Published in final edited form as:

IEEE Trans Cogn Commun Netw. 2019 ; 5: . doi:10.1109/TCCN.2019.2899871.

Deep Learning Classification of 3.5 GHz Band Spectrograms with Applications to Spectrum Sensing

W. Max Lees, Adam Wunderlich, Peter Jeavons, Paul D. Hale, and Michael R. Souryal
Communications Technology Laboratory, National Institute of Standards and Technology, Boulder, CO, USA and Gaithersburg, MD, USA

Abstract

In the United States, the Federal Communications Commission has adopted rules permitting commercial wireless networks to share spectrum with federal incumbents in the 3.5 GHz Citizens Broadband Radio Service band. These rules require commercial systems to vacate the band when sensors detect radars operated by the U.S. military; a key example being the SPN-43 air traffic control radar. Such sensors require highly-accurate detection algorithms to meet their operating requirements. In this paper, using a library of over 14,000 3.5 GHz band spectrograms collected by a recent measurement campaign, we investigate the performance of thirteen methods for SPN-43 radar detection. Namely, we compare classical methods from signal detection theory and machine learning to several deep learning architectures. We demonstrate that machine learning algorithms appreciably outperform classical signal detection methods. Specifically, we find that a three-layer convolutional neural network offers a superior tradeoff between accuracy and computational complexity. Last, we apply this three-layer network to generate descriptive statistics for the full 3.5 GHz spectrogram library. Our findings highlight potential weaknesses of classical methods and strengths of modern machine learning algorithms for radar detection in the 3.5 GHz band.

I. Introduction

Spectrum sensing, defined as the task of ascertaining spectrum usage and the activity of primary users at a given location and time, is a key component of opportunistic spectrum usage by cognitive radio networks [1], [2]. Prior work on spectrum sensing algorithms has investigated methods from classical signal detection theory, such as energy detection and matched filtering [1], [2], classical machine learning techniques such as support vector machine (SVM) and K-nearest neighbor (KNN) classifiers [3]–[7], and distributed (or cooperative) sensing approaches involving multiple sensors [1], [2], [4], [6], [8]–[10]. Recently, deep learning algorithms, such as convolutional neural networks (CNNs) and long short-term memory (LSTM) recurrent neural networks [11], which are the current state-of-the-art for many classification tasks [11], [12], have been applied to various areas in wireless communications, including spectrum sensing, e.g., [13]–[16].

In a key step towards more efficient use of radio spectrum in the United States, the U. S. Federal Communications Commission (FCC) has adopted rules for the Citizens Broadband

Radio Service (CBRS) that permit commercial wireless usage of the 3550-3700 MHz band (the “3.5 GHz band”) [17]. The CBRS architecture outlined in the FCC rules includes a spectrum access system (SAS) together with environmental sensing capability (ESC) detectors to facilitate spectrum sharing in the 3550-3650 MHz band. The purpose of the SAS is to coordinate commercial-user CBRS access so that federal incumbents are given priority access.

The primary federal incumbents in the 3.5 GHz band are shipborne and ground-based radars operated by the U.S. military [18]. The CBRS framework requires that ESC sensors detect these radars, including the SPN-43 air traffic control radar [19], also identified as Shipborne Radar 1 in [18]. ESC detection capabilities are determined by intended and unintended emissions, as well as background noise. For example, out-of-band emissions (OOBE) from an adjacent-band U.S. Navy radar, identified as Shipborne Radar 3 in [18], are prevalent [20]–[23], and could complicate SPN-43 detection.

As federal agencies collaborate with industry to refine standards and requirements for ESC detectors, sound methods for evaluation of ESC detector performance must be developed and potential limitations should be understood. Furthermore, efforts to design ESC detectors could benefit from detection algorithm comparisons and the characterization of emissions in the 3.5 GHz band.

In this paper, we address the above needs with a study of over 14,000 3.5 GHz band spectrograms recorded by a recent measurement campaign [22], [23] at two coastal locations: Point Loma, in San Diego, California and Fort Story, in Virginia Beach, Virginia. Because the hardware required to record and process spectrograms is less complex, and therefore cheaper than that required for in-phase and quadrature (I/Q) data, the use of spectrograms by ESC sensors is an attractive option. For this reason, our investigation focuses on SPN-43 detection from low-resolution spectrograms.

For the task of narrowband SPN-43 detection in a single 10 MHz channel observed by one receiver, we compare the effectiveness of thirteen detection algorithms, including eight deep learning methods, three classical machine learning approaches, and two strategies based on classical signal detection theory. In addition, for the task of wideband SPN-43 detection across multiple channels observed concurrently with one receiver, we compare the top-performing methods from the single-channel evaluation. A thorough performance evaluation utilizing a test set of unverified, human-labeled spectrograms reveals that deep learning methods outperform other approaches for SPN-43 detection. Last, we apply the best-performing deep learning method to classify the complete set of spectrograms collected in San Diego and Virginia Beach with respect to SPN-43 presence, from which we estimate SPN-43 spectrum occupancy and characterize the power of non-SPN-43 emissions.

II. 3.5 GHz Spectrograms

As described in [22], [23], 3.5 GHz band measurements were collected for a period of two months at each measurement site. The primary aim of the measurements was to acquire high-fidelity, time-domain recordings of SPN-43 radar waveforms in the 3.5 GHz band. For

this purpose, a 60-second, complex-valued (i.e., I/Q) waveform was recorded roughly every ten minutes with a sample rate of 225 MS/s, and a corresponding spectrogram was computed. The decision to retain a given waveform recording was made by comparing the spectrogram amplitudes to a threshold over the band of interest. Although only a subset of the waveforms was retained for long-term storage, the entire set of spectrograms was saved.

In total, 14,739 spectrograms were collected over the measurement campaign¹. Of these, approximately 58% were acquired in San Diego and 42% in Virginia Beach. At each measurement site, data were collected with both an omni-directional antenna and a directional, cavity-backed spiral (CBS) antenna. Roughly 45% and 55% of the spectrograms were acquired with the omni-directional and CBS antennas, respectively. The spectrograms span a 200 MHz frequency range, typically 3465-3665 MHz, and a time-interval of one minute. Each spectrogram has dimensions 134×1024, with 134 time-bins of duration 0.455 seconds and 1024 frequency-bins of length 225 MHz/1024 $\approx$ 220 kHz.

The spectrograms were computed by applying a short-time Fourier transform (STFT) and then retaining the maximum amplitude in each frequency bin (i.e., a max-hold) over each 0.455 second time-epoch. The window function for the discrete STFT was 1024 samples long, with the middle 800 points given a weight of one, and the left-most and right-most 112 points weighted with a cosine-squared taper. The STFT was implemented with 112-sample overlap between consecutive time-segments. Each spectrogram value is the maximum of 10⁵ time-averaged amplitudes, where the averaging duration is 4.55 μ s, because the STFT effectively averaged over a 1024 sample (4.55 μ s) time window.

When noted below, the spectrogram values were converted to power units (dBm) as follows. Each max-hold spectrogram value was (i) divided by 1024, the STFT window length, (ii) divided by the (site-specific) front-end gain, (iii) multiplied by a measurement instrument calibration factor, (iv) squared and divided by 2×50 (the 2 arises from the conversion between peak and root-mean-square (RMS) voltage for a narrowband signal, the 50 is for a 50 ohm load), and (v) converted to decibel-milliwatts (dBm) via the formula $10 \log_{10}(P) + 30$, where P is the power in Watts from step (iv). For some calculations, the spectrogram values were further converted to dBm/MHz by subtracting -6.2 dB/MHz, the effective noise bandwidth of the time-domain window for the STFT [22, p. 32].

