Scalable Wireless Anomaly Detection with Generative-LSTMs on RF Post-Detection Metadata

Barnes-Cook, Blake
Electrical and Computer Engineering
Virginia Polytechnic Institute
Arlington, Virginia
blakepc7@vt.edu

O'Shea, Timothy
Electrical and Computer Engineering
Virginia Polytechnic Institute
Arlington, Virginia
oshea@vt.edu

2022 IEEE Wireless Communications and Networking Conference (WCNC) | 978-1-6654-4266-4/22/$31.00 ©2022 IEEE | DOI: 10.1109/WCNC51071.2022.9771754

***Abstract*—Signal anomaly detection is commonly used to detect rogue or unexpected signals. It has many applications in interference mitigation, wireless security, optimized spectrum allocation, and radio coordination. Our work proposes a new method for anomaly detection on signal detection metadata using generative adversarial network output processed by a long short term memory recurrent neural network. We provide a performance analysis and comparison to baseline methods, and demonstrate that through the usage of metadata for analytics, we can provide robust detection, while also minimizing computation and bandwidth, and generalizing to numerous effects which differs from many prior works that focus on A.D. based signal processing on the raw RF sample data.**

***Index Terms*—Signal Processing, Anomaly Detection, Wireless, Elasticsearch, Metadata, Omnisig, Sensing, RF, Analytics, Edge Intelligence, Wireless Security**

I. Introduction

Wireless systems surround the modern world, connecting personal computing devices with cellular and WiFi technologies, personal devices with Bluetooth, emergency responders and infrastructure communications such as GMR, broadcast radio and television, telemetry and communications for air and ground vehicles, localization and timing measurement via GNSS systems, radar systems for safety and guidance, IoT systems for control of countless infrastructure devices, and near endless list of other such systems which share the radio spectrum and operate critical functions within our physical world.

Changes and disruptions to these wireless systems often underlie a wide range of physical-world events, disruption, or phenomena which hold valuable information for analysis, safety, optimization, and other applications. Today, information sources such as social media, traffic systems, camera systems, etc are already analysed for valuable real world event, trend, and safety information, but are limited in their scope of raw information available. Due to the physical broadcast nature of wireless propagation and to the pervasiveness of wireless systems around us, activity in the spectrum has a rich set of information about nearby physical events and activity which is largely untapped today. Largely this has been due to the practicality and technical complexity in providing such analytics across a wide range of raw spectrum data and today. Because each wireless system typically uses a different protocols, physical layers and access scheme, analytics used today are often narrowly focused on one band, technology, or protocol.

Advances in machine learning's ability to detect and analyse a wide range of radio access technologies and events from raw radio time-series data however is rapidly transforming the feasibility, speed, and cost associated with such wide spectrum analytics. To this end, we focus in this work anomaly and change detection over large time periods on a variety of spectrum bands. Specifically, we employ a signal detection edge-processing node, and leverage a large Elasticsearch database to summarize long-time periods of metadata describing signal emissions.

We focus on the task of providing robust anomaly detection on post-signal-detection metadata in order to rapidly identify changes, novel behaviors, or unexpected events within the pattern of life of diverse nearby wireless systems. We employ a generative model approach, using recurrent neural networks to identify unrecognized patterns on the model outputs and associated sequenced metadata. We focus on several novel datasets with labeled anomalies from nearby wireless systems and evaluate performance in terms of precision, recall, AUC metrics, and attained F1 score.

II. Background

A high level diagram is shown in figure 1 which illustrates how a software defined radio (SDR) is used to capture raw digital RF representing activity of a variety of wireless users and emitters, how a wireless activity recognition software is used to extract event summarization from these high rate data streams, and finally how activity change detection and anomaly detection can be conducted on these summary streams in order to recognize events and underlying real world phenomena.

