajendran, sofie.pollin}@esat.kuleuven.be

wannes.meert@cs.kuleuven.be, vincent.lenders@armasuisse.ch.

*Department ESAT, KU Leuven, Belgium, [‡]Department of Computer Science, KU Leuven, Belgium

[†]armasuisse, Thun, Switzerland

Abstract—Automated wireless spectrum monitoring across frequency, time and space will be essential for many future applications. Manual and fine-grained spectrum analysis is becoming impossible because of the large number of measurement locations and complexity of the spectrum use landscape. Detecting unexpected behaviors in the wireless spectrum from the collected data is a crucial part of this automated monitoring, and the control of detected anomalies is a key functionality to enable interaction between the automated system and the end user. In this paper we look into the wireless spectrum anomaly detection problem for crowdsourced sensors. We first analyze in detail the nature of these anomalies and design effective algorithms to bring the higher dimensional input data to a common feature space across sensors. Anomalies can then be detected as outliers in this feature space. In addition, we investigate the importance of user feedback in the anomaly detection process to improve the performance of unsupervised anomaly detection. Furthermore, schemes for generalizing user feedback across sensors are also developed to close the anomaly detection loop.

Index Terms—Deep learning, Anomaly detection, Crowdsourcing

I. INTRODUCTION

Radio spectrum is one of our most precious and widely used natural resources. With the advent of new wireless communication technologies, spectrum usage has become very complex resulting in airwave congestion and other interference issues [1,2]. Diverse spectrum regulations across countries have also contributed to this chaotic spectrum usage when non-standardized wireless devices cross country borders. In addition, easily available illegal wireless jammers or low-cost Software Defined Radio (SDR) devices which are capable of generating custom wireless signals are making the problem worse. Unintentional and intentional jamming of localization services such as Global Navigation Satellite System (GNSS), which is used for a wide range of applications such as automated vehicle navigation, airplane landing procedures and maritime vessel tracking, is increasing with the availability of easy jamming devices [3]–[5]. Furthermore, illegal repeaters, which are used to boost mobile coverage, can adversely affect the cell planning of mobile operators resulting in poor coverage and dropouts [6]. On the other hand, densified small cells are becoming fundamental for the new high throughput and low latency networks [7]. Automated analysis and detection of anomalous behaviors in the spectrum is becoming crucial to promote the growth of new generation wireless systems.

Wireless spectrum anomalies can occur due to reasons ranging from unintentional radio frequency (RF) leakages

to intentional spectrum misuse or jamming. For common anomaly detection use cases, a generic anomaly detector that can detect a wide range of anomalies such as *new transmitters including jammers, transmitter power variations, out of band transmissions, unexpected transmission patterns and transmission duty cycle variations* will suffice. In this paper we introduce such a generic anomaly detector incorporating user preferences to the detection process.

The wide range of wireless spectrum anomalies makes it infeasible to collect and label all types of anomalies and formulate anomaly detection as a supervised learning problem. A majority of the algorithms in literature [8]–[13] formulate wireless spectrum anomaly detection as an unsupervised learning problem where the models try to learn the normal spectrum data distributions and detect the uncommon patterns as anomalous. These models make the basic assumption that non-frequent behaviour is anomalous which is not always true. For example, transmissions in ISM bands exist that are very sparse but are not anomalies. Similarly, pirate transmitters or transmission duty cycle limit violations in a unlicensed band are anomalous even when they are very frequent. In this paper we solve wireless spectrum anomaly detection by formulating it as a crowdsourced active learning problem. More specifically we provide solutions to the following questions addressing various challenges posed by them:

- 1) How to optimize anomaly detection based on user feedback from different users in a crowdsourced network?
- 2) How to improve the anomaly detection of a particular user by using expert feedback from other users which may not have fully overlapping interests?
- 3) How to generalize user feedback across different sensors?

Some of the major challenges that need to be addressed to answer the aforementioned questions are given below.

A. Challenge 1: Finding an interpretable feature space

Contrary to other sensing contexts such as temperature and air pollution monitoring, the amount of data generated by wireless spectrum sensors is quite huge. In addition to the huge amount of data, most of the interesting techniques such as classification, prediction and anomaly detection suffer from the curse of dimensionality when addressed in the higher dimensional input data space. Furthermore, the data costs related to the transfer and storage of the spectrum sensing sensor data is huge. For instance, the crowdsourced spectrum

Fig. 1: Anomaly detection and clustering framework: The framework consists of an unsupervised data feature module and anomaly feature module for optimized feature extraction from the data collected from multiple sensors S_1 to S_N . These features are further used for active anomaly clustering and detection with user interaction.

sensing framework, Electrosense [14], enables three pipelines with very low, medium and high data transfer costs namely: Feature, Power Spectral Density (PSD) and In-phase and quadrature phase (IQ) pipeline respectively. For fulfilling the long term spectrum monitoring dream, it is essential to extract critical features and employ the feature pipeline whenever possible. First, these features should also be compressed and interpretable for further analyzing the cause for the anomaly. Second, the compressed selected feature space should promote outlier detection. Finally, the selected feature space should be the same across sensors for unifying the anomaly detection process for different users.