Figure 1 contains example spectrograms, cropped to the 3550-3650 MHz band of interest for ESC detection. The original full-bandwidth (approx. 3465-3665 MHz) versions of these cropped spectrograms are shown in Figures 3.6 and 3.14 of [23], and Figure 3.18 of [22], for the left, middle and right spectrograms, respectively. In each spectrogram, leakage from the local oscillator of the receiver is faintly visible as a vertical line at 3577 MHz (Left and Middle) and 3565 MHz (Right). The left spectrogram contains a clear SPN-43 radar emission, located at approximately 3570 MHz. Periodic radar sweeps are visible roughly every 4 seconds, corresponding to the SPN-43 antenna rotation period. The middle and right spectrograms of Figure 1 give examples of coincident Radar 3 OOB and SPN-43. In these

¹The set of 3.5 GHz spectrograms used in this work is currently designated "For Official Use Only" (FOUO) by the U.S. Department of Defense, and is not publicly available.

images, Radar 3 OOBs are visible as horizontal streaks, and weak SPN-43 emissions are visible at 3570 MHz (Middle) and 3550 MHz (Right), respectively.

The goal of this work was to create classifiers to identify SPN-43 presence in spectrograms. In order to train and test classifiers, we needed a set of SPN-43 labeled data. From the complete collection of 14,739 spectrograms, 4,491 were labeled by one of the co-authors for SPN-43 presence and Radar 3 OOBs. Note that the human-applied labels are unverified, and based on subjective visual interpretation, as we did not have access to ship locations or assigned frequencies during or after the measurements. Nearly 74% (3,318) of the labeled spectrograms were selected for labeling because they correspond to captures that triggered retention of a recorded waveform. Since this subset suffers from selection bias, an additional 1,173 spectrograms were randomly selected for human-labeling to provide a more diverse set of labeled cases.

III. Classifiers

This section gives implementation details for the thirteen classifier models that we evaluated for SPN-43 detection. Specifically, we compared two methods based on classical signal detection theory, three classical machine learning algorithms, and eight deep learning algorithms to the task of detecting SPN-43 in spectrograms. Because we had a high degree of certainty about the data labels, we focused mainly on supervised learning methods.

For our classical signal-detection algorithms, we chose standard energy detection and sweep-integrated energy detection, which combines data integration with energy detection. Further details can be found later in Section III-A. The classical machine learning algorithms that were evaluated included Support Vector Machines (SVM), K-Nearest Neighbor (KNN) classifiers, and Gaussian Mixture Models (GMM). Details are provided in Section III-B.

The deep learning methods included six well-known CNN architectures: VGG-16 and VGG-19 [24], ResNet-18 and ResNet-50 [25], the Inception-V1 network [26], and DenseNet-121 [27]. In addition, we implemented a CNN and an LSTM [28] of our own design. For further details on the well-known architectures, see Section III-C. Details about our CNN and LSTM models are provided in Section III-E and Section III-D, respectively.

The deep learning algorithms were implemented using the open-source TensorFlow™ Python library running on an Nvidia® DGX™ workstation with four Tesla® V100 Graphics Processing Unit (GPU) cards². Details on standard deep learning layers and other elements, including their formulas, can be found in the textbook by Goodfellow et al. [11].

Because the ESC detection task requires SPN-43 detection in each 10 MHz channel, the classifiers were first designed for the reduced task of detecting SPN-43 in a single 10 MHz channel. For this purpose, spectrograms were divided into 10 MHz-wide channels centered

²Certain commercial software and hardware products are identified to fully specify our implementation. This does not imply endorsement by the National Institute of Standards and Technology or that the software and hardware are the best available for the purpose.

at multiples of 10 MHz, e.g., 3550 MHz, 3560 MHz. A single instance of each machine learning classifier was trained using a random sample of cases across all 10 MHz channels covering the 3550-3650 MHz band; further details on the training set are provided in Section IV-B. Subsequently, to classify all 10 MHz channels over the 3550-3650 MHz band, copies of each classifier instance were applied in parallel, i.e., the same previously-trained classifier instance was applied to each channel. In our preliminary investigations, we explored the possibility of using a single multi-channel detector rather than multiple single-channel detectors in parallel, but found that the latter improved detection performance. Note that the SPN-43 classification results for the eleven channels covering 3550-3650 MHz were not fused for further decision-making.

A. Classical Signal Detection Methods

From the family of classical signal detection methods, we evaluated two different approaches suitable for incoherent detection from spectrograms. Note that techniques intended for coherent detection from in-phase and quadrature (I/Q) data were not possible with our data. The first signal detection method was standard energy detection [29]–[31]. The second method consisted of energy detection combined with data integration [32, Sec. 1.4.6], a method commonly used in radar signal processing to increase signal-to-noise ratio (SNR). Since this second method integrated the data over radar sweeps, we call it sweep-integrated energy detection.

Energy detection is a classical strategy based on the assumption that a signal of interest can be detected based on the total energy across a given time and frequency range. The total energy across the entire input is summed. If a given detection threshold by this summation is exceeded, the signal is decided to be present.

To improve the performance of energy detection, the whole 10 MHz channel was not used. Instead, only the 3 middle spectrogram columns (approximately 660 kHz) of each 10 MHz channel were aggregated for energy detection. This range was chosen based on the results of an empirical evaluation. Because SPN-43 can generally be expected to have carrier frequencies near multiples of 10 MHz, this modification excluded confounding emissions from the rest of the channel.

For sweep-integrated energy detection, we performed energy detection on specific parts of the input that would have highest SNR in the presence of SPN-43 and, therefore, hold the most information about signal presence. Sweep-integrated energy detection is a form of data integration [32, Sec. 1.4.6], a method commonly applied in radar signal processing. First, a SPN-43 sweep template was generated, where the distance between sweep peaks was roughly 3.85 s [22], [23]. This template consisted of a square wave in which the value was one for 0.455 s every 3.85 s and zero otherwise. The point of highest cross-correlation was used to align the sweep template and the spectrogram. Finally, the aligned template was applied as a mask to extract the portion of the input to use for standard energy detection, implemented as above.

As detailed in [22], [23], the spectrogram captures were collected with different front-ends at the two measurement sites. To account for this fact, we applied site-dependent corrections

to normalize the spectrograms to dBm units, as described in Section II. This normalization was only used for these two energy detection algorithms; the machine learning methods did not require any data normalization.

Note that the two energy detection methods described above incorporate *a priori* information. First, both algorithms rely on the fact that SPN-43 typically has a carrier frequency that is a multiple of 10 MHz. Second, the sweep-integrated energy detection method uses the fact that SPN-43 radar has a known sweep period.

B. Classical Machine Learning Methods

For classical machine learning algorithms, we evaluated KNNs, SVMs, and GMMs [33], [34].

The KNN classification relies on similar data being spatially colocated within a chosen representational space. When a new sample is classified, the most common label of the k -nearest labeled samples is assigned. We evaluated the performance of all k values for $k \in \{2, 5, 9, 12\}$.

The SVM classification attempts to maximize the distance between two classes and a parameterized hyperplane that separates the two classes, again, in some representational space. A sample is assigned a label based on which side of the hyperplane it lies. In addition to this standard type of SVM, called a linear SVM, distances can be measured using a kernel function which defines distances in alternate feature spaces. We evaluated the performance of linear SVMs, as well as SVMs with Radial Basis Function (RBF), polynomial, and sigmoid kernel functions [33, Sec. 14.2].

While GMMs are unsupervised, we examined them because it was one of the methods explored in [4]. A GMM fits a mixture of normal distributions with the expectation maximization algorithm. New samples are assigned the label with the highest posterior probability. Because the task was binary detection, the GMM was fit with a mixture of two normal distributions.

For each of the above algorithms, we evaluated the performance on the full channel and with two forms of data preprocessing. For the entire channel as input, the input was $134 \times 46 = 6164$, because there were 134 time steps and 46 frequency bins per channel.

