

# AirVIEW: Unsupervised transmitter detection for next generation spectrum sensing

Mariya Zheleva<sup>1</sup>, Petko Bogdanov<sup>1</sup>, Timothy Larock<sup>1</sup>, and Paul Schmitt<sup>2</sup>

<sup>1</sup>Department of Computer Science, University at Albany SUNY, {mzheleva, pbogdanov, tlarock}@albany.edu

<sup>2</sup>Center for Information Technology Policy, Princeton University, pschmitt@cs.princeton.edu

**Abstract**—The current paradigm of exclusive spectrum assignment and allocation is creating artificial spectrum scarcity that has a dramatic impact on network performance and user experience. Thus, governments, industry and academia have endeavored to create novel spectrum management mechanisms that allow multi-tiered access. A key component of such an approach is deep understanding of spectrum utilization in time, frequency and space. To address this challenge, we propose AirVIEW, a one-pass, unsupervised spectrum characterization approach for rapid transmitter detection with high tolerance to noise. AirVIEW autonomously learns its parameters and employs wavelet decomposition in order to amplify and reliably detect transmissions at a given time instant. We show that AirVIEW can robustly identify transmitters even when their power is only 5dBm above the noise floor. Furthermore, we demonstrate AirVIEW’s ability to inform next-generation Dynamic Spectrum Access by characterizing essential transmitter properties in wide-band spectrum measurements from 50MHz to 4.4GHz.

## I. INTRODUCTION

Current spectrum management is based on exclusive allocation and assignment of radio spectrum to a given technology and operator. As a result some popular frequency bands, such as cellular, become over-saturated, while others, like UHF TV, remain underutilized. This results in artificial spectrum scarcity that leads to high cost for communication services, decreased network performance and deteriorated user experience. At the same time, underutilized bands provide an opportunity for more efficient, shared spectrum access that has brought together policymakers, industry and academia to set an agenda for next-generation spectrum management [2].

A critical enabler of future spectrum access is deep understanding of spectrum utilization, both long-term as well as instantaneous. This goal entails *adaptive sensing* in a wide frequency range (i.e. 30MHz-6GHz [12]) followed by *autonomous characterization* to extract actionable knowledge from spectrum data and inform shared access. Our focus is on the latter: unsupervised characterization of wide-band spectrum scans that can pin-point the number of transmitters and their time-frequency properties without prior knowledge of transmitters’ behavior. Depending on the learning objectives, spectrum characterization can be performed in the time [3] or frequency domain [22] and should be regarded as a stream processing task, whereby data is analyzed as it arrives. Wide-band spectrum measurement assumes frequency-domain analysis and can be regarded at three resolutions as illustrated in Fig. 1(left): (i) at a single spectrum sweep, (ii) at a single




Fig. 1. (left) Illustration of a transmitter comprised of multiple transmissions; (right) Transmitter characterization metrics.

transmission, (i.e. an active region over continuous sweeps), and (iii) at a transmitter, defined as multiple transmissions reconciled into a single transmitter. Fig. 1(right) summarizes key metrics that must be produced by spectrum analysis.

Current approaches to spectrum characterization perform detection of idle and occupied bands [20, 23, 18], however, they do not cater to deeper functional spectrum analysis that teases out transmitters and their temporal and frequency usage patterns. Recent work has focused on cyclo-stationary analysis [17, 9] or signature matching [15]. While such approaches provide deeper understanding of spectrum occupancy, they are limited to predefined transmitter types and are, thus, not well-suited for arbitrary transmitter characterization. Other work [25] tackles unsupervised transmitter detection, however the approach requires batch-processing of data and, thus, cannot be employed for stream-processing of spectrum scans.

To address these limitations we design AirVIEW, an algorithm for unsupervised, wide-band, high-sensitivity spectrum analysis. AirVIEW operates at the granularity of a single sweep, where it employs wavelet decomposition [14] of power spectral density (PSD) to separate transmitter edges from the noise floor. Single-sweep frequency bands identified as occupied, are further reconciled through a two-step procedure that pinpoints longitudinally-active transmitters. Our wavelet decomposition approach provides an alternative expression of the measured PSD that presents a powerful construct for signal analysis [11], as it (i) can be used to effectively amplify the signal transitions while (ii) suppressing the inherent noise in spectrum measurements. While the benefits of wavelet decomposition in spectrum sensing have been observed before [16, 18, 24], all of these works operate on a single spectrum sweep and are thus not able to reconcile transmitters. Furthermore, none of them presents a fully-unsupervised approach that is able to adaptively tune its parameters to the observed spectrum dynamics and noise floor. Thus, the key advantages of AirVIEW over existing methods are that it is (i)

*robust* to noise, (ii) *unsupervised*, and thus, able to characterize arbitrary transmitters and (iii) *rapid* as it can extract actionable spectrum information even from a single sweep.

We evaluate AirVIEW’s performance on both synthetically generated and real-world ground truth transmissions, and show that it can robustly detect the temporal and frequency characteristics of transmitters even when their power is as low as 5dBm above the noise floor. We carry out a wide-band measurement campaign (50MHz-4.4GHz) in an urban location in close proximity to an airport and a military airbase. We then employ AirVIEW to analyze collected spectrum scans and demonstrate its utility for informing next generation Dynamic Spectrum Access (DSA) technology and spectrum policy through rapid characterization of (i) idle and occupied frequency bands, (ii) number of incumbents in each band, (iii) their temporal characteristics and (iv) the predictability of incumbents’ behavior. This paper makes several contributions:

- **Novelty:** We design AirVIEW, a high-sensitivity, unsupervised transmitter characterization algorithm that is robust to noise and is able to detect arbitrary transmitters.
- **Scalability:** Since it requires a single pass over the data, AirVIEW scales to wide spectrum scans at high temporal resolution, as demonstrated in our real-world evaluation.
- **Accuracy:** Our empirical evaluation demonstrates the accuracy of AirVIEW on synthetic and real spectrum traces.
- **Applicability:** We demonstrate AirVIEW’s applicability for unsupervised characterization of arbitrary transmitters on real-world, wide-band spectrum scans (50MHz-4.4GHz).

## II. RELATED WORK

Prior work on spectrum analysis can be largely subdivided in *activity detection* and *detailed transmitter characterization*. Activity detection is performed on a single sweep to determine which bands are active and which are idle. Such methods, however, do not provide further reconciliation of transmitter activity in multi-sweep spectrum measurements. The latter has been recently identified as an important advantage and tacked in several works for detailed transmitter characterization that utilize either supervised or unsupervised techniques.

**Energy-based activity detection.** Spectrum characterization has been actively explored in the past, however, the literature is limited in methods that provide detailed, robust and unsupervised transmitter characterization. Traditional approaches to spectrum characterization identify idle and occupied bands by the use of power thresholding [21, 13, 6], edge detection [18, 20] and compressive sensing [19, 4, 10]. While these approaches are computationally light-weight, they are only able to determine which parts of the spectrum are idle and which occupied, however, they are unable to attribute longitudinally-active bands to a single transmitter’s operation. Thus, such approaches cannot facilitate detailed transmitter characterization for next generation spectrum management. Furthermore, these existing techniques operate on the raw PSD measurements, which limits their efficiency in noisy and low SNR regimes. Our work addresses these challenges by transforming the measured PSD signal in the wavelet

coefficient domain, which reveals the underlying transmission structure while reducing the effect of noise. Furthermore, we develop a two-step transmission reconciliation technique that combines detected occupied bands into longitudinally-active transmitters to facilitate detailed spectrum characterization.