Our contribution: We introduce novel data driven feature extraction modules for both normal and anomalous wireless spectrum input data. The proposed training process helps to generate a common feature space across sensors without compromising the model's anomaly detection capabilities of individual sensors.

B. Challenge 2: Detecting meaningful anomalies

Anomalies can be outliers in the normal feature space or process variations. Anomalies can also be some frequent events hidden inside the feature space, such as transmitter duty cycle variations, which prevents us from putting a simple outlier threshold in the feature space. It is also challenging to separate anomalies in an unsupervised feature space across a large number of crowdsourced sensors. The developed framework should also detect a change in the process generating the input spectrum data and retrain the unsupervised model in case of large input distribution variations.

Our contribution: An unsupervised anomaly detection phase followed by a semi-supervised phase is introduced to detect anomalies from frequent as well as less frequent features. Additionally, the proposed framework incorporates provisions for automated model retraining and threshold adjustments of the unsupervised detection phase enabling improved anomaly detection across sensors.

C. Challenge 3: Anomaly knowledge sharing across sensors

Large variations of sensed wireless spectrum from sensors at different physical locations and regions makes anomaly knowledge sharing across sensors challenging. Knowledge about common anomaly types should be shared across all sensors and frequency bands for fast anomaly detection and generalization of the cause of anomalies in different locations.

Our contribution: We introduce an active anomaly clustering process, as shown in Figure 1, to group similar anomalies across sensors from different frequency bands with varying signal-to-noise ratio (SNR)s. The specialized active clustering helps to learn relevant local level distances to separate even closely located clusters if they show different behaviour. We show that knowledge acquired through pairwise queries on a single sensor can be effectively used to group similar anomalies from other sensors.

D. Challenge 4: Incorporating user feedback

In a crowdsourced network anomalies vary based on user or application perspective. One user's signal might be another's anomaly. So it is important to curate and present anomalies to the users based on their preferences across multiple sensors. In addition, the user anomaly feedback can aid the global anomaly detection process by semi-supervised learning. Thus the main challenge is to devise a scheme that can make use of expert labeling from minimal interaction with the users. The labeling process should neither be labour intensive for the users nor prone to much accuracy variations with some level of expected incorrect labeling.

Our contribution: An active anomaly detection after interactive clustering is proposed, which makes use of label propagation on the clusters, for anomaly detection. Anomaly detection over clusters helps the framework to learn anomalies from a few labeled instances (as anomalous or not) in each cluster and even easily store anomalies based on user preferences. Furthermore, the proposed anomaly detection model parameters can be adjusted to combat incorrect answering both in the active clustering and anomaly detection phases.

The rest of the paper is organized as follows. Relevant anomaly detectors proposed in the literature and presented in Section II. Section III explains the anomaly detection framework along with the training and detection algorithms. The dataset used for analyzing the performance of the proposed framework is presented in Section IV. The performance results and the advantages of the proposed model are discussed in Section V. The paper is concluded in Section VI.

II. RELATED WORK

The anomaly detectors proposed in literature can be mainly brought down to two major classes: unsupervised and semi-supervised anomaly detection.

A. Unsupervised anomaly detection

Various general anomaly detectors are proposed in literature such as One class Support Vector Machine (OSVM), Isolation Forest (IFO) [15] and Local Outlier Factor (LOF) [16] which makes the basic assumption that outliers represent potential anomalies. These models normally perform poor with increasing input data dimensionality. An anomaly detector for Dynamic Spectrum Access (DSA) is presented in [8], where distributed power measurements via cooperative sensing are used for anomaly detection. The proposed detector is limited to authorized user anomaly detection only, for the specific case of DSA. Similarly [9] makes use of Hidden Markov Models (HMM) on spectral amplitude probabilities that can detect interference on the channel of interest again in the DSA domain. Recently in [10], the authors presented a Recurrent Neural Network (RNN) anomaly detector based on predictive modeling of raw IQ data. The authors used a Long Short Term Memory (LSTM) model for predicting the next 4 IQ samples from the past 32 samples and an anomaly is detected based on the prediction error. Even though this model works on raw physical layer data which requires no expert feature extraction, it is still not sufficiently automated and generic for practical anomaly detection.

In [11], the authors extend this prediction idea on spectrograms and test the model on some synthetic anomalies. A reconstruction based anomaly detector based on vanilla deep autoencoders is presented in [12]. In Spectrum Anomaly Detector with Interpretable FEatures (SAIFE) [13] we proposed an Adversarial Autoencoder (AAE) [17] based anomaly detector which fills the shortcomings of all the aforementioned methods and can be trained in a semi-supervised fashion for extracting interpretable features, localize anomalies in the spectrum along with data compression and signal classification abilities. An adapted version of SAIFE is used for the initial unsupervised anomaly detection part of the proposed framework in this paper.