For the preprocessing, the goal was to reduce the dimensionality of the input. In the first form of preprocessing, we applied energy aggregation across frequencies in the 10 MHz channel by summing the energy in each time step, which reduced our input vector length to 134. In the second form of preprocessing, we used prior information about SPN-43 for dimensionality reduction. Namely, because the carrier frequency for SPN-43 is typically a multiple of 10 MHz, the middle two frequency bins were extracted from the entire channel, resulting in an input size of $134 \times 2 = 268$.

To signify which input we used for a given evaluation, we use subscripted model names. For example, KNN models are denoted as KNN_{6164} , KNN_{134} , and KNN_{268} , for the full, first preprocessed, and second preprocessed input forms, respectively.

C. Standard Deep Learning Classifiers

We evaluated the six standard deep learning CNNs listed above. We chose these standard networks because most of them have been winners of a well-known competition using the ImageNet [35] benchmark dataset; Liu et al. [36] cited these networks as milestones in image classification.

VGG-16 and VGG-19 [24] were the first two models we evaluated. VGG-16 contains 13 convolutional layers, 3 fully connected layers, and 5 pooling operations. VGG-19 contains 16 convolutional layers, 3 fully connected layers, and 5 pooling operations. The VGG networks are powerful CNNs that, unlike more recent networks, only consist of convolutional layers, fully connected layers, and pooling operations.

Next, we evaluated ResNet-18 and ResNet-50 [25]. ResNet-18 consists of 17 convolutional layers, 2 pooling operations, and 1 fully connected layer. ResNet-50 consists of 49 convolutional layers, 2 pooling operations, and 1 fully connected layer. The ResNet networks introduce residual connections between blocks of convolutional layers which sum the outputs of a block with the input of the block. The benefit of these residual connections is the increased flow of information through backpropagation.

In addition, we evaluated Inception-V1 [26], which consists of 59 convolutional layers, 16 pooling operations, and 7 fully connected layers along with 2 local-response normalization operations. This network was designed to handle different-sized features within the input feature space at each level of convolutional processing. As a result, it can handle a broad range of tasks with minimal tuning.

The final network we evaluated was DenseNet-121 [27], which contains 120 convolutional layers, 5 pooling operations, and 1 fully connected layers. DenseNet has slightly worse performance on most benchmarks. However, the number of parameters required to achieve similar performance compared to other standard networks is significantly smaller. For each convolutional layer in a convolutional block, DenseNet concatenates the outputs and then downsamples using 1×1 convolutions.

See the papers cited above for further details.

D. Long Short-Term Memory Recurrent Neural Network

Figure 2 (right) summarizes our LSTM architecture. In order to effectively use the LSTM, we split the 10 MHz channel along the time axis to create sequential slices of approximately 0.455 seconds. Each of the time-slices was fed into the LSTM cell one at a time along with the previous output of the LSTM cell. This is known as a residual connection. Motivations for using residual connections in LSTMs include protection against the vanishing gradient problem in backpropagation [37] as well as greater network expressivity [25]. Dropout was used between LSTM cells with a probability of 50%.

After all of the time slices were fed into the LSTM, the output of the last cell was passed on to a fully-connected layer with 50 neurons. Next, a bias was added to the output of the 50-neuron fully-connected layer and a ReLU was applied. The output of the ReLU was then

passed to a fully-connected layer of size 1 and a bias was added. Lastly, a sigmoid activation was applied to the output to generate the prediction, a continuous-valued number between zero and one. The LSTM was trained using stochastic gradient descent with crossentropy loss; further details on training are given in Section IV-B.

E. Convolutional Neural Network

Figure 2 (left) summarizes our CNN architecture, henceforth referred to as CNN-3 for its three layers, as explained below. First, the 10 MHz channel was passed through an average pooling operation with a window size of 10×2 , resulting in a down-sampled spectrogram with time and frequency dimensions reduced by factors of 10 and 2, respectively. The down-sampled spectrograms were then passed to a convolutional layer with 20 filters of size 3×3 and stride 1×1 . Zero-padding was not used. Subsequently, a bias (i.e., constant) was added to the filter activations and a rectifier linear unit (ReLU) was applied to the resulting output. The output of the ReLU step consisted of 20 activation maps for each of the convolutional layer's filters.

Next, the activation maps were averaged together to create a single averaged-activation map using an operation we call channel-average pooling. Specifically, this operation averages co-located values across all channels and averages them to output a single channel.

Equivalently, this operation can be described as a single-filter 1×1 convolutional layer [38] in which all weights are equal and sum to 1. To our knowledge, this averaging step has not been suggested previously in the CNN literature although it does resemble other channel pooling methods [39].

Note that above channel-average pooling operation is distinct from conventional average pooling, which takes localized areas from within a single channel, where a channel is a depth slice of the output from a convolutional layer, and outputs the same number of channels with average values for these areas. We found that using the averaged activations rather than the individual activation maps showed empirical improvements in accuracy.

The output of the channel-average pooling operation was passed into a fully-connected layer containing 150 neurons. A bias was added to the output of this fully-connected layer and a ReLU was applied. The output of the ReLU was then fed through a dropout step, with a dropout probability of 50%. Subsequently, the output from the dropout step was fed into another fully-connected layer containing a single neuron followed by a bias. Finally, the biased output was passed through a sigmoid activation function to produce the prediction, a continuous-valued number between zero and one. CNN-3 was trained using stochastic gradient descent with cross-entropy loss; further details on training are given in Section IV-B.

IV. Classifier Performance Assessment

As noted in Section II., a set of 4,491 spectrograms were labeled for SPN-43 presence and Radar 3 OOB. This collection of labeled data was partitioned into two disjoint sets: one for training and one for testing. In this section, after explaining how the two sets were selected, we present performance results for each classifier. Since ESC detection will only be required

for the 3550-3650 MHz portion of the 3.5 GHz band, the training and testing sets described below were limited to the eleven 10 MHz channels covering the range 3545-3655, with each channel centered on multiples of 10 MHz.

A. Test Set Composition

The sample of labeled spectrogram data was potentially biased in two respects. First, because the data collection was observational, at only two geographic locations for two months each, the respective proportions of different emission types did not necessarily reflect those in the whole population of possible field measurements, i.e., the distribution of 3.5 GHz emissions at all coastal locations under all conditions. Second, as mentioned in Section II, nearly 74% of the labeled spectrograms were selected for labeling because they corresponded to captures that triggered retention of a recorded waveform. Consequently, the set of labeled data suffered from a selection bias that resulted in a disproportionate number of labeled spectrograms with high-amplitude emissions.

Due to the above potential biases in the labeled data set, and due to the need for sufficient testing of important sub-groups, a stratified sampling approach was utilized to construct the test set. Specifically, a test set was randomly selected from the set of labeled data with approximately equal proportions of spectrograms across emission categories (SPN-43, Radar 3 OOB, Both SPN-43 and Radar 3 OOB, Neither), measurement locations (Virginia Beach and San Diego), and antenna type (Omni-directional and CBS). In addition, the maximum number of spectrograms containing multiple SPN-43 emissions were included. Table I shows the proportions for each category in the most general test set, denoted Test Set A. Note that the proportions are not exactly equal because the random test set generation program had to satisfy a hierarchy of preferences that did not typically lead to a perfect solution. Also, observe that roughly 50% of the cases contained SPN-43 and 50% of the cases did not contain SPN-43. To assess classifier performance without the presence of Radar 3 OOB, a subset of Test Set A, called Test Set B, was used; see Table I.

The data stratification used for the test set was selected to ensure that the full gamut of test cases were represented, including variations in measurements due to channel effects, receiver reference level, antenna type, measurement location, etc. Our aim was to carry out a rigorous evaluation of models by including an adequate representation of all cases that are likely to be observed in the field.