**Wavelet-based activity detection.** Prior work has theoretically justified wavelet analysis for spectrum activity detection [16, 18, 24]. Existing approaches, however, (i) have not been considered for transmitter reconciliation, (ii) have not been employed on real-world data and (iii) are not able to autonomously tune the wavelet analysis parameters to the measured spectrum dynamics and noise floor. This renders them inapplicable for unsupervised detection of arbitrary transmitters. In contrast, our proposed method adaptively learns the parameters for wavelet coefficient analysis, performs robustly on noisy, real-world signals and reconciles occupied bands over multiple consecutive spectrum sweeps.

**Supervised and unsupervised characterization.** Another related body of work is signature-based characterization which requires prior knowledge of transmitter activity patterns for detection [9, 15]. While such techniques enable detailed spectrum characterization, they suffer inherent limitations in the number and types of transmitters they can detect, which makes them unfeasible for wide-band characterization of arbitrary transmitters. Our work departs from these early paradigms by developing a robust and unsupervised technique for high-sensitivity transmitter characterization. Prior unsupervised approaches [25] characterize scans using Rayleigh-Gaussian mixture models, however, the method operates on batch as opposed to streaming mode. This renders it inapplicable for real-time transmitter detection at scale.

## III. METHODOLOGY

In spectrum characterization, we differentiate between a *transmission* and a *transmitter* (Fig. 1). We define a *transmission* as a single continuously-occupied time-frequency block. A *transmitter* is, thus, a set of transmissions that are caused by the activity of the same radio-emitter. Under this definition, a broadcast transmitter will be characterized with a single transmission, whereas a TDMA, FDMA or frequency-hopping transmitter will be comprised of multiple transmissions. The task of spectrum characterization can, thus, be split in two key sub-tasks: (i) robust identification of transmissions and (ii) efficient grouping of transmissions into a transmitter. Thus, AirVIEW operates in three stages. As data arrives from the spectrum sensor, AirVIEW performs single-sweep transmission identification (§III.B). A single transmission, however, may span multiple consecutive sweeps, thus the second step in AirVIEW is to combine aligned single-sweep transmissions into a multi-sweep transmission (§III.D). Finally, AirVIEW combines similar-band transmissions into a transmitter (§III.E). In what follows, we first provide the necessary background on wavelet decomposition of spectrum scans. We then detail our proposed approaches to detect transmissions, and group them into transmitters.




Fig. 2. Wavelet decomposition with key features employed by AirVIEW.

**A. Preliminaries.** Let  $p_t(f)$  be the PSD over discrete frequencies  $f \in F$  at time  $t$ , also referred to as a single-time sweep or simply a sweep. Given a wavelet basis, the one-dimensional wavelet decomposition is a function  $W : \mathcal{R}^{|F|} \rightarrow \mathcal{R}^{|F|}$  that maps  $p_t(f)$  to a set of real wavelet coefficients  $w_t$  of the same dimensionality. The coefficients can logically be thought of as a binary-tree hierarchy of increasing scales  $s \in [0, \lceil \log |F| \rceil]$  (Fig. 2), where the 0-th scale corresponds to the root of the hierarchy and the maximum scale corresponds to the hierarchy’s leaves. Thus,  $w_t(s, j)$ , are the coefficients at a given scale  $s$  of the tree, where  $j$  is the index of the coefficient at that scale. For the analysis in this paper we use the *Haar function*, as the wavelet basis, where the mother wavelet is a “square-shaped” function that is piece-wise constant [11]. Our basis choice is informed by the underlying shape of power spectrum measurements characterized with sharp, vertical changes between active and idle frequencies. To construct the Haar wavelet decomposition, one computes a sequence of averages across scales according to the mother wavelet (details in [11]). The decomposition  $W$  is lossless as the original signal  $p_t(f)$  can be exactly reconstructed by a reverse reconstruction function  $W^{-1}$  using the wavelet coefficients  $w_t$ . A lossy reconstruction of the original signal can also be obtained by only considering some of the wavelet coefficients [5]. Low-scale coefficients (close to the root) capture more drastic changes of the underlying signal, whereas high-scale coefficients (close to the leaves) capture finer changes (i.e. jitter due to noise). Thus, a lossy signal reconstruction based on low-scale coefficients will produce a coarse representation of the signal.

The multi-scale product  $\pi_s$  at level  $s$  is defined as  $\pi_s = \langle W^{-1}(w_t(s, \cdot)), W^{-1}(w_t(s+1, \cdot)) \rangle$ , where the right-hand side denotes the element-wise product of the signal reconstructions  $W^{-1}(w_t(s, \cdot))$  using only coefficients at level  $s$  and  $s+1$ . Such multi-scale products have been shown to “elucidate” edges of transmissions while suppressing the effect of noise [18, 16, 24]. A set of products and their corresponding scales in the wavelet tree are presented in Fig. 2. Local maxima in low-scale products (e.g.  $\pi_1$ ) clearly delineate frequency regions in which an abrupt power change occurs, however they lack “spatial” resolution in that they cannot pinpoint the exact frequencies at which a transmission starts. Conversely, high-scale products are less coarse, albeit more susceptible to noise. The medium scale product  $\pi_i$  in the figure offers a good trade-off between robustness to noise, and frequency resolution and is a good candidate for precise detection of transmission edges. Intuitively, our detection algorithm identifies the most

---

### Algorithm 1: Single-sweep transmission detection

---

**Input:**  $p_t(f)$ ,  $s$ ,  $\beta$   
**Output:** A set of transmission intervals  $\mathcal{D}_t = \{[f_i, f'_i]\}$   
1: Compute  $W(p_t(f))$   
2: Compute  $\pi_s(f)$  and  $\hat{\pi}_s(f) = |\pi_s(i) - \pi_s(i-1)|, i \in F$   
3: Detect  $X^\beta = \{x_i\}$ : local maxima of  $\hat{\pi}_s(f)$  exceeding  $\tau_t = \mu_t + \beta\sigma_t$   
4: Compute the average power  $p_t(r_i)$  of frequency regions  $R^\beta = \{r_i\}$  between local maxima  $X^\beta$   
5:  $\mathcal{D}_t = \emptyset$   
6: **for**  $\forall r_i$  in descending power  $p_t(r_i)$  **do**  
7:     **if** Neighbors  $N(r_i)$  of  $r_i$  not in detected  $\mathcal{D}_t$  **then**  
8:          $\mathcal{D}_t = \mathcal{D}_t \cup r_i$   
9:     **end if**  
10: **end for**  
11: RETURN  $\mathcal{D}_t$

---

significant local maxima in the multi-scale product of a carefully-selected coefficient scale to determine the start and end frequencies of transmitter activity captured in a single sweep. The scale and threshold for local maxima’s significance are adaptively determined in an unsupervised manner (§III.C).