A. Software Radio

Software radios have been widely used as general purpose front-ends to allow capture of raw radio timeseries data, and to move application or protocol specifics into software implementation on a range of computing platforms. While many SDR platforms may be used for this approach, we

Research funded in part by Virginia Center for Innovative Technology

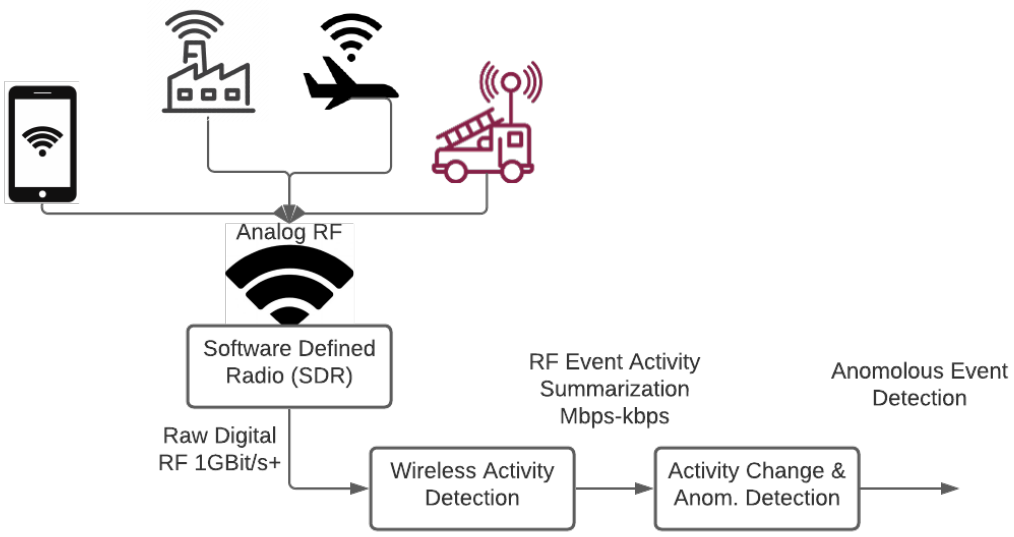

Fig. 1: High level architecture for RF anomaly and change detection.

primarily use a USRP [1] B200 mini for our work, which constrains us to instantaneous capture of around 40 MHz (1.3 GBits) of digital RF data between 100 MHz and 6 GHz over the USB3 bus.

B. High Speed RF Event Recognition

While not the primary focus of this work, event recognition on wideband digital RF data-streams has been a rapidly evolving area of applied research and practice in recent years. Recent competitions such as [1] have explored optimal algorithmic approaches to detecting and classifying RF events in high rate I/Q data-streams and provided open datasets to advance research in this area. Approaches from computer vision and other object recognition domains have been adapted and specialized for this task and can use performant convolutional neural network architectures to provide fast and accurate edge summarization of high rate datastreams into compact metadata formats. This is an important enabler for wireless anomaly detection as heterogeneous, fast, efficient, and accurate wireless signal detectors have often been impractical previously. The latest generation has seen speedups over 100x in recent years using DL based approaches and has gained significant sensitivity often of over 8dB compared with prior approaches such as naive energy detection. [2]. Enabling efficient summarization of RF activity at scale is a powerful enabler for intelligent behavior of spectrum access systems in the future.

C. RF Anomaly & Change Detection

Detection of changes in spectrum activity or anomalous signal behaviors which deviate from a normal pattern of life, are often key events of interest. For example disappearing cell-tower signals, appearance of new unauthorized cellular signals, immediate surges or dips in traffic on push-to-talk (PTT) or ground mobile radio (GMR) networks, or activity on new channels, are all high-level spectrum events which hold significant real-world meaning and may warrant changes or reactions in wireless network configurations, spectrum enforcement actions, or other real world reactions. Deriving these high level detections from a post-detection metadata stream is an appealing approach because it allows for scale and the fusion of many sensors, limiting the bandwidth requirements, storage requirements, and distributing computing to make use of edge processing efficiently. A number of prior works have looked at wireless anomaly detection as a specific short-time task on the highly computationally complex task of processing full-rate input baseband radio signals, such as in [3], [4], [5]. These approaches have demonstrated promising results, but also seem to scale in complexity per-sensor and per-sample rate in a largely unmanageable way, requiring raw radio sample data in order to perform baseline and out-of-distribution detection, limiting their viability or scale for fusing inputs and requiring significant edge-computation to deploy to single sensors.