B. Training

As stated in Section III, the machine learning classifiers were first trained on 10 MHz channels for single-channel detection and then these already trained single-channel detection instances are connected in parallel for multichannel detection over the full spectrogram. The training set consisted of 10 MHz channels randomly extracted from the collection of labeled spectrograms not in Test Set A, which included cases collected at both measurement locations, with both antenna types, and all receiver reference levels. A total of 4,285 channels were randomly selected for training, where half contained SPN-43 and half did not.

All deep learning methods used uniform Xavier initializations [40] for convolutional layers; truncated mean initializations with means of zero, standard deviations of one, and truncation

at two standard deviations above and below the mean for fully connected layers; and zero initializations for bias layers. Both Adagrad [41] and Adam [42] optimizers were compared during training with no empirical advantage between the two. The learning rate was fixed at 0.0001 for all algorithms and the cross-entropy loss function was used.

To assess the sensitivity of CNN-3 and the LSTM training to weight initialization, we trained both algorithms 100 times with different, randomly-generated initializations. For this test, we used the Adagrad optimizer. We used cross-entropy for the loss function. Each training instance was run using the same training set for 1,000 epochs (an epoch consists of one pass through all training examples). Figure 3 summarizes the performance of CNN-3 and the LSTM on Test Set A over the 100 training initializations. In this figure, multichannel detection performance for each training initialization is summarized with an empirical free-response receiver operating characteristic (FROC) curve, shown in light gray. Appendix A reviews FROC curves, which can be used to summarize multichannel detection performance. The minimum and maximum bounds over all 100 FROC curves are shown in bold. Note that these bounds are not necessarily the same as any individual FROC curve.

From the plots in Figure 3, we can draw three conclusions. First, the CNN-3 distribution is much tighter than the LSTM distribution. Second, the best LSTM and CNN-3 instances performed similarly. Last, these plots emphasize the necessity of testing classifier performance over multiple training initializations before settling on a particular set of weights. The results presented below were generated using the CNN-3 and LSTM instance with the largest area under the FROC curve (FROC-AUC).

C. Performance Evaluation

Single-channel and multichannel detection performance was assessed using receiver operating characteristic (ROC) and FROC curves, respectively. See Appendix A for an introduction to ROC and FROC curves. In the single-channel ROC evaluations, the eleven 10 MHz channels covering 3550-3650 MHz in each spectrogram were tested, and the results were aggregated on a per-channel basis. On the other hand, the multichannel FROC evaluations holistically assessed detection performance over the entire 3550-3650 MHz frequency range by aggregating detection results on a per-spectrogram basis.

Single-channel performance of all thirteen algorithms is evaluated. For conciseness, however, the multichannel performance is only evaluated for the top-performing algorithm of each class for Test Set A. Single-channel detection performance is summarized by estimates of the area under the ROC curve (ROC-AUC), where a higher number indicates better performance. Only the top-performing SVM, KNN, and GMM models are shown in the single-channel analysis included here. Specifically, the top-performing models were the linear SVM₆₁₆₄, KNN₆₁₆₄ with $k=9$, and GMM₂₆₈.

Table II gives estimates of the ROC-AUC on Test Set A and B. Each AUC point estimate was estimated nonparametrically by computing the area under the empirical ROC curve. This is mathematically equivalent to the normalized Mann-Whitney U statistic [43]. Ninety-five percent confidence intervals were estimated using the nonparametric method of DeLong et al. [44] together with the logit transformation method recommended by Pepe [43, p. 107].

The top-performing algorithms for each algorithm category are in bold. Note that because the classifiers were applied to the same test set, the confidence intervals are correlated.

Table III gives estimates of the FROC-AUC on Test Set A and B. In this table, FROC-AUC was normalized to make it a number between 0 and 1; the normalization factors for test set A and B were 10.28 and 10.36, respectively. See Appendix A for a discussion of our rationale for FROC-AUC normalization. The AUC point estimates were estimated nonparametrically by computing the area under the empirical FROC curves. Ninety-five percent confidence intervals for FROC-AUC were estimated using the percentile bootstrap method [45], where the bootstrapping was stratified to maintain the proportions in Table I. Again, because the classifiers were applied to the same test sets, the confidence intervals are correlated.

From the ROC-AUC results, we see that each best-in-class machine learning method decidedly outperformed sweep-integrated energy detection. However, there was not a statistically significant difference between CNN-3, Inception-v1, and SVM₆₁₆₄. The FROC-AUC results yield a similar conclusion. Figures 4 and 6 show empirical ROC and FROC curves for CNN-3, Inception-v1, the linear SVM, and sweep-integrated energy detection (SI-ED) on Test Set A. These plots support the AUC conclusions, but provide additional insight into differences in performance across false positive rates.

The detection performance for Test Set B, a proper subset of Test Set A that excluded Radar 3 OOB, is summarized in Figure 5 and Figure 7 and with the respective columns in Table II and III. Comparing these results to those for Test Set A, we see that the removal of Radar 3 OOB only yielded a slight improvement in sweep-integrated energy detection for low false-positive rates. To elucidate this finding, Figure 8 shows estimated probability density functions (PDFs) for the power in each 660 kHz-wide channel used by the sweep-integrated energy detector for SPN-43-absent and SPN-43-present cases. Each PDF was estimated using the kernel density estimation method with a Gaussian kernel and a bandwidth of one. The peaks in the SPN-43-absent distributions correspond to the receiver noise floor, which varied with measurement location and receiver reference level [22], [23], [46]. From these plots, we see that the SPN-43-absent PDF for Test Set A has a fatter tail between -85 dBm and -70 dBm than for Test Set B, which is consistent with Radar 3 OOB being present in set A. However, there is only a very slight shift in the peaks of each distribution for sets A and B. The very small differences between the distributions for sets A and B help to explain the small improvement in sweep-integrated energy detection performance on Test Set B.

D. Speed Evaluation

Because ESC systems are required to detect SPN-43 within a small time window, information about the speed at which each model performed detection is relevant. To test detection speed for each algorithm, a single data sample was loaded into memory along with any model parameters. The model was timed while performing detection on this single sample 200,000 times. The average detection time for each model on the single sample is listed in Table IV. Note that only the top-performing algorithms from each category are listed.

In order to ensure the timings were measured fairly, all measurements were performed on an Nvidia® DGX™ workstation. All timings were performed solely on the CPU during times when the workstation was not otherwise in use. While hardware specifics may change the values of the times listed in Table IV, we believe the relative ordering will remain consistent across most hardware platforms.

As expected, the sweep-integrated energy detector was faster than the machine learning methods. CNN-3 was more than twice as fast as SVM₆₁₆₄ and 4.5 times quicker than Inception-v1. In the light of the results from the previous subsection, we can conclude that CNN-3 provides top-tier detection accuracy for significantly less computational cost than the other machine learning methods.

E. Detection Examples

To gain further insight into classifier performance, we examined spectrograms in which there was no consensus between machine learning and sweep-integrated energy detection. For this analysis, as a representative machine learning algorithm, we chose CNN-3. Both CNN-3 and sweep-integrated energy detection were applied with a decision threshold corresponding to a false-positive rate of 0.05 on Test Set A. Three notable examples of spectrograms in which sweep-integrated energy detection and CNN-3 differed, denoted Example 1, Example 2, and Example 3, respectively, are shown in Figure 9.

Example 1, shown in Figure 9 (Left), contains high-power Radar 3 OOB and no SPN-43 emissions. In this case, sweep-integrated energy detection incorrectly detected SPN-43 in every 10 MHz channel, i.e., every channel was a false-positive. On the other hand, CNN-3 had no false-positives.

Example 2, shown in Figure 9 (Middle), contains low-power Radar 3 OOB and no SPN-43 emissions. In this spectrogram, frequency-banding is evident, which may be due to multi-path fading. For this example, energy detection correctly labeled all channels as negative for SPN-43. However, CNN-3 incorrectly detected SPN-43 in several channels, resulting in false-positives. This example illustrates a potential weakness of CNN-3 that could be explored in future work.