B. Semi-supervised anomaly detection

A few semi-supervised anomaly detection algorithms exist in literature such as Lightweight on-line detector of anomalies (LODA) [18,19]. In [18], the authors make use of an unsupervised ensemble-based algorithm which is further enhanced

Section	Layer	Type	Kernel	No. of filters
Encoder	1,2,3	Conv+maxpool	(3,3)	32, 32, 32
	4_1	Fully connected (softmax)	8 (nlabels)	-
	4_2	Fully connected (linear)	50 (nfeatures)	-
Decoder	1,2,3	Conv+upsample	(3,3)	32, 64, 64
	4	Conv	(3,3)	1
	5	Fully connected (linear)	6*300	-

TABLE I: Data and anomaly feature module layers.

with an active learning phase in [19]. A semi-supervised accuracy-at-the-top loss function is used to update weights of LODA based on user feedback. The base anomaly detection algorithm is not limited to LODA and the authors also show that the active anomaly detection phase can be incorporated with tree based algorithms too [20]. On the other hand in [21], the authors formulate anomaly detection as an optimization problem and propose a generalization using both labeled and unlabeled instances. Recently, in [22] the authors proposed Semi-Supervised Detection of Outliers (SSDO), a semi-supervised anomaly scoring algorithm which relies on clustering and user labeling.

III. ANOMALY DETECTION FRAMEWORK

We address all the challenges mentioned in the introduction by developing an unsupervised anomaly detection module followed by active anomaly clustering and detection modules as shown in the framework Figure 1. The framework consists of four major components: a data feature module, an anomaly feature module, an active clustering module and an active anomaly detector module which are explained in detail in the following subsections.

A. Module-1: Data feature generator and unsupervised anomaly detection

The data feature module extracts optimized compressed features from the sensed spectrogram data generated by each sensor. We adapt SAIFE [13], an AAE based unsupervised anomaly detector, with a different training procedure to address *Challenge 1* mentioned in the introduction. The layers used for the encoder and decoder of both data and anomaly feature modules are summarized in Table I. It has been already shown that SAIFE achieves very good anomaly detection accuracy when the anomaly is defined as a less frequent event. In addition, SAIFE can be trained in a semi-supervised fashion to extract interpretable results such as the signal position and bandwidth. The model can also achieve anomaly localization in the frequency space.

For this specific crowdsourced problem and further to help addressing *Challenge 3*, the SAIFE encoder and decoder modules are adapted to address the data variability across the sensors. We make use of a global data feature encoder which helps to obtain a *uniform feature space* across sensors for further clustering. In addition, a global encoder also allows to *share knowledge* across sensors and *generalize* the feature space better, as a single sensor might not witness all wireless signals. Along with a common generalized data feature space, the model should also achieve good initial anomaly detection capabilities. To achieve this, we make use of a sensor specific decoder module which prevents the model from averaging

out sensor specific anomalies. For example, if we consider FM radio stations at various locations their centre frequencies might be different. If the decoder is trained globally, the model will perceive that an FM transmission is present at all centre frequencies preventing it from detecting an anomalous pirate transmitter at a particular location.

An initial anomaly detection is performed after generating the data features using the standard anomaly detection score [13] based on the reconstruction loss (R_l), discriminator losses (D_{lcont} and D_{lcat}) and classification error as summarized below:

$$A_{score} = (R_l > (\mu_{R_{lt}} + n * \sigma_{R_{lt}})) \\ \vee ((\mu_{D_{lcont}} - n * \sigma_{D_{lcont}}) > D_{lcont} > (\mu_{D_{lcont}} + n * \sigma_{D_{lcont}})) \\ \vee ((\mu_{D_{lcat}} - n * \sigma_{D_{lcat}}) > D_{lcat} > (\mu_{D_{lcat}} + n * \sigma_{D_{lcat}})) \\ \vee (Class_{Encoder} \neq Class_{input}) \quad (1)$$

A simple n-sigma threshold is employed on the reconstruction and discriminator loss based on the mean (μ) and standard deviation (σ) values from the training data. An input data frame is classified as anomalous if A_{score} is *True*. The data feature module filters out a large number of normal input spectrum regions enabling the anomaly feature module to learn much more anomaly specific features as shown in Figure 1.

B. Module-2: Anomaly feature generator

This module captures the anomalous spectrum features that the data feature module is trained not to capture (as the data feature module extracts only frequent data features). The difference between the input and the reconstructed signal ($X - \hat{X}$) of all the anomalies detected by the data feature module is fed to the anomaly feature extractor module. We make use of the difference signal as these would represent only anomalies without normal data. This module is global to all the sensors which creates a common anomaly feature space across sensors for future clustering. An AAE similar to the one used in the data feature module is also employed in this module. The whole training process of the data and anomaly feature modules are summarized in Algorithm 1. Both the data and anomaly feature use the same encoder and decoder layers as summarized in Table I. *Tanh* activation is used for all convolutional layers and maxpooling and upsampling uses a factor of 2.

C. Module-3: Active clustering module

The extracted anomaly and data features are then fed to an interactive clustering module for further processing. The aim of the clustering module is to group similar anomalies and normal behaviour which are present in different frequency bands across different sensors with minimal user interaction. We make use of the data features along with anomaly features (F_a, F_d), to enable users to cluster even anomalies from the normal feature space. Constraint-based Repeated Aggregation and Splitting (COBRAS) [23], a query and time efficient interactive clustering algorithm is used for achieving clustering.