**B. Single-sweep transmission detection.** Algorithm 1 presents our single-sweep transmission detection. The method takes as an input the measured PSD signal  $p_t(f)$ , the product scale for detection  $s$  and a threshold scaling parameter  $\beta$ . AirVIEW adaptively learns  $s$  and  $\beta$  in a pre-processing step detailed in §III.C. Our algorithm first computes the wavelet decomposition  $W(p_t(f))$  (Step 1) and the multi-scale product  $\pi_s(f)$  at scale  $s$  (Step 2). We note that depending on the selected scale, the sign of the product value corresponding to a transmitter edge might be positive or negative. Since our edge detection is concerned only with the magnitude but not the sign of the product, we perform local maxima analysis on the absolute pairwise difference between consecutive product values  $\hat{\pi}_s(f) = |\pi_s(i) - \pi_s(i-1)|, i \in F$ . Next (Step 3), we consider all local maxima  $x_i$  of the pairwise differences and keep only the  $\beta$ -outlier ones, i.e. those whose pairwise differences exceed the absolute threshold  $\tau_t$  of the current sweep  $t$ , calculated as  $\tau_t = \mu_t + \beta\sigma_t$ . Here  $\mu_t$  and  $\sigma_t$  are the mean and standard deviation of the pairwise difference of the multiscale products  $\hat{\pi}_s(f)$  at scale  $s$  for sweep  $t$ .  $\beta$  is a threshold-scaling parameter that determines how many standard deviations from the mean should pairwise differences be in order to be ranked as a outlier local maxima.

The positions of outlier local maxima  $X^\beta$  partition the frequency domain into non-overlapping frequency regions  $R^\beta = \{r_i\}$  (e.g. the local maxima in  $\pi_i$  from Fig. 2 partition the scan in three regions  $r_0, r_1$  and  $r_2$ ). We compute the average power  $p_t(r_i)$  from the original PSD signal for all regions  $r_i$  in  $R^\beta$  (Step 4) and consider the regions by decreasing average power (Steps 5-10). First, we initiate the set of detected regions  $\mathcal{D}_t$  (Step 5). Then, for each region  $r_i$ , we check if any of its immediate preceding and succeeding regions, termed neighbors  $N(r_i)$ , are already in the set of output transmissions  $\mathcal{D}_t$ . If the neighbors are not in  $\mathcal{D}_t$  then we include  $r_i$  in  $\mathcal{D}_t$ , else we proceed to the next region by decreasing average power. This sequence of steps is based on the premise that we have successfully detected transmission edges as outlier local maxima and that they will outline regions

of transmission (high power) and noise (low power). Finally, we return the set of detected transmissions  $\mathcal{D}_t$ .

The complexity of Algorithm 1 is  $O(F \log F)$ , since the Haar wavelet decomposition can be computed by averaging with re-use in the hierarchy of size  $F$  and due to the need to sort the adjacent multi-scale product differences (Step 6). The two key parameters in our method are the threshold scaling parameter  $\beta$  and the scale of analysis  $s$ , as these parameters determine AirVIEW's accuracy. In what follows, we describe our unsupervised approach to optimally estimate  $s$  and  $\beta$ .

**C. Alignment-driven parameter estimation.** Our approach exploits a natural domain regularity, namely that multi-sweep transmissions tend to occupy the same band. Hence, the intuition behind our parameter learning is: *given a short interval of sweeps, find the parameter setting for  $\beta$  and  $s$  that results in maximally-aligned detected transmissions across time.*

Two important challenges arise with our approach for optimal selection of the parameters  $\beta$  and  $s$ . First, the magnitude of values in the multiscale product will depend on the inherent characteristics of the measured PSD (including the degree of signal oscillations and the average transmitter power) and the selected scale  $s$ . Thus, in order to adequately select the top-ranked local maxima (Step 3 in Algorithm 1), we need an adaptive and data-driven approach for selection of the corresponding threshold scaling parameter  $\beta$  for AirVIEW to be accurate for arbitrary SNRs and signal variations. Second, low-scale products amplify the edges, but lose the location specificity of the transmissions' edges (e.g.  $\pi_1$  from Fig. 2). High-scale products are more specific in the exact position of edges, though more susceptible to noise (e.g.  $\pi_{\log|F|-1}$  from Fig. 2). Hence, there is a trade-off between edge position specificity and the accuracy of transmission detection that can be controlled by the careful selection of the product scale. We extensively investigate the trade-offs of  $\beta$  and  $s$  selection in both synthetic data with controlled SNR and in real-world spectrum traces (§IV). Our analysis shows that an approach that uses a fixed  $\beta$  and  $s$  leads to a sub-optimal detection performance. Thus, we design an unsupervised method for optimal selection of  $\beta$  and  $s$ , informed by the properties of the underlying PSD signal and formalized as follows.

Let  $D_t^i(f_s^i, f_e^i)$ ,  $i \in [1, |\mathcal{D}_t|]$  be a single detected active band in time  $t$ . Here,  $f_s^i$  and  $f_e^i$  are the start and end frequency of this detected band and  $\mathcal{D}_t$  is the set of all active single-sweep bands  $D_t^i$  at time  $t$ . Intuitively, a well-aligned multi-sweep detection will result in conserved  $f_s^i$  and  $f_e^i$  (or as close as possible) across consecutive sweeps. Following this intuition, we formalize the level of alignment of transmissions detected in consecutive sweeps as the symmetric bi-directional average Jaccard similarity between maximally-aligned detected bands (see Eq. (1)). Here  $|\mathcal{D}_t|$  is the number of detected bands at time  $t$ ,  $D_t^i \cap D_{t-1}^j$  is the number of frequencies in which the two bands overlap, and  $D_t^i \cup D_{t-1}^j$  is the number of frequencies in the union of the two bands. Intuitively, the definition of  $J_t$  averages the best possible Jaccard similarities for every band  $t$  with its predecessor  $t - 1$  and vice versa. Since, individual

Jaccard similarities are constrained within  $[0, 1]$ , so is  $J_t$ .

$$J_t = \frac{1}{2} \left( \frac{1}{|\mathcal{D}_t|} \sum_{D_t^i \in \mathcal{D}_t} \max_{D_{t-1}^j \in \mathcal{D}_{t-1}} \frac{D_t^i \cap D_{t-1}^j}{D_t^i \cup D_{t-1}^j} + \frac{1}{|\mathcal{D}_{t-1}|} \sum_{D_{t-1}^j \in \mathcal{D}_{t-1}} \max_{D_t^i \in \mathcal{D}_t} \frac{D_{t-1}^j \cap D_t^i}{D_{t-1}^j \cup D_t^i} \right), \quad (1)$$

Next, for a spectrum scan comprised of  $T$  sweeps, we calculate the overall detection alignment  $\mathcal{J}$  as the average  $J_t$ :

$$\mathcal{J} = \frac{1}{T-1} \sum_{t=2}^T J_t \quad (2)$$

$\mathcal{J}$  also varies between 0 and 1, where a maximally-aligned detection corresponds to  $\mathcal{J} = 1$ . We use the so-defined detection alignment  $\mathcal{J}$  for parameter estimation. Given a multi-sweep spectrum scan  $P_T^F$  over  $F$  frequency bins and  $T$  spectrum sweeps, we seek  $(\beta_o, s_o)$  that maximizes  $\mathcal{J}$ :  $(\beta_o, s_o) = \arg \max_{\beta, s} \mathcal{J}(P_T^F, \beta, s)$ . Note that in the definitions of  $\mathcal{J}$ ,  $J_t$  and  $\mathcal{D}_t$  we deliberately omitted  $\beta$  and  $s$  for simplified notation, however, all bands are detected by Algorithm 1, and thus, require the parameters as input.