Last, Example 3, shown in Figure 9 (Right), contains an evident SPN-43 emission at 3630 MHz and a faint SPN-43 emission at 3600 MHz. The vertical line at 3577 MHz is local oscillator leakage, which appears bright due to the tight grayscale window. In this example, sweep-integrated energy detection failed to detect both SPN-43 radars, which were correctly detected by CNN-3.

V. Applications

Distributions estimated from field measurements for SPN-43 spectrum occupancy and SPN-43-absent power density are potentially informative to both federal regulators and commercial industry, as they may be relevant to ESC requirements [47], [48] and ESC development efforts. Namely, occupancy distributions may be relevant to a requirement that the channel be vacated for a fixed time-interval after incumbent signals have been detected

[47]. Distributions of SPN-43-absent power may be relevant to ESC developers since some detection strategies may result in unacceptably high false-alarm rates for channels with higher levels of non-SPN-43 emissions. In addition, field observations of ambient power levels are relevant to ESC certification testing [48], since they could inform selection of background noise levels.

As explained in Section II, a total of 14,739 spectrograms were collected in San Diego and Virginia Beach, of which 4,491 were human-labeled for SPN-43 presence. In this section, we describe how one of the best-performing classifiers from Section IV-C, CNN-3, was applied to classify the unlabeled spectrograms for SPN-43 presence and how the complete set of spectrograms was then used to estimate distributions for SPN-43 spectrum occupancy and for power density when SPN-43 was absent. To classify unlabeled spectrograms for SPN-43 presence, we chose a decision threshold to generate the CNN-3 prediction output from each 10 MHz channel. Below, we give details on the selected decision threshold, which was different for each application to accommodate dissimilar preferences between true-positive and false-positive rates.

The findings given here are a partial selection from a larger set of descriptive statistics that is provided in an accompanying technical report [46]. It should be emphasized that because these results are derived from spectrum observations at only two geographic locations for two months each, the reader should be careful not to draw overly-general conclusions.

A. Channel Occupancy Statistics

For the goal of estimating SPN-43 channel occupancy, we chose the CNN-3 decision threshold to control the false-positive rate. Specifically, the decision threshold was selected to correspond to a false-positive rate of 0.01 on Test Set A; this operating point corresponds to a true-positive rate of 0.97. Note that false-positives lead to a positive bias in occupancy estimates and a negative bias in vacancy estimates. Thus, because we controlled the false-positive rate, our occupancy and vacancy estimates are conservative and liberal, respectively.

As stated in Section II, spectrograms were collected roughly every ten minutes. This sampling interval was not exact due to hardware restrictions, like the rate at which data was saved to disk. Despite this fact, to simplify our estimates of vacancy and occupancy time-intervals, we assumed that the captures were exactly ten minutes apart. To calculate the length of time a 10 MHz channel was either occupied by SPN-43 or vacant, we ordered the spectrograms by their capture time-stamp and then counted the number of consecutive vacant and occupied observations. The counts were then multiplied by 10 minutes to estimate durations. Note that this approach could not resolve changes in SPN-43 occupancy that occurred less than 10 minutes apart.

Figure 10 shows histograms of occupied and vacant time-intervals measured in minutes for the 10 MHz channel centered at 3550 MHz in San Diego. Specifically, the occupancy histogram lists the number of time-intervals for which SPN-43 was continuously-observed for the specified duration. For example, the 3550 MHz channel was continuously occupied for 30-40 minutes nine times during the two-month measurement period in San Diego. Similarly, the vacancy histogram lists the number of time-intervals for which SPN-43 was

not present for the specified duration, e.g., there were ten SPN-43 vacancies with durations of 50-60 minutes. Only time-intervals below 120 minutes are shown. Of the observed time-intervals, 24 occupancies and 138 vacancies exceeded 120 minutes.

To gain a better understanding of how often a channel was occupied, we estimated the occupancy ratio, i.e., the amount of time the channel was occupied by SPN-43 divided by the total observation time. Table V lists the estimated occupancy ratio for channels where SPN-43 was observed in San Diego and Virginia Beach, respectively.

B. Power Density Distributions

For the aim of estimating the distribution of power density when SPN-43 was absent, we chose the CNN-3 decision threshold to control the number of missed SPN-43 detections (false-negatives). Specifically, the decision threshold was selected to correspond to a true-positive rate of 0.98 on Test Set A; this operating point corresponds to a false-positive rate of 0.02. Because missed detections lead to the inclusion of SPN-43 emissions in estimates of the SPN-43-absent power density, and because the power density of SPN-43 emissions is typically quite high, missed detections are expected to add a positive bias. Therefore, to avoid such a bias, we controlled the rate of missed detections with the potential expense of additional false-positives, which shrank our sample size for SPN-43-absent observations.

After classification of the unlabeled spectrograms, the channels found to contain SPN-43 were discarded, and the empirical cumulative distribution function (CDF) for the power density was estimated from the set of spectrogram values (converted to dBm/MHz) in the 220 kHz-wide frequency-bin at the center of each 10 MHz channel (the expected location for SPN-43). Figure 11 shows examples of the empirical complementary CDF (CCDF), equal to one minus the CDF, for the SPN-43-absent power density for frequencies where SPN-43 was observed in San Diego and Virginia Beach. These plots can be used to quickly read off percentiles associated with the upper tails of each distribution. Namely, the 90th and 99th percentiles correspond to the power density values where the CCDF is equal to 0.1 and 0.01, respectively. Nonparametric simultaneous 95% confidence bands were estimated for the empirical CCDFs using a method based on the Dvoretzky-Kiefer-Wolfowitz inequality [49, Thm. 7.5]. The steep decrease in the CCDFs corresponds to the noise floor of the receiver, which varied with measurement location and receiver reference level [22], [23], [46]. The tail heaviness indicates the prevalence of non-SPN-43 emissions, such as Radar 3 OOB.

VI. Summary and Discussion

Accurate detection of radar in the 3.5 GHz band with ESC sensors is needed to both protect federal incumbent systems and to enable economical commercial utilization of the band. In this paper, for the task of single-channel SPN-43 radar detection from spectrograms collected with a single receiver, we investigated the effectiveness of thirteen detection algorithms, including eight deep learning methods, three classical machine learning approaches, and two energy detection strategies. Furthermore, for the task of wideband SPN-43 detection across multiple channels observed concurrently with one receiver, we compare the top-performing methods from the single-channel evaluation. The detection

algorithms were trained and tested with a set of nearly 4,500 unverified, human-labeled spectrograms collected at two coastal locations.

For ESC networks to be effective, they will require a large numbers of detectors. Detection algorithms that are fast and computationally inexpensive are ideal for such networks because they reduce the costs for the large number of individual sensors. For this reason, we chose to explore the use of spectrograms for detection, as both generating and processing spectrograms requires relatively inexpensive hardware. Algorithms for detecting emissions from other data representations such as I/Q recordings may result in higher detection accuracy but also require more expensive hardware to generate and process. Nonetheless, the use of I/Q recordings is an area for potential future work.

Our evaluations provide estimates for the level of performance that can potentially be expected from an ESC detector across a range of real-world conditions. In particular, our test set included cases with realistic channel conditions and Radar 3 OOB. Despite these advantages, there are several drawbacks to using real data. First, using real-world data with unknown signal and noise components prevents accurate assessment of detection performance as a function of SNR. This type of analysis can provide information about emission strength that must be present for detectors to succeed. In future work, we plan to evaluate detection strategies using synthetic data generated from our field measurements in which the SNR can be controlled.