COBRAS makes use of the concept of super-instances: sets of instances that are assumed to belong to the same cluster in

Algorithm 1 Data and anomaly features module training

Input: $\mathbf{X}$: a set of spectrograms from different sensors

Output: a model that can generate data and anomaly features, (F_a, F_d)

- 1: Train data feature encoder and decoder to reduce reconstruction loss of $\mathbf{X}$ along with other losses
- 2: Global encoder: Freeze encoder weights
- 3: Sensor i specific decoder D_i : Train only the decoder part of the model specific to each sensor i using the data X_i from corresponding sensors.
- 4: Freeze weights of the data feature module
- 5: Determine detection threshold n , for the n -sigma cutoff of SAIFE
- 6: Initial unsupervised anomaly detection with SAIFE
- 7: Train anomaly feature module on the error signal $X_i - \hat{X}_i$ from all sensors i
- 8: Freeze weights of the anomaly feature module

Fig. 2: COBRAS clustering system relying on interaction in the form of pairwise queries.

the unknown target clustering. COBRAS starts with a super-instance that contains all instances and gradually refines the clustering based on pairwise queries asking whether the shown anomalies belong to the same cluster or not. In each iteration COBRAS performs two steps. First, it picks the largest super-instance from the cluster and splits it into new super-instances. A constraint based procedure is used for splitting the largest super-instance. The algorithm splits the largest super-instance into two new temporary super-instances based on 2-means and then queries the user, the relation between their medoids [24]. Medoids are similar in concept to means, but are restricted to the data points in the dataset or cluster whose average dissimilarity with all the cluster points are minimal. In our case, the error input images ($X_i - \hat{X}_i$) for the corresponding medoid features are shown to the user, asking whether they belong to the same cluster or not, as shown in Figure 2. In this paper the term *query* when mentioned related to clustering refers to the process of asking the user whether a pair belongs to the same cluster or not. If they can be linked, COBRAS considers the super-instance as pure, otherwise the clustering process is repeated on other super-instances. In the second merging step the relation between newly created clusters are determined and clusters are merged if they belong to the same group based on constraints. If the relation between two

Fig. 3: SSDO anomaly detection example: After achieving initial clustering using COBRAS, the algorithm performs unsupervised anomaly scoring and anomaly label propagation. In this example labels are provided saying clusters A and C are normal and Clus-B is anomalous (anomalies in red).

clusters is not known, they are queried again for user response. Detailed algorithm steps can be found in the original paper.

D. Module-4: Active anomaly detector

The active anomaly detector is an integral module of the entire framework which solves the *Challenge 3* and *Challenge 4* mentioned in the introduction. This module helps to label the clusters generated by COBRAS as anomalous or not. As mentioned in the introduction less frequent events detected by SAIFE could be non anomalous. In addition, the active anomaly detection also helps to store user specific anomaly preferences when anomaly labeling is done by a single user, still giving the flexibility to share these anomaly label knowledge across users. For instance, in a crowdsourced network the clusters which are labeled as anomaly by a majority of users can be considered as a global anomaly.

An anomaly score is generated on the clustered instances using SSDO [22], which can work in a semi-supervised setting. SSDO performs semi-supervised anomaly detection in two major steps after the clustering step as shown in Figure 3. First, it does an unsupervised anomaly scoring of instances based on three measures: (i) the point deviation of an instance from its cluster centre, (ii) the cluster deviation of a selected cluster from other clusters and (iii) the cluster size. In addition, SSDO supports score updating based on user anomaly labels with two controllable parameters α and k . α controls the impact of the user provided labels versus the labels from the clustering process and k determines the depth of label propagation for a single labeled instance. These parameters are very useful in controlling the speed of the entire process

Algorithm 2 Active Anomaly detection

Input: (F_a, F_d)

Output: Anomalous spectrograms

- 1: Active clustering on (F_a, F_d) using COBRAS
 - 2: Active anomaly score generation using SSDO
 - 3: Present final anomalies to users
-

Algorithm 3 Model retraining (Section IV-B)

- 1: R_{la} = Reconstruction loss of the $X_i - \hat{X}_i$ from the anomaly feature module
 - 2: **if** $R_{la} > \text{threshold}(2\sigma)$ **then**
 - 3: Add input to the anomaly training set
 - 4: **while** True **do**
 - 5: Retrain anomaly feature module
 - 6: Adjust data feature module detection threshold based on new labeled anomalies not detected by the data feature SAIFE and other anomaly clusters
 - 7: Retrain data and anomaly feature module in case of low anomaly scores for majority of clusters
 - 8: Sleep(threshold_days)
-

and helps in dealing with noisy labels from users. In this paper the term *query* when mentioned related to anomaly detection refers to the process of asking the user whether an instance is anomalous or not. The whole anomaly detection process is summarized in Algorithm 2. The reader is encouraged to read the original paper [22] for detailed equations of these measures.

E. Users and expert feedback

As shown in the framework Figure 1, users also play an important role in the clustering and anomaly detection process. Electrosense, already leverages crowdsourced sensors for large scale spectrum monitoring. In the newly proposed framework, expert user feedback is utilized both in the clustering and anomaly detection phase. By expert feedback, we refer to feedback from the wireless domain experts: either who develop and train the model or reputed experts among users. As the clustering process helps to share knowledge across sensors, we propose to involve experts in the clustering process. We also analyze the effect of incorrect labeling or feedback in detail in Section V. The active anomaly detection after clustering could be made open to all users as even with labeling errors we can store them as user specific anomaly preferences and use it globally across sensors based on majority voting. We would like to eventually enhance the spectrum monitoring experience for users with more user engagement in the anomaly detection process, by making this framework a part of Electrosense.