Since we do not know anything about the properties of the function  $\mathcal{J}(P_T^F, \beta, s)$ , the simplest approach to its maximization is to discretize the space of values of  $\beta$  and perform a scan over parameter value combinations. We employ this approach in a training phase, in which AirVIEW learns the optimal  $(\beta, s)$  for a small sub-scan, and then applies this combination for the remainder of the spectrum characterization campaign.

The complexity of parameter learning is  $O(BTF \log^2 F)$ , where  $B$  is the number of discrete values  $\beta$  considered in the parameter estimation,  $T$  is the number of sweeps in the training scan, a factor of  $O(F \log F)$  is added for each invocation of the single-sweep detection (Algorithm 1), and finally an extra factor of  $\log F$  is added for the possible number of scales  $s$  in the scan, since the wavelet tree height is logarithmic in  $F$ . As we demonstrate in §IV, a small number of sweeps  $T$  and relatively coarse granularity for  $\beta$  (i.e. small  $B$ ) are sufficient to robustly learn  $\beta$  and  $s$ . In addition, our empirical evaluation suggest that  $\mathcal{J}(\beta, s)$  behaves similar to a concave function w.r.t. both of its parameters, so a simple hill-climbing approach can reduce the factor of  $O(B \log F)$  to a constant without compromising quality if frequent and fast parameter estimation is warranted in non-stationary scenarios.

Several important questions arise with our approach to unsupervised parameter estimation. *Is the alignment of transmitter detection a good proxy metric for accuracy of transmitter detection? How long do we need to sense the spectrum before AirVIEW is able to robustly learn the optimal  $(\beta, s)$ ?* Our evaluation (§IV) provides empirical answers to these questions.

**D. Multi-sweep transmission detection.** AirVIEW's multi-sweep transmission detection, detailed in Algorithm 2, takes as an input a continuous stream of PSD sweeps  $p_t(f)$ ,  $s$ ,  $\beta$ , and an additional temporal smoothing parameter  $\lambda$ . The algorithm reports detected transmissions in contiguous sweeps  $T = \{(D, t)\}$ , where  $D$  is the frequency interval of an instantaneous transmission at time  $t$ . We begin by initializing

---

**Algorithm 2: Multi-sweep transmission detection**

---

**Input:**  $p_t(f)$ ,  $s$ ,  $\beta$ ,  $\lambda$   
**Output:** A set of temporal transmissions  $\mathcal{T} = \{T = \{(D, t)\}\}$   
1: Initialize temporal transmissions  $\mathcal{T} = \emptyset$   
2: **for** Consecutive  $t$  consider PSD signal  $p_t(f)$  **do**  
3:  $\mathcal{D}_t \leftarrow \text{SingleSweep}(p_t(f), s, \beta)$  {Alg. 1}  
4: **for**  $\forall D_t \in \mathcal{D}_t$  **do**  
5: **if**  $D_t$  matches a temporal transmission  $T \in \mathcal{T}$  **then**  
6: extend  $T$  by  $\text{smooth}(D_t, T, \lambda)$   
7: **else**  
8: Start a new active transmission  $T \leftarrow (D, t)$   
9: Add  $T$  to the set of active transmissions  $\mathcal{T}$   
10: **end if**  
11: **end for**  
12: Report transmissions in  $\mathcal{T}$  that were not extended in  $t$   
13: **end for**

---

the list of active transmissions (Step 1). We then process the consecutive temporal PSD signals (Steps 2-11). For each time  $t$  we detect all transmissions  $D_t \in \mathcal{D}_t$  (Step 3) and process them one at a time (Step 4). We match each transmission interval  $D$  to the list of active transmissions  $\mathcal{T}$  and if its intersection with  $T$  is at least half the span of  $D$  we declare it a match and add a temporally-smoothed version of  $D$  to  $T$  (Step 6). If no match is found, we initiate a new active transmission  $T$  and add it to the set of active transmissions  $\mathcal{T}$  (Step 8). After all instantaneous transmissions  $D_t$  are processed, we report transmissions in  $\mathcal{T}$  that were not extended in time  $t$  (Step 12) and proceed to the next scan  $p_{t+1}(f)$ .

In order to reconcile an instantaneous transmission  $D_t$  with its matched temporal transmission  $T$ , we consider the edges of  $D_t$  and those of all preceding transmissions in  $T$  in function  $\text{smooth}(D_t, T, \lambda)$  (Step 6). Let  $(f_i, \tau_i)$  be the left edge positions of all transmissions in  $T$  including that of  $D_t$  at time  $t$ . We compute the time-decayed weighted average of those edge positions as  $\bar{f} = \lceil \sum_i w_i f_i / \sum_i w_i \rceil$ , where  $w_i = e^{-\lambda|t-\tau_i|}$  is exponentially time-decaying weight giving preference to more recent instantaneous transmissions in  $T$  and  $\lambda \geq 0$  is an exponential smoothing parameter. We perform similar smoothing average  $\bar{f}'$  for the right edge. A large smoothing parameter  $\lambda$  makes the contribution of past transmissions negligible and hence preserves the detected edges of  $D_t$  without smoothing. Alternatively, when  $\lambda = 0$  all past edges in  $T$  are weighted equally and the edge averages are unweighted means of all past edges, while values in between result in exponentially-decaying importance of past transmissions. Note, that we use the independently detected edges of  $D_t$  for smoothing, but we report the smoothed versions of the instantaneous transmissions. As demonstrated in our experiments temporal smoothing using small non-zero values of  $\lambda$  helps minimize the “shifting” of consecutive transmissions caused by noise and low-scale  $s$  product detection.

**E. Transmitter reconciliation.** As illustrated in Fig. 1, a single longitudinally-active transmitter may be comprised of multiple transmissions. We employ a simple approach that combines transmissions into a transmitter based on the transmissions’ relative frequency position. For all multi-sweep transmissions  $T$ , we take the union of individual single-sweep transmission intervals and group  $T$ s whose union intervals

overlap by at least 90% of their extent. Note, that this is a simple solution that cannot handle frequency-hopping transmitters and is oblivious to other transmitter properties such as power level and transmissions inter-arrival time. Handling such cases is important for detection of arbitrary transmitters (i.e. frequency-hopping), however, it is beyond the scope of this paper. We leave such extensions for future work and focus on frequency-aligned transmitters in our evaluation.