Second, capturing a desired ratio of observations for different subgroups is difficult in the field. Consequently, our dataset had a limited number of cases available to construct training, validation, and test sets with sufficient representation of the subgroups listed in Table I. To ensure that our training and testing sets contained enough cases for each relevant subgroup, we chose to not use a separate validation set [33, Sec. 1.4.8]. Instead, in an attempt to avoid over-fitting, we stopped training the deep learning models after a fixed number of epochs.

Our evaluations on real-world data demonstrated that machine learning methods offer superior detection performance compared to methods based on energy detection. The excellent detection performance of CNN-3 allowed us to estimate spectrum occupancy statistics and power distributions for non-SPN-43 emissions with a much higher accuracy than would have been otherwise possible. In particular, using energy detection to classify the unlabeled data would have resulted in many more false-positives and missed detections at a given decision threshold, which would lead to biased estimates. Namely, false-positives lead to overestimation of spectrum occupancy, whereas missed detections result in an overestimation of spectrum vacancy and add a positive bias to non-SPN-43 power distributions. As explained in Section V, occupancy statistics and power distributions may have value to both ESC developers and spectrum regulators. A complete set of occupancy statistics and power distributions for each channel in the 3.5 GHz band is provided in a technical report [46].

Appendix A

ROC and FROC Curves

We present a brief introduction to two types of graphical plots that can be used to evaluation detection performance: the receiver operating characteristic (ROC) curve and a related generalization, called the free-response ROC (FROC) curve. Further background on ROC curves can be found in [43], [50], [51] and details on FROC curves are given in [52], [53]. Although ROC curves are commonly utilized in signal processing and machine learning, FROC curves are lesser-known, since they have been primarily applied in radiology to evaluate lesion detection performance. In the context of multichannel spectrum sensing for cognitive radio, a notable application of FROC curves to classifier performance evaluation is the work of Collins and Sirkeci-Mergen [54].

A. Binary Signal Detection: The ROC Curve

For a binary signal detection task, the aim is to use a data observation to decide whether or not a signal is present. Each decision results in one of four possible outcomes: true-positive (TP), false-positive (FP), true-negative (TN), or false-negative (FN). These outcomes give rise to four conditional probabilities (or rates). In the engineering literature, the TP rate, FN rate, and FP rate are commonly called the “detection”, “miss” and “false-alarm” probabilities, respectively. For a given decision threshold, binary classification performance is fully described by the FP and TP rates. Namely, the FN rate is equal to one minus the TP rate, and the TN rate is one minus the FP rate.

A useful way to summarize detection performance is the ROC curve, defined as the plot of TP rate versus FP rate, over all decision thresholds [43], [50]. An example of an ROC curve is shown in Figure 12 (left). When comparing ROC curves, better classifier performance is indicated by a higher curve that is closer to the upper left corner. Namely, for a perfect classifier, there exists a threshold where the TP rate is equal to one with a FP rate of zero. By contrast, for a useless classifier, the ROC curve is equal to or below the diagonal dashed “chance line” shown in Figure 12 (left) for which the TP rate is equal to the FP rate at all decision thresholds [43].

ROC curves possess three properties that make them particularly useful. First, they fully characterize binary classifier performance over all decision thresholds, which enables evaluation and comparison of classifiers that may be deployed at various operating points (thresholds) [51]. Second, ROC curves are invariant under strictly-increasing transformations of the decision variable [43]. Thus, classifiers with decision variables on different ordinal scales can be compared via ROC curves. Third, ROC curves are independent of signal prevalence [51]. This implies that ROC curves can be used to assess classifiers that may be deployed in environments with different signal prevalences.

A commonly-used summary measure for binary classification performance is the area under the ROC curve (ROC-AUC). ROC-AUC takes values between zero and one, with higher values indicating better performance. ROC-AUC can be interpreted as the average TP rate, averaged uniformly over all FP rates. Alternatively, ROC-AUC can be interpreted as a

probability. Namely, given randomly-selected signal-absent and signal-present cases, ROC-AUC is the probability that the signal-present case is rated higher [43, Sec. 4.3].

In this paper, we use the so-called “empirical” nonparametric estimators for the ROC curve and its area. For details on these estimators, see [43, Sec. 5.2]. In addition to having simple implementations, the empirical estimators are nonparametric and unbiased.

B. Multiple Signal Detection and Localization: The FROC Curve

The FROC curve [52], [53] is generalization of the ROC curve designed to summarize classifier performance for a combined detection and localization task in which multiple detection decisions are made for each observation. An example of such a task arises in multichannel spectrum sensing, where the aim is to detect one or more signals and localize them in frequency. After specifying a criterion for correct signal localization, it is possible to determine if a detection result is a correctly-localized TP or a FP. Suppose that each correctly-localized TP detection occurs with the same probability, called the signal-detection fraction. The FROC curve is defined as the plot of signal detection fraction versus the mean number of false-positives per observation, plotted over all decision thresholds; an example FROC curve is shown in Figure 12. Like the ROC curve, an FROC curve closer to the upper left corner of the graph indicates better classifier performance. To estimate the FROC curve, we use the usual empirical estimator [52].

When the number of detection decisions is bounded, the abscissa (x-axis) of the FROC curve is bounded, and the area under the FROC curve (FROC-AUC) is a well-defined summary measure. For example, multichannel spectrum sensing typically aims to assess spectrum occupancy for a fixed number of frequency channels. Because the maximum abscissa value for an empirical FROC curve depends on the maximum number of possible FP decisions in the test set, the maximum empirical FROC-AUC can be different for dissimilar test sets. Therefore, in this paper, to enable straightforward comparisons between test sets, we normalize FROC-AUC values to fall between zero and one. The normalization factor depends on the maximum number of possible FP decisions in the test set.

C. When to Use Which Curve?

Because ROC and FROC curves are designed for different, but related tasks, they provide complementary insights into classifier performance. In particular, ROC curves focus solely on signal detection for a single decision, regardless of signal localization. Thus, for the problem of multichannel spectrum sensing, ROC curves are best suited to low-level assessment of classifier performance for a single channel. Such evaluations may be particularly useful for classifier development. On the other hand, FROC curves assess both detection and signal localization when multiple decisions must be made. For this reason, they are better suited to classifier performance evaluation for the full multichannel spectrum sensing task. If FROC curves are not a good match to the task and associated preferences, one can consider variations of FROC curves that weight TP and FP decisions differently; for further details on FROC variants and their generalizations, see [55].