III. The Problem

Here, we focus on the problem of robust anomaly detection and change detection in post detection RF emission metadata datasets, where a single application agnostic signal detection and classification engine is deployed to the edge which is optimized for efficiency and wideband RF activity summarization to metadata. Here, as shown in figure 1, a front-end radio or SDR is used to obtain a wide range of radio frequency bands, is processed by a wireless emission or activity detections, and then an anomaly or change detection module is used to further

[1]https://www.ettus.com/all-products/usrp-b200mini-i-2/

TABLE I: Anomaly Datasets.

Dataset Format				Training Set	Testing Set	
Dataset Name	Anomaly Type	Signal Types	Sample Duration	Training Samples	Normal Samples	Anomalous Samples
Energy	Feature	LTE	30 seconds	10240(3.5 days)	8(4 min)	2(1 min/1,056 anom. signals)
SNR	Feature	LTE	30 seconds	10240(3.5 days)	10(5 min)	2(1 min/336 anom. signals)
TimeOfDay	Feature	LTE	30 seconds	10240(3.5 days)	10(5 min)	2(1 min/240 anom. signals)
Bandwidth	Position	FM	30 seconds	10240(3.5 days)	8(4 min)	2(1 min/66 anom. signals)
Hopper	Position	FM + DMR	30 seconds	10240(3.5 days)	15(7.5 min)	75(32.5 min/185 anom. signals)
LowerP25	Dip	P25	30 seconds	10240(3.5 days)	90(45 min)	90(45 min of reduced signals)
NoP25	Dip	P25	30 seconds	10240(3.5 days)	90(45 min)	90(45 min of no signals)
NoLTE	Dip	LTE	30 seconds	10240(3.5 days)	90(45 min)	90(45 min of no signals)

analyse the distributions of raw detections from one or more sensors and performs the role of identifying which emissions are out-of-distribution, anomalous, or present large significant changes of interest. Our focus in this work is to develop a highly effective machine learning model for detection of out-of-distribution activity in this post-detection signal activity metadata and in some cases to perform change detection or detection of changes in underlying signal anomalies which change the temporal behavior of these signals in the post detection signal metadata. In table I we outline a number of real world wireless anomalies of interest which we focus our approach on, which represent a range of changes in signal characteristics across cellular and GMR spectrum bands including changes such as received signal strength, signal quality, signal activity over time, carrier presence or absence and a range of other anomalies which such a general purpose signal metadata anomaly detection engine should be able to rapidly flag on a wide range of bands and locations.

IV. Representing Metadata

Recently the Signal Metadata Format (SigMF) [6] has been widely adopted as an open source standard for storage of both raw RF data with annotations and metadata, as well as for pure metadata only storage of descriptions of signals or RF environments. We use this format to store and process metadata acquired from sensors. These records are stored both in pure JSON file formats and an Elasticsearch database.

A. JSON File

In this format, SigMF records are stored in groups according to the time the sensor recorded the signals. In order to analyze each signal in the JSON file each group needs to be read and parsed. Due to this inability to filter, each dataset in this format is meant to be used as standalone training scenario. These datasets are useful for quick model validation as they can be heavily compressed and shared. Sharing these files as output from sensors is not feasible.

B. Elasticsearch Database

Each signal is stored separately in the Elasticsearch database. The nature of Elasticsearch allows the querying of data from any time period with filters corresponding to channels. New signal data can also be sent to a particular index, updating the dataset in real time. This allows training and testing on live data with any combination of channels, bandwidth filters, and sensor selection. Querying Elasticsearch does require more throughput and computational overhead than the JSON file based approach.

V. Anomaly Types and Associated Datasets

To perform testing on the models, we prepared several different datasets containing various types of anomalies. These datasets are derived from a base dataset of 17.5 million signals captured over a period of 7 days. In this case, ground truth background data is used from real world recordings, while in each case a form of synthetic anomaly is added to this data to ensure known ground truth. Table I describes each dataset with the anomalous feature generated, associated anomaly type, anomalous signal type, duration of each sample, total amount of training samples, total benign testing samples, and anomalous testing samples with the total amount of anomalous signals across the samples.