## IV. EVALUATION

In this section we evaluate the accuracy of AirVIEW in detecting transmissions in both real and synthetic data sets and in comparison with baseline methods.

**A. Implementation, data and baseline approaches.** Our current implementation of AirVIEW is a single-core Java program and all experiments are executed on commodity desktop machines. Our eventual system implementation will, however, make use of recent advances in fast parallel wavelet decomposition for general purpose and specialized architectures such as FPGA and GPUs [7]. In addition, we plan to bring the decomposition and detection “closer to the sensor” for on-sensor and collaborative spectrum characterization.

We evaluate AirVIEW on both synthetically-generated data, in which we control the signal-to-noise ratio, and in a real-world spectrum scans of TV channels, in which we have ground truth position of transmitters. We use two baselines for comparison. *Naive* is a thresholding scheme that detects outlier local maxima in the original PSD  $p_t(f)$  signal as opposed to the multi-scale product  $\pi_s$ . Extraction of the transmissions is done in a similar way to AirVIEW, once edges are detected. A second baseline *Denoised* follows the general idea of lossy wavelet reconstruction by maintaining the most important coefficients based on magnitude and scale [5].

**B. Robustness to noise.** One of the main challenges in wide-band spectrum characterization is detection of transmission in a noisy environment, where noise is introduced both due to the environment and imperfections of the sensor. Hence, we first focus on evaluation of AirVIEW for varying signal-to-noise ratios (SNR). To control SNR, we synthesize realistic spectrum scans of similar characteristics to those we capture using sensors and vary the SNR by decreasing the mean signal power to levels very close to that of the noise. Similar to the noise in radio frequency signals [8] we add additive white Gaussian noise using a normal distribution with mean  $-110dBm$  (our “noise floor”) and variance  $4.0dBm$ . We randomly select regions in time and frequency in which we inject transmissions of a desired power mixed with Gaussian noise of the same magnitude as non-transmission regions.

Since we know the ground-truth position of instantaneous transmission in this synthetic data, our evaluation seeks to quantify how closely the real transmissions are recovered by the competing techniques. Let  $\mathcal{D} = \{D_i\}$  be the set of detected instantaneous transmissions at time  $t$  ( $t$  is omitted for notation simplicity). Also, let  $\mathcal{A} = \{A_i\}$  denote the set of actual ground truth transmissions at the same time. We define a true positive






Fig. 3. (a) TPR and (b) FPR of two instantiations of AirVIEW ( $s = 4/\beta = 2$  and  $s = 4/\beta = 4$ ) and the competing techniques Naive and Denoised at varying signal levels with noise floor fixed at  $-110\text{dBm}$ .

rate (TPR) measure based on how well we “cover” actual transmissions  $A_i$  by detected transmissions  $D_i$ , as follows:

$$TPR = \frac{1}{|\mathcal{A}|} \sum_{A_i \in \mathcal{A}} \max_{D_j \in \mathcal{D}} \frac{|A_i \cap D_j|}{|A_i|} \quad (3)$$

where  $|A_i \cap D_j|$  denotes the length of the intersection interval common to  $A_i$  and  $D_j$ . Intuitively, TPR will be 1 if all actual transmissions are completely “covered” by detected transmissions. We also want to measure the rate of false detection. The false positive rate is similarly defined as:

$$FPR = \frac{1}{|\mathcal{D}|} \sum_{D_i \in \mathcal{D}} \min_{A_j \in \mathcal{A}} \frac{|D_i| - |D_i \cap A_j|}{|D_i|} \quad (4)$$

An ideal method will have a TPR of 1 and FPR of 0.

Fig. 3 presents the performance of AirVIEW on synthetic data with varying SNR. We compare the TPR and FPR of two variations of AirVIEW ( $s = 4, \beta=2$  and 4), Naive and Denoised. At higher SNR ( $-85\text{dBm}$ ,  $-88\text{dBm}$ ,  $-91\text{dBm}$ ), Naive and Denoised are able to detect transmitters with high accuracy. However, as the signal approaches the noise floor, their performance deteriorates quickly resulting in rapid increase of FPR and decrease in TPR. AirVIEW’s TPR and FPR, on the other hand, remain robust to noise for very low signal to noise ratios (Figs. 3(a), 3(b)). It is important to note that the slightly lower TPR of AirVIEW in high SNR cases is due to its lack of specificity of edge position for low scales  $s$ , which triggers AirVIEW to designate as active slightly larger frequency ranges than the actual ones. As a result AirVIEW is more conservative in placing the edges compares to Naive and Denoised. Overall, AirVIEW exhibits satisfactory performance with high SNRs and significantly outperforms its counterparts Naive and Denoised when the signal to noise ratio is low.

**C. Effect of the scale  $s$  and threshold  $\beta$ .** Beyond the robustness to noise we also study TPR and FPR for varying  $\beta$  and  $s$ . Fig. 4 summarizes the TPR and FPR achieved by AirVIEW when detecting transmissions in both synthetic and real-world PSD data. In synthetic data (Fig. 4(a), 4(b)), there is a clear trend of reaching perfect TPR and FPR when using relatively low scales ( $s = 2, 4$ ) and  $\beta > 2$  for  $s = 2$  and

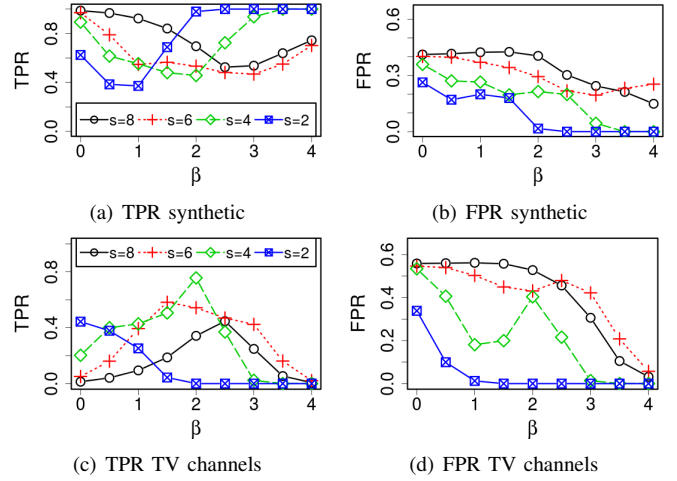


Fig. 4. (a),(b): Average accuracy of detection in synthetic data using various parameter setting combinations of  $s$  and  $\beta$ . Over 1000 runs, 2 transmitters with mean power of  $-88$  were randomly placed in the data and detection was applied using 4 different scales. (c),(d): Similar evaluation in real-world sensed TV band spectrum for constantly transmitting stations of known bands.