IV. DATASETS

Spectrogram data from the PSD pipeline of four electrosense sensors over a period of 56 days (starting from 8 Nov, 2018) are collected for the analysis of the proposed framework. The data is collected through the Electrosense API¹ with a spectral resolution of 10 kHz and time resolution of 60

¹<https://electrosense.org/open-api-spec.html>

seconds. The selected sensors are located in Spain, Denmark, Switzerland and Slovenia and data from the same frequency bands are collected from all these sensors as given in Table II. These sensors are low-cost RTLSDRs configured at a sampling rate of 2.4 MS/s with omni-directional antennas which are deployed indoors. The sensors follow sequential scanning of the spectrum with an Fast Fourier Transform (FFT) size set to 256 giving a frequency resolution close to 10 kHz. With a FFT size of 256 and sensor ADC bit-width of 8, we get an effective bitwidth of 12 resulting in a theoretical dynamic range of 74dB. Practical dynamic range depends on the ADC frontend stages and the noise level, which may vary between 60 to 65dB. Five FFT vectors are averaged for reducing the thermal noise of the receiver.

Three synthetic signals (i) *scont*: single continuous signal with random bandwidth, (ii) *randpulses*: random pulsed transmissions on the given band and (iii) *wpulse*: pulsed wideband signals covering the entire frequency with SNR from -20 to +20 dB are used as anomalies. These synthetic anomalies are injected to the original collected data which help us to have ground truth clusters (three clusters) and labels of the anomalies for further analysis.

A. Feature extraction model training

The dataset mentioned in the previous section is split into two subsets, a training and a testing subset, with equal number of vectors. A seed is used to generate random mutually exclusive array indices, which are then used to split the data into two ascertaining the training and testing sets are entirely different. The train and test datasets for performance analysis on the Electrosense dataset are selected from different, non-adjacent, time periods (7 days for training and the other 49 days for testing) to make sure that there are no identical points in the training and test dataset. Spectrogram inputs each with a vector size 6x300 (time-frequency bins) are used as input to the model. The model is trained in an unsupervised fashion to reduce the mean squared error between the input and decoder output and a semi-supervised fashion to learn the the continuous features and class labels. The adversarial networks as well as the autoencoder are trained in three phases: the reconstruction, regularization and semi-supervised phase as mentioned in [17]. The Adam optimizer [25], a first-order gradient based optimizer, with a learning rate of 0.001 is used for training in all the phases. In the semi-supervised phase the model is trained to learn the class, position and bandwidth of the input signal by training it on 20% of the labeled samples from the training set. The model is trained for 500 epochs which takes around one hour of training time on a x86 PC with Nvidia GeForce GTX 980 Ti graphics card. Once trained, the profiled model inference time for 500 input spectrograms is 0.0456 seconds on the same hardware resulting in a processing time of around 0.91 microseconds per input vector. The data feature modules (Module-1) along with the unsupervised anomaly detection are deployed on the sensors to reduce the involved data transfer costs by only sending features to the centralized processor. All the other anomaly detection modules mentioned in the framework are executed in the central processing backend.

Count	Frequencies (MHz)	Expected technology/modulations
1	86-108	FM
2	192-197	Broadcasting (DAB/Microphones)
3	198-228	Broadcasting (DAB/Microphones)
4	790-801	Tactical links/TV Broadcasting
5	801-810	Tactical links/TV Broadcasting
6	811-821	Tactical links/TV Broadcasting
7	933-935	Mobile (GSM-900)
8	955-960	Mobile (GSM-900)

TABLE II: Electrosense dataset frequency bands.

B. Model retraining

The anomaly and data feature module should be incrementally trained in the following scenarios. An increase in the anomaly detector reconstruction loss (R_{la}) indicates that new unseen anomalies are getting fed to the anomaly feature module. This demands retraining of the anomaly feature module. Similarly, a large number of non-anomalous clusters after expert labeling process might either indicate there are large process variations in the sensor output or the selected low frequency events are not anomalous to the users. This demands retraining of the feature extractor modules along with the initial unsupervised threshold update. The framework can also enable users to label portions of the spectrum that are anomalous but not detected by the initial unsupervised detection phase, resulting in a possible threshold reduction for incorporating more normal behaviours to the anomaly space. Summary of the training phase can be found in Algorithms 1, 2 and 3.

V. RESULTS AND DISCUSSION

The framework is evaluated based on the unsupervised anomaly detection performance, clustering performance and finally based on the anomaly detection accuracy. We answer the following questions based on the analysis:

Q1: How good is the initial unsupervised anomaly detection performance compared to state-of-the-art (SoA) models with the new interpretable common feature space and training process (Challenge 1)?

Q2: Does the common feature space along with active clustering help to share knowledge of anomalies across sensors (Challenge 3)?

Q3: Can the model effectively detect meaningful anomalies even in the feature space with active learning (Challenge 2)? Can the model achieve this with a few labels and work reasonably well even with incorrect feedback (Challenge 4)?