A. Feature Based Anomalies

Feature based anomalies add signals to a specific time period that have anomalous feature values compared to the rest of the dataset. These features are anything that the sensor can detect and note in the SigMF entry. Anomaly examples include abnormally high signal to noise ratio, high receive signal strength, signals occurring at a strange time of day, or variations in the confidence of the type of signal detected.

B. Position Based Anomalies

Position based anomalies are similar to feature based anomalies, though differ in that the only anomalous features are changes in either signal bandwidth and frequency (i.e. they have moved in their spectral location). This may occur for instance if a radio changes its rate or bandwidth configuration, or a signal appears at a new frequency or channel location, or hops over multiple frequencies or locations. Detection is treated somewhat different for these types of anomalies as a 2D density feature is produced in order to aggregate windows of detections for some time interval.

C. Dip Based Anomalies

Dip based anomalies feature a reduction or complete absence in observations of certain signals during anomalous time periods. Signal features generally remain constant (e.g. RSSI, SNR, BW), however observation frequency or present dips

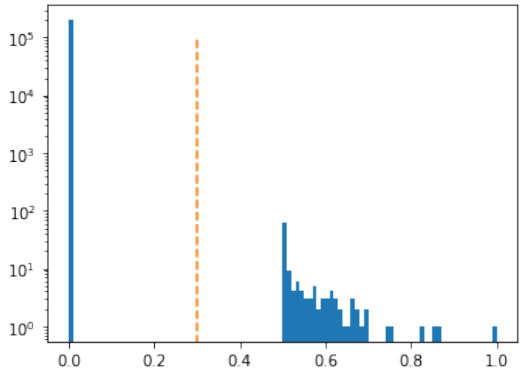

Fig. 2: Normalization Result of Preprocessing for LSTM GAN.

to abnormal levels indicating the disappearance, failure, or reduction in activity from a certain radio emitter.

VI. Anomaly Detection Methods

We train and test models using the Pytorch Lightning [7] software platform. Our Pytorch datasets are customized for this application and models are designed to retrieve and process specific frequency bands, sensors, and time periods within an Elasticsearch database containing numerous sensors across numerous bands spanning months of metadata signal detection and classification data. Additionally, some subsets of these records have been annotated with "known" anomalies of various types within the data, which correspond to a variety of types of anomalous RF activity within these datasets.

A. Adversarial Dual Autoencoder with LSTM Based Scoring

1) Preprocessing & Features: Preprocessing begins by computing a two dimensional density feature vector of signals in a 400x500 shape with signal bandwidth and signal center frequency as the axes for each channel. Bandwidth is binned using a logarithmic bin spacing to fit max bandwidth while maintaining resolution while frequency is linearly binned and limited to a minimum and maximum value appropriate for the dataset and band. Density functions are evaluated here over 30 second time windows, however other time resolutions are possible. Density reflects the total on-time of each bin over that window or interval. Density values are normalized with a minimum of 0.5, and a maximum of 1.0, while bins that do not have any signal values remain at zero. Figure 2 shows the normalization result of a particular time segment.

Because thousands of these density functions or equivalent-images are required in order to train the model, the dataset quickly becomes too large to initialize at our desired resolution and time period. It can be seen in figure 2, each image consists of mostly zeros post-normalization. We therefore leverage sparse tensors to store and represent the data, and to decrease data storage by several orders of magnitude. This allows us to download and initialize datasets of over 7 days of raw SigMF emission data extremely rapidly and with significant storage, memory and bandwidth reductions vs the full density functions and even more so the original digital signal rate by many orders of magnitude.