$\beta > 3$  for  $s = 4$ . Higher scales are more affected by noise and thus do not exhibit as favorable TPR vs. FPR trends for any  $\beta$ . We collect our real-world groundtruth data over 25MHz in the UHF TV bands. This scan contains two active broadcasts, each of which we annotate as a single transmission with smooth edges in frequency. Our results in Fig. 4(c) and 4(d), indicate that, similar to the synthetic data, lower scales ( $s = 2, 4$ ) and slightly lower thresholds  $\beta < 2$  achieve the optimal TPR vs. FPR regime. Of note is that our real-world benchmark achieves slightly worse TPR and FPR as compared to synthetic data. This behavior can be attributed to the noisy nature of the sensed transmitters that causes some single-time scans to be detected as idle due to fading. At the same time, these scans were annotated as occupied.

**D. Alignment-driven parameter estimation.** Our results from Fig. 4 show that AirVIEW’s performance is dependent on  $\beta$  and  $s$ . Thus, we design a data-driven approach to parameter estimation, which we presented in §III.C. Some key questions arise with our approach. First, we use the detection alignment  $\mathcal{J}$  as a criteria to select the optimal  $(\beta, s)$ . In this section we justify this choice by demonstrating that detection alignment is a good proxy metric for detection accuracy. Second, we show that, indeed maximal detection accuracy occurs at different  $\beta$  as  $s$  increases, which empirically demonstrates the need and benefits of adaptive parameter selection. Third, we quantify the amount of spectrum sweeps necessary for AirVIEW to robustly learn the optimal  $(\beta, s)$ .

1) *Detection alignment is a good proxy for accuracy:* We use two metrics in the following evaluation: alignment and accuracy. Our definition of alignment is as in Eq. (2). We define accuracy as:

$$Accuracy = \frac{1}{|\mathcal{D}|} \sum_{D_i \in \mathcal{D}} \max_{A_j \in \mathcal{A}} \frac{|D_i \cap A_j|}{|A_j|} \quad (5)$$

Intuitively, our definition of accuracy encapsulates both TPR and FPR, making it an appropriate metric for overall




Fig. 5. Accuracy and alignment at mean transmit power of  $-103\text{dBm}$  for various  $(\beta, s)$  combinations. We vary scale from 8-9 in (a) to 6-7 in (b). On each graph we vary  $\beta$  from  $-1$  to  $4$  in increments of  $0.2$ . Across all  $(\beta, s)$  combinations the accuracy and alignment follow a similar trend, which indicates that alignment is a good proxy for accuracy.

evaluation of AirVIEW’s performance. Accuracy will be low if we fail to detect transmissions (i.e. our TPR is low) or if we falsely detect transmissions (i.e. our FPR is high). Before we delve in our evaluation, we note that accuracy of transmitter detection can only be calculated if we have ground-truth data for transmitter activity. It is thus impossible to use accuracy for unsupervised transmitter detection (i.e. without prior knowledge of transmitter activity). Alignment, on another hand, does not require ground-truth and is thus ideal for unsupervised parameter estimation. Nevertheless, it is essential to evaluate the ability of alignment-based parameter estimation to gain maximal accuracy. Thus, our evaluation necessarily requires tight control over transmitter configuration and SNR regime, and thus, uses synthetic data.

Our evaluation scans are comprised of 30 sweeps, each of which contains 1024 PSD values. Within each  $30 \times 1024$  scan, we instantiate ten randomly-positioned, non-overlapping transmitters. Noise values in these synthetic scans were drawn from a normal distribution with a mean of  $-109\text{dBm}$  and a standard deviation of  $2.0$ , whereas the transmitter values were drawn from normal distributions with decreasing means and standard deviation of  $2.0$ . These cases present increasingly-challenging, yet realistic scenarios. Specifically, all scenarios are challenging due to the high variation of the generated values. The last two cases are particularly challenging due to the low power of the generated transmitters. The scenarios are realistic, since the assigned standard deviations are informed by real-world spectrum measurements in the UHF band.

Fig. 5 presents accuracy and alignment for our most challenging scenario with transmitter’s mean power of  $-103\text{dBm}$  (only  $6\text{dBm}$  above the noise floor). We experimented with decreasing scales from 8-9 to 5-6, while varying beta from  $-1$  to  $4$  in regular increments of  $0.2$ . In interest of space, the figure only presents two scale combinations, however, the trends were similar in the other results. Across all scales, we see that accuracy and alignment follow similar trends, which indicates that regardless of the selected  $(\beta, s)$  combination, *alignment is always a good proxy for accuracy*. We also note that for this particular scenario AirVIEW chooses scale 8-9 and  $\beta = 0$  as optimal, since at this  $(\beta, s)$  combination the detection is maximally-aligned.

We extend this analysis to multiple SNR regimes in Fig 6. The figure shows the maximum accuracy (red continuous line)




Fig. 6. Maximum and achieved accuracy with decreasing transmitter PSD. Alignment is a good proxy for accuracy across all PSD regimes. Achieved accuracy is high, persistent and close to the max accuracy.

Fig. 7.  $\beta$  at max accuracy and  $\beta$  at max alignment for increasing scales. As  $s$  grows, max accuracy and max alignment correspond to decreasing  $\beta$ s. This underlines the importance of adaptive parameter estimation.

and accuracy of the maximally-aligned detection (dotted blue line) for increasing mean PSD of the generated transmitters. The *max accuracy* is the largest achieved accuracy over all  $(\beta, s)$  combinations. The *accuracy at max alignment* is the accuracy achieved with the optimal  $(\beta, s)$  learned by AirVIEW’s unsupervised parameter estimation. Fig. 6 allows several important conclusions. First, since the accuracy at max alignment is well-coordinated with the max accuracy, we can conclude that alignment is a good proxy for accuracy across various SNR regimes. Second, we show that AirVIEW’s accuracy is high, persistent and close to the maximal accuracy with decreasing transmitter power. Third, AirVIEW is highly-accurate even in the most challenging scenarios, where the transmitters’ PSD is only  $6\text{dBm}$  higher than the noise floor and the signal variance is as large as  $2.0\text{dBm}$ .

2) *Benefits of adaptive selection of  $(\beta, s)$* : We now demonstrate the need for adaptive selection of parameters  $(\beta, s)$ . Our results are presented in Fig. 7, that plots the optimal  $\beta$  as the scale increases from 4-5 to 8-9. Red continuous line presents  $\beta$  at max accuracy, while blue dotted line presents  $\beta$  at max alignment. As the scale increases (i.e. towards the leaves), the optimal  $\beta$  decreases for both max accuracy and max alignment. Thus, in order to maintain optimal characterization performance, AirVIEW needs to adaptively select  $(\beta, s)$  for a given spectrum sensing environment. A fixed selection of  $(\beta, s)$  will result in sub-optimal accuracy.

3) *Training period*: Finally, we evaluate the effects of training duration. Let  $P_T^F$  be a PSD spectrum scan comprised of  $T$  sweeps and  $F$  frequency bins in each sweep. Let  $Accuracy_t$  be the transmitter detection accuracy over the entire scan duration  $T$  when parameter estimation was performed on a subset of all the sweeps of size  $t$ . Here,  $Accuracy_t$  is calculated as in Eq. (5). We define *stability* as the ratio of  $Accuracy_t$  over  $Accuracy_T$ . Intuitively, this stability measure captures the relative difference in detection accuracy when training on partial vs. complete spectrum scan.