A. Q1: Feature extraction module performance

The trained data feature module's performance is compared against various SoA algorithms such as OSVM, IFO [15] and LODA [18]. The average anomaly detection accuracy and false positive rates of these algorithms over 8 frequency bands, 3 anomalies and four sensors for a cutoff $n = 2$ are plotted in Figure 4. It can be noticed that SAIFE performs better than all algorithms with good detection accuracy and low false positive rates. The false positive rate of other algorithms can be reduced by increasing the n value of the n -sigma cutoff threshold as shown in Figure 5 at the cost of detection accuracy. These

Fig. 4: Averaged true anomaly detection accuracies (left) and false positive rates (right) for different algorithms on 8 frequency bands from four Electrosense sensors with n -sigma cutoff $n = 2$.

Fig. 5: Averaged true anomaly detection accuracies (left) and false positive rates (right) for different algorithms on 8 frequency bands from four Electrosense sensors. n -sigma cutoff $n = 4$ for IFO and LODA, $n = 6$ for OSVM and SAIFE with $n = 2$ for comparison.

results show that unsupervised threshold selection, even with full anomaly ground truth knowledge, is difficult highlighting the need for a further semi-supervised phase. To systematically compare different algorithms, the detection accuracy for a constant false alarm rate of 5% and 10% are plotted in Figure 6. It can be noticed that SAIFE performs superior to all algorithms. Even with the new training process the performance of the model is consistent with the original model mentioned in [13] with anomalies like *scont* which simulates new random continuous transmissions in different bands and sensor locations. This shows that the proposed training process is able to detect new pirate transmitters as discussed in Section III-B. In addition, SAIFE provides an interpretable feature space when compared to other algorithms.

B. Q2: Clustering performance

The average clustering performance of COBRAS for three anomalies (three ground truth clusters) on 8 different frequency bands and four sensors for different SNRs is plotted in Figure 7. The clustering performance measure plotted is Adjusted Rand Index (ARI) [26] which is a symmetric similarity measure between two clusterings, in this case between the produced anomaly clustering and ground truth. As mentioned in Section IV, we make use of samples with three added synthetic anomalies, providing three clusters as our ground truth. The results shown are averaged scores from 10 runs where COBRAS is allowed to query from data from all sensors. The clustering accuracy of high SNR anomalies (20

Fig. 6: Averaged true anomaly detection accuracies for different algorithms on 8 frequency bands from four Electrosense sensors for a constant false alarm rate of 5% (left) and 10% (right).

Fig. 7: Average clustering ARI with queries answered from all four sensors. Results are averaged from 10 runs of COBRAS.

Fig. 8: Average clustering ARI with queries answered from only one sensor for different SNRs. Results are averaged from 10 runs of COBRAS.

Fig. 9: Average clustering ARI with queries answered for different sensor counts. Results are averaged from 10 runs of COBRAS and over 3 SNR values (0, 10 and 20 dB).

Fig. 10: Average clustering ARI with queries answered with different incorrect answering probabilities (P_i). Results are averaged from 10 runs of COBRAS and over 3 SNR values (0, 10 and 20 dB).

and 10 dB) reach above 80% with tens of queries while the low SNR anomalies require queries in hundreds to reach the same. The visibility of anomalies in spectrograms are poor for lower SNRs below 0 dB, which prevents even expert users from detecting these anomalies visually. In the case of -10 dB SNR, the clustering ARI values only reach 11% after 400 querying steps with COBRAS.

In order to test the generalization capability across different sensors the clustering performance is plotted again with only answering queries of anomalies from a single sensor. The results shown in Figure 8 confirms that a few more queries are required even at high SNR anomaly conditions to reach the 80% mark when compared to answering queries from all sensors as expected. This validates that the global anomaly feature space across sensors helps to generalize similar anomalies across sensors. After a few queries, the clustering performance

Fig. 11: Average anomaly detection accuracy with correct labeling with $\alpha, k = 2, 200$. α controls the impact of the user provided labels versus the labels from the clustering process and k determines the depth of label propagation for a single labeled instance. Results are averaged from 10 runs.

saturates and it doesn't keep on improving as in the previous multi-sensor querying scenario. This indeed indicates the need to enable interactive querying from a few different sensors to improve the clustering performance. To further validate this, we include the average ARI across all SNRs (from 0 to 20 dB) by enabling querying from multiple sensors. The results plotted in Figure 9 shows that the clustering accuracy gets better with a fewer queries with more sensors and it keeps on increasing at a faster rate when compared to lower sensor counts.

Further, the clustering performance is evaluated for understanding the effect of wrong query answers which might happen especially in a crowdsourced expert-feedback paradigm. The average ARI for various probabilities of incorrect answering (P_i) are shown in Figure 10. It can be seen that the clustering performance is significantly affected by incorrect answering. Above $P_i = 0.2$ the clustering performance is largely affected. At very high incorrect answering rates it is better to switch to an unsupervised clustering algorithm or adjust the α and k parameters of the anomaly detection algorithm (SSDO) to give low priority to the clustering process.

