

SEMI-SUPERVISED EIGENBASIS NOVELTY DETECTION

DAVID R. THOMPSON^{1,2}, WALID A. MAJID², COLORADO J. REED², AND KIRI L. WAGSTAFF²

ABSTRACT. We present a semi-supervised online method for novelty detection and evaluate its performance for radio astronomy time series data. Our approach uses sparse, adaptive eigenbases to combine 1) prior knowledge about uninteresting signals with 2) online estimation of the current data properties to enable highly sensitive and precise detection of novel signals. We apply Semi-Supervised Eigenbasis Novelty Detection (SSEND) to the problem of detecting *fast transient* radio anomalies and compare it to current alternative algorithms. Tests based on observations from the Parkes Multibeam Survey show both effective detection of interesting rare events and robustness to known false alarm anomalies.

1. INTRODUCTION

Recent discoveries in high time resolution radio astronomy data have drawn attention to a new class of sources. *Fast transients* are rare pulses of radio-frequency energy lasting from microseconds to seconds that might be produced by a variety of exotic astrophysical phenomena [6, 5, 12, 13]. For example, X-ray bursts, neutron stars, active galactic nuclei, and extraterrestrial intelligence (ETI) are all potential sources of short-duration transient radio signals. Such events are often discovered serendipitously in data collected for other purposes. These transients are generally faint and subtle, so improved detection algorithms can directly benefit all such commensal monitoring. Existing detection approaches rely on a *dispersed pulse* model of the signal shape. This paper presents a new method for analyzing real-time high-resolution radio astronomy data that operates without this model assumption. Therefore, it can potentially detect a far broader class of anomalous events in real time, as well as unexpected events that do not match a known profile.

We have formulated fast transient monitoring as a time series statistical anomaly detection problem [17, 18]. The main challenges of our domain are:

- **High dimensionality:** Signals of interest span multiple antenna power measurements that could include hundreds of time steps and frequency channels.
- **Real time processing:** With the exception of a few dedicated surveys, most high time resolution data is too voluminous to archive. Therefore, events must be detected in real time to select only the most interesting candidates for storage and later exhaustive analysis.
- **Nonstationarity:** Background noise characteristics change over time on medium to long scales, manifesting as narrow-band noise or large-scale gain fluctuations that change with hardware and observing conditions. Detection of anomalous “fast” signals should be robust to these effects.
- **False alarms:** Certain known classes of events, such as momentary Radio Frequency Interference (RFI), are not astronomically interesting but are easily mistaken for fast transients. It is important to avoid flagging these events as novel to avoid filling the detection buffer with these false alarm events. Further, false alarms waste valuable astronomer time in reviewing the results.

This work proposes a new solution that learns a low-dimensional linear manifold for describing the “normal” data. The novelty of our approach lies in combining basis vectors learned in an unsupervised, online fashion from the data stream with supervised basis vectors learned in advance

¹david.r.thompson@jpl.nasa.gov

² Jet Propulsion Laboratory, California Institute of Technology.

from known false alarms. We thereby achieve adaptive, data-driven anomaly detection that also exploits prior domain knowledge about signals that may be statistically anomalous but are not scientifically interesting (and should therefore be ignored). We identify truly interesting anomalies by compressing and reconstructing the data [9] using the combined basis. High reconstruction error indicates a signal that does not match the learned profile of the normal data. The unsupervised component uses the incremental method of Ross et al. [14, 19], an efficient online algorithm that can run in real time.

We evaluated semi-supervised novelty detection using data from the Parkes Multibeam Survey. This data set was originally collected to search for pulsars, which are astronomical sources that emit radio pulses at regular periods. However, several non-pulsar anomalies have recently been discovered in this dataset [3], making it a compelling test case. We found that by explicitly filtering known false alarm patterns, semi-supervised anomaly detection yields significantly better performance than state-of-the-art transient detection methods. This method shows promise for use in current and future astronomical surveys, including data to be collected by the Square Kilometre Array, a radio telescope currently under development that will be 50 times more sensitive than any existing instrument.

We presented the core idea of Semi-Supervised Eigenbasis Novelty Detection (SSEND) at the October 2011 NASA Conference on Intelligent Data Understanding [22]. This paper enhances SSEND by using local context to isolate false alarm signals from their incidental background. This minor change significantly improves precision (see Section 3.2 and Figure 7). We also show a sparse PCA formulation which offers better interpretability for known uninteresting signals (Section 3.2). More generally, this sparse version demonstrates how the basic SSEND approach can incorporate alternative basis learning techniques.

2. RELATED WORK

Generic approaches to anomaly detection are data-driven: they typically learn a representation of the “normal” or uninteresting data, then identify any observations that do not match this model. One such method is one-class support vector machine (SVM) classification [20], in which an SVM is trained only on examples from the normal class and then detects any new data belonging to a different, previously unobserved class. More recent efforts seek to include user-labeled examples. Blanchard et al. [2] propose a semi-supervised technique that trains a classifier using two kinds of data: labeled data known to be normal and an additional unlabeled sample that may contain anomalous data. Both approaches aim to train a binary classifier that labels new items as either “normal” or “anomalous.” The Blanchard technique further accommodates an upper limit on the false anomaly detection rate. Our approach differs from these methods in that it specifically incorporates known examples of false alarms to further improve the system’s precision. In addition, our system is designed for online operation rather than batch processing of previously collected data.