We run AirVIEW on synthetic scans. Each scan is comprised of 120 sweeps, each containing 1024 PSD values. Within each  $120 \times 1024$  scan we injected 40 randomly-located, non-overlapping transmitters. Each of these transmitters spans 80 frequency bins and 5 sweeps. The noise values in this scan were drawn from a normal distribution with mean  $-109\text{dBm}$  and standard deviation of  $2.0\text{dBm}$ . The transmitter values were




Fig. 8. Parameter estimation with increasing number of training sweeps.




Fig. 9. Effect of smoothing on transmitter detection: (left) no smoothing ( $\lambda = \infty$ ), and (right) max smoothing ( $\lambda = 0.0$ ).

drawn from a distribution with mean  $-85\text{dBm}$ . Fig. 8 presents average and standard deviation of stability across 100 runs for an increasing number of training sweeps. AirVIEW learns the optimal  $(\beta, s)$  with as few as 40 training sweeps. Any additional training further stabilizes the parameter estimation as indicated by the decrease of standard deviation.

**E. Detection in time.** In Section III.D we introduced a method to “smooth” the detected edges of individual transmissions accounting for the “fuzzy” nature of transmissions over time caused by noise in measurements. We qualitatively present the effect of smoothing in a synthetic transmitter in Fig. 9. To generate this data, we first choose a starting location and bandwidth for the transmission and introduce “drift” patterns in the transmission location in the frequency domain. In each instantaneous transmission, we randomize the position of the edges from a fixed *model* position according to a normal distribution. To introduce drift, we bias the random offsets in a positive or negative direction for fixed windows of time. Without any smoothing (Fig. 9(left),  $\lambda = \infty$ ) the transmissions shift frequently and disagree due to the drift. As we decrease  $\lambda$  (more pronounced smoothing) the instantaneous transmissions start to agree over time (Fig. 9(right)). It is worth noting that this scenario presents a particularly challenging instance and real-world scans require moderate levels of smoothing.

## V. WIDE-BAND SPECTRUM ANALYSIS WITH AIRVIEW

**Measurement and analysis.** We collect a wide-band spectrum trace from  $50\text{MHz}$  to  $4.4\text{GHz}$  in an urban area in close proximity to an airport and a military airbase. Our measurement platform consists of two USRP N210s; one with WBX daughterboard ( $50\text{MHz} - 2.2\text{GHz}$ ) and one with SBX daughterboard ( $400\text{MHz} - 4.4\text{GHz}$ ). The two USRPs were connected through a splitter to a single multi-polarized, wide-band antenna that operates in the range of  $30\text{MHz}$  to  $6\text{GHz}$ . We used a quad-core Lenovo laptop with  $16\text{GB}$  of RAM running Linux as a host for our software defined radio setup. We implemented a spectrum scanner on top of Gnuradio and configured it to sweep the target spectrum in  $25\text{MHz}$  steps with dwell time of 6 seconds. At each step, the scanner ran with a sample rate of  $25\text{Msps}$  and calculated the PSD with FFT size of 1024. This measurement campaign produced a total of  $134\text{GB}$  of data, where each  $25\text{MHz}$  chunk was swept 140,000 times. We run AirVIEW on the above spectrum trace with  $s = 6$  and  $\beta = 3$  and  $\lambda = 0.5$ . Our transmitter reconciliation sets the minimum frequency overlap to 90%.

**Wide-band characterization with AirVIEW.** Our analysis shows that 74% of all bands were idle. For the active 26% of

the spectrum, we show the number of detected transmitters per  $25\text{MHz}$  band in Fig. 10(left). 26% of all the bands have at least one transmitter. AirVIEW also finds 157 transmitters in the bands  $225\text{MHz} - 328\text{MHz}$  and  $335\text{MHz} - 400\text{MHz}$ , allocated for federal/military use [1]. We find a large number of transmitters, 122, in the amateur radio bands ( $420\text{MHz} - 450\text{MHz}$ ). Finally, AirVIEW robustly identifies TV transmitters, whereby with two exceptions, it finds between 1 and 4 transmitters in a  $25\text{MHz}$  chunk. In two of the bands,  $547.5\text{MHz} - 572.5\text{MHz}$  and  $659.5\text{MHz} - 684.5\text{MHz}$  we observe 14 and 12 transmitters, respectively, which is due to the noisy nature of the measured TV channels triggering the imperfections in our transmitter reconciliation technique.

We also demonstrate AirVIEW’s utility for wide-band analysis of temporal transmitter characteristics. We measure the transmitters’ active time, cycle and gap, where cycle is defined as the time between the beginnings of consecutive transmissions and the gap, as the time between the end of one transmission and the beginning of the next. These metrics allow us to quantify the fraction of time a band is available for secondary access, and whether the temporal behavior of primary users is predictable. Fig. 10 (middle, right) present average (black) and standard deviation (green envelope) of cycle and gap of detected transmitters. Each value is averaged over all transmitters in a  $25\text{MHz}$  chunk. In terms of active time (not shown in interest of space), broadcasts, such as TV, are active either for the entire 6 seconds or appear as 3-second transmissions, due to intermittent deterioration in the scanned signal. Furthermore, they exhibit no variance in the cycle or gap. Other bands such as some federal/military bands ( $225\text{MHz} - 328\text{MHz}$  and  $335\text{MHz} - 400\text{MHz}$ ) and the amateur radio band ( $420\text{MHz} - 450\text{MHz}$ ) have smaller active times, however, their cycle size and gaps vary drastically and will, thus, be hard to predict. For example, the federal/military bands have a typical active time of less than a second and a large gap of 1–2 seconds between consecutive transmissions. However the variance of the cycle and gap are as high as 0.31 and 0.25, which makes the DSA opportunity hard to predict.

**Discussion.** AirVIEW enables wide-band unsupervised spectrum analysis and can thus play an essential role in DSA technology, spectrum policy and spectrum enforcement. In support of *DSA technology*, AirVIEW can answer what is the secondary access opportunity across various frequency bands. In support of *spectrum policy*, AirVIEW can answer how efficiently is spectrum utilized by mapping the number of transmitters in a wide frequency range along with the size and predictability of their active time. *Spectrum enforcement* engines can leverage AirVIEW’s single-pass detection to pinpoint transmitter patterns that deviate from the expected or the historically-observed patterns in a particular frequency range and alert spectrum enforcement authorities.

## VI. DISCUSSION AND CONCLUSION

Real-world spectrum sensing and automated analysis have emerged as key challenges towards opportunistic access. Spectrum scans, however, are plagued with high signal variation,






Fig. 10. Wide-band spectrum analysis with AirVIEW. (left) Transmitter count indicates that 74% of all bands are unused. (mid.) Cycle and (right) gap analysis shows that some intermittently-occupied bands have predictable incumbent patterns while in others the incumbent activity is very unpredictable.

which makes automated analysis extremely challenging. This reality is further aggravated when transmitters are sensed with low power level. To address these challenges and enable robust and rapid detection of low-power transmitters in noisy spectrum scans, we design AirVIEW. AirVIEW is the first mechanism for unsupervised transmitter characterization that enables a single-pass transmitter detection with high sensitivity to low-power signals. We explored key trade-offs in parameter selection and demonstrated AirVIEW's ability to robustly detect synthetic and ground-truth transmitters. We also employed AirVIEW to create a map of spectrum use and characterize the DSA opportunity over a wide range of frequencies.