C. Q3: Anomaly detection performance

The clustered features are fed to the SSDO algorithm for further anomaly labeling as explained in Section III-D. As the signals fed to the detection phase are initially detected anomalies from the data feature module, we randomly select one of the three synthetic anomalies as non-anomalous and then continue the detection procedure to investigate if we can learn that it is not an anomaly. A random instance is chosen from each cluster and shown to the user for confirming as anomalous or not and the process is continued in a round *robin-fashion* for all clusters. From Figure 11, showing the averaged anomaly detection accuracy across all sensors and SNRs, it is clear that with a few labels or queries itself (each query provides single anomaly label), the model can

Fig. 12: Average anomaly detection accuracy with incorrect cluster and anomaly labeling probability $P_i = 0.05$. Results are averaged from 10 runs.

achieve a good detection accuracy above 90%. The speed of convergence is affected by the values of α and k , mentioned in the Section III-D, as expected. With increasing queries, the number of false positives and false negatives keeps on decreasing improving the precision and recall values, as shown in Figure 11. The anomaly detection performance is also analyzed for various parameter values with a incorrect labeling probability of $P_i = 0.05$ both in the clustering and anomaly detection phase. It can be seen that even with a low starting clustering accuracy of 40% at $P_i = 0.05$ the anomaly detection process is able to reach close to 80% accuracy with 500 queries. With wrong labeling the convergence is slower which is also an expected behaviour. More advanced strategies such as *selective repeated-labeling* [27] can be deployed for improving the detection accuracy under noisy conditions.

VI. CONCLUSION

Detecting wireless spectrum anomalies is a challenging problem due to the complexity of wireless spectrum usage and the electromagnetic environment itself. In this paper we have proposed an active framework for wireless spectrum anomaly detection with distributed crowdsourced sensors. The model performs an unsupervised initial anomaly detection and outperforms SoA algorithms. This unsupervised anomaly detection is followed by a semi-supervised active anomaly detection phase where it provides above 95% detection accuracy on the selected anomalies. The active anomaly detection phase is also resilient to noisy labeling and it is shown that the model achieves good accuracy close to 80% with only hundreds of queries. The model also provides good anomaly clustering performance which will be used in future for informing users about relevant anomalies based on user preferences.

ACKNOWLEDGEMENTS

This research was sponsored in part by the department of Science and Technology, armassuisse and the North Atlantic

Treaty Organization (NATO) Science for Peace and Security Programme under grant G5461.

REFERENCES

- [1] P. Denisowski, "Recognizing and resolving lte/catv interference issues," *White Paper, Rohde and Schwarz*, 2011.
- [2] "Viavi," <https://www.rcrwireless.com/20161101/network-infrastructure/interference-hunting-fundamental-quality-service>, accessed: 2018-07-01.
- [3] S. Thombre, M. Z. H. Bhuiyan, P. Eliardsson, B. Gabriellsson, M. Pattinson, M. Dumville, D. Fryganiotis, S. Hill, V. Manikundalam, M. Pölöskey *et al.*, "Gnss threat monitoring and reporting: Past, present, and a proposed future," *The Journal of Navigation*, vol. 71, no. 3, pp. 513–529, 2018.
- [4] "STRIKE3, Initial Findings from the STRIKE3 GNSS Interference Monitoring Network," <https://www.gps.gov/governance/advisory/meetings/2018-05/dumville.pdf>, accessed: 2018-07-01.
- [5] M. Strohmeier, M. Schfer, R. Pinheiro, V. Lenders, and I. Martinovic, "On perception and reality in wireless air traffic communication security," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 6, pp. 1338–1357, June 2017.
- [6] "Telstra," <https://exchange.telstra.com.au/illegal-mobile-repeaters/>, accessed: 2018-07-01.
- [7] M. Matinmikko, M. Latva-Aho, P. Ahokangas, S. Yrjölä, and T. Koivumäki, "Micro operators to boost local service delivery in 5g," *Wireless Personal Communications*, vol. 95, no. 1, pp. 69–82, 2017.
- [8] S. Liu, Y. Chen, W. Trappe, and L. J. Greenstein, "Aldo: An anomaly detection framework for dynamic spectrum access networks," in *IEEE INFOCOM 2009*, April 2009, pp. 675–683.
- [9] W. Honghao, J. Yunfeng, and W. Lei, "Spectrum anomalies autonomous detection in cognitive radio using hidden markov models," in *2015 IEEE Advanced Information Technology, Electronic and Automation Control Conference (IAEAC)*, Dec 2015, pp. 388–392.
- [10] T. J. O'Shea, T. C. Clancy, and R. W. McGwier, "Recurrent neural radio anomaly detection," *arXiv preprint arXiv:1611.00301*, 2016.
- [11] N. Tandiya, A. Jauhar, V. Marojevic, and J. H. Reed, "Deep predictive coding neural network for rf anomaly detection in wireless networks," *arXiv preprint arXiv:1803.06054*, 2018.
- [12] Q. Feng, Y. Zhang, C. Li, Z. Dou, and J. Wang, "Anomaly detection of spectrum in wireless communication via deep auto-encoders," *The Journal of Supercomputing*, vol. 73, no. 7, pp. 3161–3178, 2017.
- [13] S. Rajendran, W. Meert, V. Lenders, and S. Pollin, "Saife: Unsupervised wireless spectrum anomaly detection with interpretable features," *IEEE International Symposium on Dynamic Spectrum Access Networks*, 2018.
- [14] S. Rajendran, R. Calvo-Palomino, M. Fuchs, B. V. den Bergh, H. Corbobs, D. Giustiniano, S. Pollin, and V. Lenders, "Electrosense: Open and big spectrum data," *IEEE Communications Magazine*, vol. 56, no. 1, pp. 210–217, Jan 2018.
- [15] F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation-based anomaly detection," *ACM Transactions on Knowledge Discovery from Data (TKDD)*, vol. 6, no. 1, p. 3, 2012.
- [16] M. M. Breunig, H.-P. Kriegel, R. T. Ng, and J. Sander, "Lof: identifying density-based local outliers," in *ACM sigmod record*, vol. 29, no. 2, ACM, 2000, pp. 93–104.
- [17] A. Makhzani, J. Shlens, N. Jaitly, I. Goodfellow, and B. Frey, "Adversarial autoencoders," *arXiv preprint arXiv:1511.05644*, 2015.
- [18] T. Pevný, "Loda: Lightweight on-line detector of anomalies," *Machine Learning*, vol. 102, no. 2, pp. 275–304, 2016.
- [19] S. Das, W. K. Wong, T. Dietterich, A. Fern, and A. Emmott, "Incorporating expert feedback into active anomaly discovery," in *2016 IEEE 16th International Conference on Data Mining (ICDM)*, Dec 2016, pp. 853–858.
- [20] S. Das, W.-K. Wong, A. Fern, T. G. Dietterich, and M. A. Siddiqui, "Incorporating feedback into tree-based anomaly detection," *arXiv preprint arXiv:1708.09441*, 2017.
- [21] N. Görmitz, M. Kloft, K. Rieck, and U. Brefeld, "Toward supervised anomaly detection," *Journal of Artificial Intelligence Research*, vol. 46, pp. 235–262, 2013.
- [22] V. Vercruyssen, W. Meert, G. Verbruggen, K. Maes, R. Bäumer, and J. Davis, "Semi-supervised anomaly detection with an application to water analytics," in *Proceedings/IEEE International Conference on Data Mining*, IEEE, 2018.
- [23] T. Van Craenendonck, S. Dumančić, E. Van Wolputte, and H. Blockeel, "Cobras: Interactive clustering with pairwise queries," in *International Symposium on Intelligent Data Analysis*. Springer, 2018, pp. 353–366.