2) Model Architecture : The GAN portion of this network is based on a adversarial dual autoencoder (ADAE) approach which is shown to be very effective in detecting anomalous regions within brain MRIs in density distributions or images [8]. In this model, the image is first passed through the generator and discriminator autoencoders (i.e. the ADAE). Image features and regions that are recognized (i.e. in-distribution) from training process are present in the output of the ADAE, while out of distribution unexpected features are excluded from this reconstruction. The output is then evaluated based on the known populated indices of the input data, sorted by the lowest value. For 10 of the lowest values, a kernel of size 3x3 is extracted from the ADAE output image centered around the values. This is then compiled into a list to pass into a recurrent LSTM portion of the network, which will memorize temporal patterns of the extracted kernels or modes and map them to index values during training. While testing, the outputs of the LSTM are compared with the correct indices, and the MSE of the difference is used as the index anomaly score.

The goal of this network is to identify areas of the ADAE output that were expected to be reconstructed but were not. Given that both the input and output are sparse, using a basic loss function such as MSE directly on the full-dimension 2D density input and output vectors does not yield a useful anomaly score, which motivates this index based approach.

B. Feature Based Anomaly Detection using LSTM

1) Preprocessing : Preprocessing for the feature based anomaly model first involves determining which features (RSSI, SNR, energy estimate, etc.) to include from the dataset. Selected features are then averaged for all signals in the analyzed time period for each channel (LTE, FM, etc.). If a given channel has captured no signals in the time period selected, all features in the associated array will have a value of 0.

Since this data only includes a set number of features for each time period analyzed, it does not benefit from sparse tensors. This data also takes up significantly less space than the LSTM GAN dataset.

2) Model Architecture : This network is composed of a single LSTM layer that predicts the average of each feature for a specific time period. The feature arrays of 9 previous time windows are compiled into a list and processed by the LSTM layer. The output is a predicted feature array that is compared with the retrieved feature array of the current time period. Each feature is individually evaluated using MSE to create an anomaly score for that feature.

This model was designed to work in conjunction with the LSTM GAN. Due to the normalization process required by the ADAE network, some details of the data would be lost (e.g. low intensity modes in the distribution) and anomalies such as the absence of modes which are deemed to be regular in the input. This would make the network effectively blind to many feature changes, such as those represented in the dip based

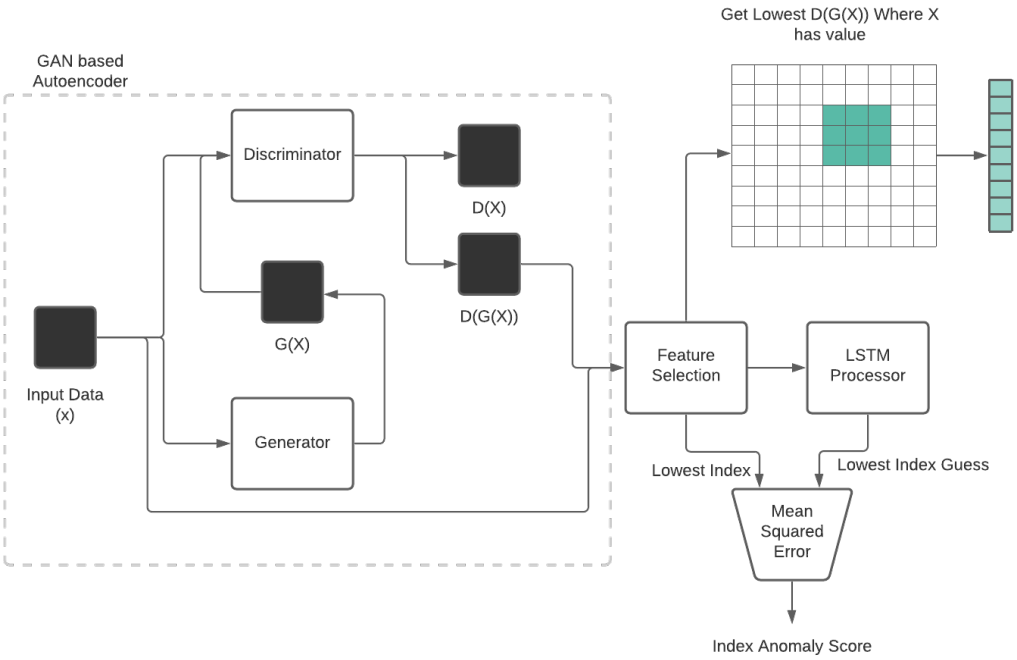

Fig. 3: Architecture for the LSTM GAN.