References

- [1]. Yucek T and Arslan H, "A survey of spectrum sensing algorithms for cognitive radio applications," *IEEE Commun. Surveys Tuts*, vol. 11, no. 1, pp. 116–130, 2009.
- [2]. Axell E, Leus G, Larsson EG, and Poor HV, "Spectrum sensing for cognitive radio: State-of-the-art and recent advances," *IEEE Signal Process. Mag.*, vol. 29, no. 3, pp. 101–116, 2012.
- [3]. Bkassiny M, Li Y, and Jayaweera SK, "A survey on machine-learning techniques in cognitive radios," *IEEE Commun. Surveys Tuts*, vol. 15, no. 3, pp. 1136–1159, 2013.
- [4]. Thilina KM, Choi KW, Saquib N, and Hossain E, "Machine learning techniques for cooperative spectrum sensing in cognitive radio networks," *IEEE J. Sel. Areas Commun.*, vol. 31, no. 11, pp. 2209–2221, 2013.
- [5]. Zhang K, Li J, and Gao F, "Machine learning techniques for spectrum sensing when primary user has multiple transmit powers," in *Proc. IEEE Int. Conf. Comm. Sys.*, pp. 137–141, 2014.
- [6]. Lu Y, Zhu P, Wang D, and Fattouche M, "Machine learning techniques with probability vector for cooperative spectrum sensing in cognitive radio networks," in *Proc. IEEE Wireless Comm. and Netw. Conf.*, pp. 16, 2016.
- [7]. Jiang C, Zhang H, Ren Y, Han Z, Chen K-C, and Hanzo L, "Machine learning paradigms for next-generation wireless networks," *IEEE Wireless Communications*, vol. 24, no. 2, pp. 98–105, 2017.
- [8]. Quan Z, Cui S, Poor HV, and Sayed AH, "Collaborative wideband sensing for cognitive radios," *IEEE Signal Process. Mag.*, vol. 25, no. 6, 2008.
- [9]. Quan Z, Cui S, Sayed AH, and Poor HV, "Optimal multiband joint detection for spectrum sensing in cognitive radio networks," *IEEE Trans. Signal Process.*, vol. 57, no. 3, pp. 1128–1140, 2009.
- [10]. Farrag M, Muta O, El-Khamy M, Furukawa H, and El-Sharkawy M, "Wide-band cooperative compressive spectrum sensing for cognitive radio systems using distributed sensing matrix," in *Proc. IEEE Vehicular Tech. Conf.*, pp. 1–6, 2014.
- [11]. Goodfellow I, Bengio Y, and Courville A, *Deep Learning*. MIT press, 2016.
- [12]. LeCun Y, Bengio Y, and Hinton G, "Deep learning," *Nature*, vol. 521, pp. 436–444, 5 2015. [PubMed: 26017442]
- [13]. O'Shea T and Hoydis J, "An introduction to deep learning for the physical layer," *IEEE Trans. Cogn. Commun. Netw.*, vol. 3, no. 4, pp. 563–575, 2017.
- [14]. O'Shea TJ, Roy T, and Clancy TC, "Over-the-air deep learning based radio signal classification," *IEEE J. Sel. Topics Signal Process.*, vol. 12, no. 1, pp. 168–179, 2018.
- [15]. Rajendran S, Meert W, Giustiniano D, Lenders V, and Pollin S, "Deep learning models for wireless signal classification with distributed low-cost spectrum sensors," *IEEE Trans. Cogn. Commun. Netw.*, vol. 4, pp. 433–445, 9 2018.
- [16]. Zhang C, Patras P, and Haddadi H, "Deep learning in mobile and wireless networking: A survey," preprint arXiv:1803.04311, 2018.
- [17]. "Citizens Broadband Radio Service." Code of Federal Regulations. Title 47, Part 96, 6 2015.
- [18]. "An assessment of the near-term viability of accommodating wireless broadband systems in the 1675–1710 MHz, 1755–1780 MHz, 3500–3650 MHz, 4200–4220 MHz and 4380–4400 MHz bands." National Telecommunications and Information Administration, 10 2010.
- [19]. "Operation and maintenance instructions, organizational level, radar set AN/SPN-43C." Naval Air Systems Command, Technical Manual, EE216-EB-OMI-010, vol. 1, 9 2005.
- [20]. Cotton MG and Dalke RA, "Spectrum occupancy measurements of the 3550–3650 megahertz maritime radar band near San Diego, California," *Tech. Rep. TR 14-500*, National Telecommunications and Information Administration, 1 2014.
- [21]. Sanders FH, Carroll JE, Sanders GA, and Cohen LS, "Measurements of selected naval radar emissions for electromagnetic compatibility analyses," *Tech. Rep. TR 15-510*, National Telecommunications and Information Administration, 10 2014.
- [22]. Hale P, Jargon J, Jeavons P, Lofquist M, Souryal M, and Wunderlich A, "3.5 GHz radar waveform capture at Point Loma," *Tech. Note 1954*, National Institute of Standards and Technology, 5 2017.

- [23]. Hale P, Jargon J, Jeavons P, Lofquist M, Souryal M, and Wunderlich A, "3.5 GHz radar waveform capture at Fort Story," Tech. Note 1967, National Institute of Standards and Technology, 10 2017.
- [24]. Simonyan K and Zisserman A, "Very deep convolutional networks for large-scale image recognition," preprint arXiv:1409.1556, 2014.
- [25]. He K, Zhang X, Ren S, and Sun J, "Deep residual learning for image recognition," in Proc. IEEE Conf. Computer Vision and Pattern Recognition, pp. 770–778, 2016.
- [26]. Szegedy C, Liu W, Jia Y, Sermanet P, Reed S, Anguelov D, Erhan D, Vanhoucke V, and Rabinovich A, "Going deeper with convolutions," in Proc. IEEE Conf. Com. Vision Pattern Recog, pp. 1–9, 2015.
- [27]. Huang G, Liu Z, Van Der Maaten L, and Weinberger KQ, "Densely connected convolutional networks," in Proc. IEEE Conf. Computer Vision and Pattern Recognition, vol. 1, p. 3, 2017.
- [28]. Hochreiter S and Schmidhuber J, "Long short-term memory," Neural Comput, vol. 9, no. 8, pp. 1735–1780, 1997. [PubMed: 9377276]
- [29]. Urkowitz H, "Energy detection of unknown deterministic signals," Proc. IEEE, vol. 55, no. 4, pp. 523–531, 1967.
- [30]. Abdulsattar MA and Hussein ZA, "Energy detection technique for spectrum sensing in cognitive radio: a survey," Int. J. Computer Networks & Communications, vol. 4, pp. 223–224, 9 2012.
- [31]. Atapattu S, Tellambura C, and Jiang H, Energy Detection for Spectrum Sensing in Cognitive Radio. Springer, 2014.
- [32]. Richards MA, Fundamentals of Radar Signal Processing. McGraw-Hill, 2005.
- [33]. Murphy KP, Machine Learning: A Probabilistic Perspective. MIT press, 2012.
- [34]. Zaki MJ, Meira W Jr, and Meira W, Data Mining and Analysis: Fundamental Concepts and Algorithms. Cambridge University Press, 2014.
- [35]. Deng J, Dong W, Socher R, Li L-J, Li K, and Fei-Fei L, "ImageNet: A Large-Scale Hierarchical Image Database," in Proc. IEEE Conf. Computer Vision and Pattern Recognition, pp. 248–255, 2009.
- [36]. Liu L, Ouyang W, Wang X, Fieguth P, Chen J, Liu X, and Pietikäinen M, "Deep learning for generic object detection: A survey," preprint arXiv:1809.02165, 2018.
- [37]. Kim J, El-Khamy M, and Lee J, "Residual LSTM: Design of a deep recurrent architecture for distant speech recognition," arXiv preprint, 2017 <https://arxiv.org/abs/1701.03360>.
- [38]. Lin M, Chen Q, and Yan S, "Network in network," preprint arXiv:1312.4400, 2013.
- [39]. Huang Y, Sun X, Lu M, and Xu M, "Channel-max, channel-drop and stochastic max-pooling," in Proc. IEEE Conf. Computer Vision and Pattern Recognition Workshops, pp. 9–17, IEEE, 2015.
- [40]. Glorot X and Bengio Y, "Understanding the difficulty of training deep feedforward neural networks," in Proc. 13th Int. Conf. Artificial Intelligence and Statistics, pp. 249–256, 2010.
- [41]. Duchi J, Hazan E, and Singer Y, "Adaptive subgradient methods for online learning and stochastic optimization," Journal of Machine Learning Research, vol. 12, no. Jul, pp. 2121–2159, 2011.
- [42]. Kingma DP and Ba J, "Adam: A method for stochastic optimization," arXiv preprint arXiv:1412.6980, 2014.
- [43]. Pepe MS, The Statistical Evaluation of Medical Tests for Classification and Prediction. Oxford Univ. Press, 2003.
- [44]. DeLong ER, DeLong DM, and Clarke-Pearson DL, "Comparing the areas under two or more correlated receiver operating characteristic curves: A nonparametric approach," Biometrics, vol. 44, pp. 837–845, 9 1988. [PubMed: 3203132]
- [45]. Efron B and Tibshirani RJ, An Introduction to the Bootstrap. Boca Raton, FL: Chapman & Hall/CRC, 1993.
- [46]. Lees WM, Wunderlich A, Jeavons P, Hale PD, and Souryal MR, "Spectrum occupancy and ambient power distributions for the 3.5 GHz band estimated from observations at Point Loma and Fort Story," Tech. Note 2016, National Institute of Standards and Technology, 9 2018.