AirVIEW presents a significant departure from existing spectrum characterization and activity detection techniques across several criteria. First, by its operation on an alternative (wavelet-based) representation of raw spectrum scans, AirVIEW is able to combat the adversary effects of low-power and high-variability signals that has been a historic challenge in classical techniques. Second, classical spectrum characterization and activity detection methods operate in the raw PSD domain and are, thus, sensitive to a changing noise floor. AirVIEW overcomes this limitation by adopting a data-driven approach to parameter estimation that allows adaptive parameter tuning in the face of a changing noise floor to maintain high detection accuracy. Third, AirVIEW employs post-processing techniques that robustly reconcile single-sweep into multi-sweep transmissions for holistic spectrum characterization. Last but not least, this capability is achieved in a fully-unsupervised fashion that does not require prior knowledge of spectrum properties or extensive training.

While AirVIEW achieves robust transmission detection, there are several open problems that shape a compelling future research agenda. Most notably, the current implementation of AirVIEW features a simple heuristic to group transmissions into transmitters that is only informed by the size and overlap of transmissions in frequency. This approach fails to detect frequency-hopping incumbents and is oblivious to transmissions' power level and co-occurrence. A more robust technique is necessary for a single-pass arbitrary transmitter detection.

In light of these important future directions, AirVIEW lays the foundations for next generation DSA technology and spectrum policy. We enable robust, single-pass detection in wide-band spectrum measurements that is able to answer thus-far open questions related to number of incumbents and the opportunity they grant for secondary spectrum access.

## VII. ACKNOWLEDGEMENTS

This work was supported through a NSF CISE Research Initiation Initiative (CRII) grant CNS-1657476.

## REFERENCES

- [1] National Telecommunications and Information Administration, Federal Government Spectrum Use Reports 225MHz – 5GHz. <https://www.ntia.doc.gov/page/federal-government-spectrum-use-reports-225mhz-5ghz>.
- [2] NSF Workshop on Spectrum Measurement Infrastructures, Illinois Institute of Technology, Chicago, IL, USA. [http://www.cs.albany.edu/~mariya/nsf\\_smsmw/](http://www.cs.albany.edu/~mariya/nsf_smsmw/), April 6-7 2016.
- [3] E. Azzouz and A. Nandi. *Automatic modulation recognition of communication signals*. Springer Science & Business Media, 2013.
- [4] E. Candes and M. Wakin. An introduction to compressive sampling. *Signal Processing Magazine, IEEE*, 25(2):21–30, March 2008.
- [5] K. Chakrabarti, M. Garofalakis, R. Rastogi, and K. Shim. Approximate query processing using wavelets. *The VLDB Journal—The International Journal on Very Large Data Bases*, 10(2-3):199–223, 2001.
- [6] F. F. Digham, M. S. Alouini, and M. K. Simon. On the energy detection of unknown signals over fading channels. *IEEE Transactions on Communications*, 55(1):21–24, Jan 2007.
- [7] J. Franco, G. Bernabé, J. Fernández, and M. E. Acacio. A parallel implementation of the 2D wavelet transform using CUDA. In *2009 17th Euromicro International Conference on Parallel, Distributed and Network-based Processing*, pages 111–118. IEEE, 2009.
- [8] A. Goldsmith. *Wireless communications*. Cambridge Univ Pr, 2005.
- [9] S. S. Hong and S. R. Katti. Dof: A local wireless information plane. *SIGCOMM Comput. Commun. Rev.*, 41(4):230–241, Aug. 2011.
- [10] J. Laska, W. Bradley, T. Rondeau, K. Nolan, and B. Vigoda. Compressive sensing for dynamic spectrum access networks: Techniques and tradeoffs. In *DySPAN'11*, Aachen, Germany, May 2011.
- [11] S. Mallat. *A wavelet tour of signal processing*. Academic press, 1999.
- [12] M. McHenry. Spectrum measurements requirements survey. [http://www.cs.albany.edu/~mariya/nsf\\_smsmw/docs/NSF\\_SMIW\\_survey\\_results.pdf](http://www.cs.albany.edu/~mariya/nsf_smsmw/docs/NSF_SMIW_survey_results.pdf), April 6-7 2016.
- [13] M. McHenry, E. Livsics, T. Nguyen, and N. Majumdar. Xg dynamic spectrum access field test results [topics in radio communications]. *IEEE Communications Magazine*, 45(6):51–57, June 2007.
- [14] Y. Meyer. *Wavelets - Algorithms and applications*. 1993.
- [15] S. Rayanchu, A. Patro, and S. Banerjee. Airshark: Detecting non-WiFi RF Devices Using Commodity WiFi Hardware. *IMC '11*, pages 137–154. ACM, 2011.
- [16] B. M. Sadler and A. Swami. Analysis of multiscale products for step detection and estimation. *IEEE Transactions on Information Theory*, 45(3):1043–1051, 1999.
- [17] P. Schmitt, D. Iland, M. Zheleva, and E. Belding. HybridCell: Cellular connectivity on the fringes with demand-driven local cells. In *IEEE INFOCOM '16*, San Francisco, CA, USA, April 2016.
- [18] Z. Tian and G. Giannakis. A wavelet approach to wideband spectrum sensing for cognitive radios. In *CROWNCOM*, Mykonos Island, Greece, June 2006.
- [19] Z. Tian and G. B. Giannakis. Compressed sensing for wideband cognitive radios. In *IEEE ICASSP'07*, Honolulu, HI, April 2007.
- [20] L. Yang, W. Hou, L. Cao, B. Y. Zhao, and H. Zheng. Supporting Demanding Wireless Applications with Frequency-Agile Radios. *NSDI'10*, San Jose, California, 2010.
- [21] Y. Yuan, P. Bahl, R. Chandra, P. A. Chou, J. I. Ferrell, T. Moscibroda, S. Narlanka, and Y. Wu. Knows: Cognitive radio networks over white spaces. In *DySPAN 2007*, pages 416–427, April 2007.
- [22] T. Yucek and H. Arslan. A survey of spectrum sensing algorithms for cognitive radio applications. *IEEE Communications Surveys Tutorials*, 11(1):116–130, Jan 2009.
- [23] T. Yucek and H. Arslan. A survey of spectrum sensing algorithms for cognitive radio applications. *IEEE Communications Surveys Tutorials*, 11(1):116–130, 2009.
- [24] L. Zhang and P. Bao. Edge detection by scale multiplication in wavelet domain. *Pattern Recognition Letters*, 23(14):1771 – 1784, 2002.
- [25] M. Zheleva, A. Chowdhery, R. Chandra, A. Kapoor, and P. Garnett. TxMiner: Identifying Transmitters in Real-World Spectrum Measurements. In *IEEE DySPAN'15*, Stockholm, Sweden, September 2015.