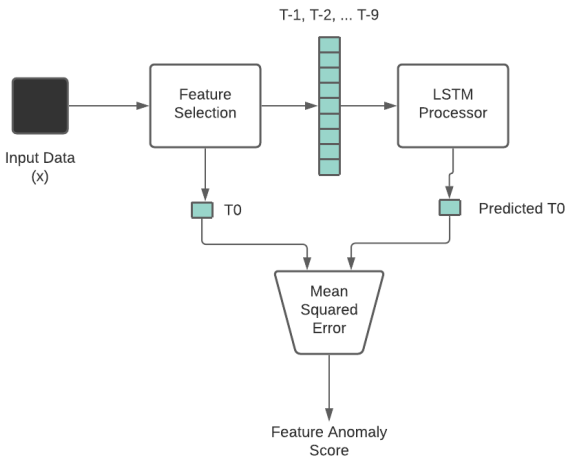

Fig. 4: Architecture for Feature LSTM.

anomaly datasets or when a signal which is regularly present disappears suddenly. This network then detects changes in features in such cases.

VII. Experiments

Both the LSTM GAN based model and the LSTM Feature based model were trained concurrently for each dataset. Training was performed for the results shown on 3 days of baseline emission data prior to the anomalous time-period and occurred for 100 epochs with a training/validation data split of 75/25 across an 8 GPU system. Once fully trained, the network is evaluated through the anomalous region in 30 second time intervals, recording anomaly scores for each time period. A post-processing method was used to then select an optimal threshold that produced the highest F1 score performance. Precision, Recall, ROC curves, and F1 scores were then generated using the anomaly threshold utilizing the scikit-learn [9] library.

VIII. Results

Results for anomaly detection performance of the proposed network approach are shown in table II. For each dataset and/or type of anomaly, the index anomaly score generated by the LSTM GAN can be seen. The best performing feature score sourced from the LSTM feature based model is also shown for each dataset, along with the feature used to generate the score values. Performance is shown using Precision, Recall, F1 Score, and the area under the ROC curve. The anomaly threshold used to generate the results is also listed for each score, which varies to some degree based on type of anomaly and band.

In each case recall is excellent for nearly all datasets for the LSTM GAN, while precision is fairly low for some. This conforms with the intended use of this network to identify suspicious sections or RF events which are out of distribution, while ignoring normal occurrences in the normal RF emission distribution. The downside of this approach is the incorrect classification of normal signals. An ensemble approach that factors in output from both networks could be used to further reduce false positives and improve performance.

TABLE II: LSTM GAN and Feature LSTM Performance

Dataset	LSTM GAN					Feature LSTM					
	Precision	Recall	F1score	AUC	Threshold	Feature Type	Precision	Recall	F1score	AUC	Threshold
Energy(LTE)	1.0	1.0	1.0	1.0	10130.33	Confidence	1.0	1.0	1.0	1.0	0.07
SNR(LTE)	0.25	1.0	0.4	0.7	8333.17	SNR Estimate	1.0	1.0	1.0	1.0	4.47
TimeOfDay(LTE)	0.67	1.0	0.8	0.95	11699.62	Count	1.0	1.0	1.0	1.0	6.31
Bandwidth(FM)	0.33	1.0	0.5	0.75	22.24	Confidence	1.0	0.5	0.67	0.75	0.04
Hopper(FM)	0.51	1.0	0.68	0.5	3.70	RSSI	0.51	0.93	0.66	0.49	0.06
Hopper(DMR)	0.51	1.0	0.68	0.5	52.77	Confidence	0.46	0.69	0.56	0.45	0.01
LowerP25(P25)	0.87	0.82	0.86	0.86	176.25	Count	0.51	0.95	0.67	0.51	0.24
NoP25(P25)	0.87	0.85	0.86	0.85	186.14	Confidence	0.78	0.85	0.81	0.80	0.05
NoLTE(LTE)	0.51	1.0	0.67	0.50	146.44	Count	0.87	0.85	0.86	0.86	16.83

IX. Future Work

Numerous opportunities remain for future work in this area, which has been relatively lightly explored in the past. These include feature space and representations conveyed from the signal detection and classification engine up front, representations of the sparse signal detection data and time-windowed density functions thereon, network architectures for out-of-distribution detection and temporal change detection thereon, as well as methods for per-band calibration and detection optimization on the metrics detected from such a system.