- [47]. Wireless Innovation Forum, Requirements for Commercial Operation in the U.S. 3550-3700 MHz Citizens Broadband Radio Service Band, Working Document WINNF-TS-0112, Version V1.5.0, 5 2018.
- [48]. Sanders FH, Carroll JE, Sanders GA, Sole RL, Devereux JS, and Drocella EF, "Procedures for laboratory testing of environmental sensing capability sensor devices," Tech. Rep. TM 18-527, National Telecommunications and Information Administration, 11 2017.
- [49]. Wasserman L, All of Statistics: A Concise Course in Statistical Inference. New York: Springer, 2004.
- [50]. Van Trees HL, Detection, Estimation, and Modulation Theory, Part I. New York: John Wiley & Sons, 1968.
- [51]. Metz CE, "Basic principles of ROC analysis," Semin. Nucl. Med, vol. 8, pp. 283–298, 1978. [PubMed: 112681]
- [52]. Bunch P, Hamilton J, Sanderson G, and Simmons A, "Free-response approach to the measurement and characterization of radiographic observer performance," J. Appl. Photogr. Eng, vol. 4, no. 4, pp. 166–171, 1978.
- [53]. Metz CE, "Receiver operating characteristic analysis: A tool for the quantitative evaluation of observer performance and imaging systems," J. Am. Coll. Radiol, vol. 3, pp. 413–422, 2006. [PubMed: 17412096]
- [54]. Collins SD and Sirkeci-Mergen B, "Localization ROC analysis for multiband spectrum sensing in cognitive radio," in Proc. IEEE Mil. Comm. Conf, pp. 64–67, 2013.
- [55]. Wunderlich A, Goossens B, and Abbey CK, "Optimal joint detection and estimation that maximizes ROC-type curves," IEEE Trans. Med. Imag, vol. 35, pp. 2164–2173, 9 2016.

Fig. 1:

Example spectrogram captures, cropped to 3545-3655 MHz. (Left) Strong SPN-43 emissions near 3570 MHz; grayscale window [-90 -50] dBm (Middle) Radar 3 OOB coincident with SPN-43 emissions near 3570 MHz; grayscale window [-95 -75] dBm (Right) Radar 3 OOB coincident with weak SPN-43 emissions near 3550 MHz; grayscale window [-95 -75] dBm.

Fig. 2:
Flowcharts for the deep learning implementations. (Left) The CNN-3 architecture. (Right) The LSTM architecture.

Fig. 3: FROC curves summarizing multichannel detection performance on Test Set A over 100 different training initializations for CNN-3 (Left) and LSTM (Right). The minimum and maximum bounds over all initializations are shown in bold.

Fig. 4:
Test Set A ROC results. Left: Full ROC curves for single channel detection. Right: Y-axis zoom of plot on left.

Fig. 5: Test Set B ROC results. Left: Full ROC curves for single channel detection. Right: Y-axis zoom of plot on left.

Fig. 6:
Test Set A FROC results. Left: Full FROC curves for multichannel detection. Right: Y-axis zoom of plot on left.

Fig. 7:
Test Set B FROC results. Left: Full FROC curves for multichannel detection. Right: Y-axis zoom of plot on left.

Fig. 8:
Estimated probability density functions for power in the frequency bins used by energy detector. Left: Test Set A. Right: Test Set B.

Fig. 9:

Example spectrogram captures, cropped to 3545-3655 MHz. (Top) Example 1: high-power Radar 3 OOB; grayscale window $[-95 -20]$ dBm. (Middle) Example 2: low-power Radar 3 OOB; grayscale window $[-95 -70]$ dBm. (Bottom) Example 3: SPN-43 emissions near 3630 MHz and a weak, barely visible SPN-43 at 3600 MHz; grayscale window $[-100 -85]$ dBm.

Fig. 10:

Histograms of the empirical length of time the 10 MHz channel centered at 3550 MHz was occupied (left) or unoccupied (right) by SPN-43 in San Diego for time intervals less than 120 minutes.

Fig. 11:

The power CCDF for emissions captured using the CBS antenna when SPN-43 was absent in San Diego (Left) and Virginia Beach (Right). Simultaneous 95% confidence bands are indicated by dashed lines.

Fig. 12: Examples of an ROC curve (Left) and an FROC curve (Right). In the ROC plot, the “chance line” is depicted with the diagonal dashed line.

TABLE I:

Test Set Compositions. Above, “R3-OOBE” stands for Radar 3 OOBE. Test Set B is a proper subset of Test Set A that does not contain cases with Radar 3 OOBE.

Set	Total	Multi SPN-43	Emissions				Location		Antenna	
			SPN-43	R3-OOBE	Both	Neither	VB	SD	Omni	CBS
A	509	109	24.56%	24.36%	26.72%	24.36%	51.28%	48.72%	48.92%	51.08%
B	249	40	50.20%	0%	0%	49.80%	50.20%	49.80%	50.20%	49.80%

TABLE II:

ROC-AUC estimates for Test Set A and B. Each entry shows a point estimate and a 95% confidence interval.

Classifier	ROC-AUC (Set A)	ROC-AUC (Set B)
CNN-3	.997, [.993, .998]	.997, [.992, .999]
LSTM	.988, [.978, .994]	.994, [.987, .998]
Inception-v1	.994, [.985, .997]	.998, [.993, .999]
ResNet-18	.913, [.889, .932]	.891, [.856, .918]
DenseNet-121	.838, [.807, .866]	.782, [.727, .829]
ResNet-50	.622, [.576, .665]	.623, [.558, .684]
VGG-16	.500, [.500, .500]	.500, [.500, .500]
VGG-19	.500, [.500, .500]	.500, [.500, .500]
SVM ₆₁₆₄	.991, [.985, .994]	.996, [.989, .998]
KNN ₆₁₆₄	.965, [.949, .976]	.981, [.959, .991]
GMM ₂₆₈	.868, [.845, .888]	.884, [.847, .913]
SI-ED	.884, [.854, .910]	.899, [.850, .933]
ED	.854, [.819, .883]	.852, [.796, .895]

TABLE III:

Normalized FROC-AUC estimates for Test Set A and B. Each entry shows a point estimate and a 95% confidence interval.

Classifier	FROC-AUC (Set A)	FROC-AUC (Set B)
CNN-3	.997, [.994, .998]	.993, [.984, .997]
Inception-v1	.990, [.980, .998]	.997, [.993, 1.00]
SVM ₆₁₆₄	.988, [.980, .993]	.994, [.988, .999]
SI-ED	.867, [.834, .899]	.881, [.835, .926]

TABLE IV:

The average detection time for each model on a single sample.

Classifier	Time (ms)
SI-ED	0.40
CNN-3	1.89
SVM ₆₁₆₄	4.43
Inception-v1	8.45

NIST Author Manuscript

NIST Author Manuscript

NIST Author Manuscript

TABLE V:

Estimated occupancy ratios for 10 MHz channels in which SPN-43 was observed. The error bands indicate 95% confidence intervals. Note that SPN-43 was observed at 3520 MHz in San Diego [22], which is outside of the CBRS band.

Measurement Site	Channel Center (MHz)	Occupancy Ratio
San Diego	3520	0.036 ± 0.002
San Diego	3550	0.287 ± 0.005
San Diego	3600	0.018 ± 0.001
Virginia Beach	3570	0.116 ± 0.004
Virginia Beach	3600	0.118 ± 0.004
Virginia Beach	3630	0.041 ± 0.003