- [24] A. Struyf, M. Hubert, P. Rousseeuw *et al.*, “Clustering in an object-oriented environment,” *Journal of Statistical Software*, vol. 1, no. 4, pp. 1–30, 1997.
- [25] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *CoRR*, vol. abs/1412.6980, 2014. [Online]. Available: <http://arxiv.org/abs/1412.6980>
- [26] L. Hubert and P. Arabie, “Comparing partitions,” *Journal of classification*, vol. 2, no. 1, pp. 193–218, 1985.
- [27] V. S. Sheng, F. Provost, and P. G. Ipeirotis, “Get another label? improving data quality and data mining using multiple, noisy labelers,” in *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2008, pp. 614–622.

Sofie Pollin obtained her PhD degree at KU Leuven with honors in 2006. From 2006-2008 she continued her research on wireless communication, energy-efficient networks, cross-layer design, coexistence and cognitive radio at UC Berkeley. In November 2008 she returned to imec to become a principal scientist in the green radio team. Since 2012, she is tenure track assistant professor at the electrical engineering department at KU Leuven. Her research centers around Networked Systems that require networks that are ever more dense, heterogeneous, battery powered and spectrum constrained. Prof. Pollin is BAEF and Marie Curie fellow, and IEEE senior member.

Sreeraj Rajendran is a post-doctoral researcher in the Networked Systems group of KU Leuven. He received his PhD from KU Leuven, Belgium in 2019 and masters in communication and signal processing from IIT, Bombay, India in 2013. He worked as a senior design engineer in the baseband team of Cadence (Tensilica) after completing his masters. He also worked as an ASIC verification engineer in Wipro Technologies during 2007-2010. His main research interests include machine learning algorithms for wireless spectrum awareness and low

power wireless sensor networks.

Vincent Lenders is the head of C4I Networks group and Cyber-Defence Campus at armasuisse. He is also the co-founder and chairman of the executive boards of the OpenSky Network and Electrosense associations. He earned his Ph.D degree (2006) and M.sc (2001) in electrical engineering and information technology both at ETH Zurich, Switzerland. He was also postdoctoral research faculty at Princeton University in the USA. Vincent Lenders is the author of more than 100 publications that appeared in peer-reviewed international conferences and journals and

is the inventor of two patents. He has received best paper awards at IEEE WONS 2012, DFRWS EU 2015, ACM CPSS 2015, and DASC 2015, and the Security Award in 2011 from the Swiss Federal Department of Defense. He is a member of IEEE, ACM, and the expert Jury of the Swiss Economic Forum. He holds various security professional and auditor certifications including CISA, CISM, CRISC (ISACA) and CISSP (ISC2).

Wannes Meert received his degrees of Master of Electrotechnical Engineering, Micro-electronics (2005), Master of Artificial Intelligence (2006) and Ph.D. in Computer Science (2011) from KU Leuven. He is currently research manager in the DTAI research group at KU Leuven. His work is focused on applying machine learning, artificial intelligence and anomaly detection technology to industrial application domains.