In this work, we observed that threshold selection per-band and in some cases per-anomaly type could produce excellent precision while preserving recall. However, more work is needed to automate this process without confirmed supervised anomaly training on new bands and without known anomaly classes. Further, expanding the approach proposed here to combined multiple anomaly scores to identify an individual anomalous signal or event in each time-window or interval is also a major focus in future work which could significantly improve future accuracy.

Our future work also includes deployment of this anomaly detection for streaming, real-time inference and detection on live data – to this end, we are extending a version of this model to subscribe to live Elasticsearch index data to provide real-time predictions of anomalous or out-of-distribution events in RF detection metadata feeds. Using this approach, it is also possible to scale rapidly to numerous sensors and bands as they are deployed in volume without unmanageable bandwidth or computational burdens. This will allow real-time RF anomaly and change-detection monitoring of numerous bands in multiple locations using a decentralized sensor model as well as a distributed data and processing model.

X. Conclusion

Broad-spectrum monitoring of numerous wireless communications technologies is an extremely nontrivial task, especially when scalability, generality, and efficiency are key. Using multiple software defined radios paired with a wireless signal detectors to identify and summarize wireless activity into signal metadata has opened up several new avenues of monitoring and detection, allowing for extensive analysis and processing on high level signal behaviors across a range of technologies. By leveraging this class of post-detector signal detection and summarization metadata in a large scale platform such as we used in Elasticsearch and our proposed model, it is now feasible to analyze extremely large volumes of RF emission activity across several channels or bands from various a range of unique locations in real-time. This provides a significant novel benefit to traditional wireless anomaly detection methods previously explored using raw RF timeseries data, as it allows for significantly more scale in the number of sensors, types of anomalies, time-scale which can be digested, and practicality of deployment and usage for real world systems. Overall, we believe this class of approach has demonstrated impactful initial performance results for a relatively new task which was previously intractable, and does so in a scalable and efficient way which provides significant value in terms of alerting on important events. Significant room remains for refining the approach, as this is a relatively new deployment model, however we believe these results represent a strong baseline approach for this task demonstrating its value and feasibility.

References

[1] N. West, T. O'Shea, and T. Roy, "A wideband signal recognition dataset," in *2021 IEEE 22nd International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*. IEEE, 2021, pp. 6–10.

[2] S. Rajasegarar, C. Leckie, and M. Palaniswami, "Anomaly detection in wireless sensor networks," *IEEE Wireless Communications*, vol. 15, no. 4, pp. 34–40, 2008.

[3] S. Rajendran, W. Meert, V. Lenders, and S. Pollin, "Saife: Unsupervised wireless spectrum anomaly detection with interpretable features," in *2018 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*. IEEE, 2018, pp. 1–9.

[4] N. Tandiya, A. Jauhar, V. Marojevic, and J. H. Reed, "Deep predictive coding neural network for rf anomaly detection in wireless networks," in *2018 IEEE International Conference on Communications Workshops (ICC Workshops)*. IEEE, 2018, pp. 1–6.

[5] T. J. O'Shea, T. C. Clancy, and R. W. McGwier, "Recurrent neural radio anomaly detection," *arXiv preprint arXiv:1611.00301*, 2016.

[6] B. Hilburn, N. West, T. O'Shea, and T. Roy, "Sigmf: the signal metadata format," in *Proceedings of the GNU Radio Conference*, vol. 3, no. 1, 2018.

[7] W. Falcon and T. P. L. team, "Pytorch lightning," 3 2019. [Online]. Available: https://www.pytorchlightning.ai

[8] H. S. Vu, D. Ueta, K. Hashimoto, K. Maeno, S. Pranata, and S. M. Shen, "Anomaly detection with adversarial dual autoencoders," 2019.

[9] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.