In contrast with statistical novelty detection, radio astronomers generally use physical models of the anticipated events. If the precise shape of the event is known in advance, matched filtering provides maximum sensitivity to detect faint transient pulses. These models reflect the fact that signals from remote astronomical sources are *dispersed*. As the signal travels through the interstellar medium that lies between the source and the observer, it encounters free electrons that absorb some of the signal’s energy and delay its propagation. This affects lower frequency components more than higher frequency components. The slight difference accumulates over long distances and ultimately causes a broadband signal to appear dispersed in time, so that the lower frequency components arrive later.

Real-time transient detection typically uses *incoherent* analysis which represents the data as a matrix of signal powers channelized into discrete time and frequency bins. The data is typically portrayed as a two-dimensional image in which the axes correspond to time and frequency. The pixel intensity shows observed power, the accumulated squared voltage received by the antenna. Figure 1 (left) shows a pulse from pulsar J0742-2822 that displays a typical dispersed “sweep.”

Dispersion manifests as a time delay Δt_{delay} that is inversely proportional to the signal’s frequency. Following [15]:

$$(1) \quad \Delta t_{\text{delay}} = 4.1\text{ms} \frac{DM k}{\Delta \nu_{\text{GHz}}^2}$$

Here $\Delta \nu$ is the difference between the frequency of the reference channel and the delayed channel. The amount of dispersion, or the Dispersion Measure (DM), correlates with the number of interfering electrons present between the source and the observer [1]. It is commonly reported in parsecs per cm^3 . For regions of constant electron density, the amount of dispersion suggests the physical distance to the source.




FIGURE 1. Examples of typical and atypical transient signals. The image at left shows a single pulse from pulsar J0742-2822, with a classic dispersed pulse profile. Such signals can be found by inverting the dispersion effect prior to matched filtering. More exotic and poorly understood phenomena, like the *peryton* signal pictured at right, do not match typical dispersion and could benefit from model-free detection strategies with fewer assumptions. This example shows a distinctive “kink” in the curved signal. The narrow horizontal lines are narrow-band interference; such behavior is time-variable but not astronomically relevant and would ideally not affect the detection decision.

Current detectors for remote transient signals are typically tailored to the known properties of dispersion. Data is exhaustively *dedispersed* using a variety of different candidate DMs [1, 4]. Dedispersion provides a detection statistic by summing across all frequency channels, applying an appropriate temporal shift at each frequency to undo the time delay of a given assumed DM. This tailored summation is equivalent to a matched filter, and increases detection sensitivity over a naive sliding window detection using all frequency channels. By seeking the maximum dedispersed sum across many potential DMs, one can characterize the signal (and roughly the distance to the source). A dedispersion search can also help separate genuine astronomical signals from Radio Frequency Interference (RFI). Broadband RFI manifests as a vertical signal with no dedispersion ($DM = 0$); the pulse originates locally and all frequencies arrive simultaneously.

This approach has proven effective for the detection of pulsars and other astronomical phenomena [7, 12, 6, 13]. It can be implemented efficiently to keep up with streaming data using FPGAs, GPUs, or other parallel architectures; dedispersion over multiple DMs is inherently highly parallelizable. The weakness of this approach, however, is that it is sensitive to only one kind of signal. While dispersion is a known phenomenon of all remote signals, some recently-discovered sources (Figure 1, right) exhibit deviations from the expected shape which renders them difficult to detect. Further, it is not known how many other exotic source types may currently be overlooked due to the detection method’s dependence on one kind of signal model. The next section presents a more flexible strategy

that could operate in parallel with dedispersion searches, providing the capability to detect both dispersed pulses and unanticipated novel events.

3. APPROACH

We describe a new approach that combines 1) prior knowledge about uninteresting signals with 2) online estimation of the current data properties to enable flexible detection of novel signals. We treat the data as a sequence of observations that arrive sequentially from the antenna. We combine n such observed data points $\mathbf{x}_i \in \mathbb{R}^d$ as columns of a $d \times n$ data matrix $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n]$. Here, d is the number of frequency channels observed at each time step. The goal is to compute a discriminant function that maps each observation to a novelty score, $f(\mathbf{x}_i) : \mathbb{R}^d \mapsto \mathbb{R}$. The discriminant value should be small for typical data but large for interesting or novel data.

3.1. Constructing an eigenbasis. We use the popular strategy of measuring the distance from the signal to a low-dimensional manifold learned from the data stream [21, 9]. We will start by describing the simpler case of novelty detection in a static (non-adaptive) subspace. We hypothesize that the “regular” data lies on a linear subspace in \mathbb{R}^d with $d' \ll d$. Subtracting the data mean $\bar{\mathbf{x}}$ yields a translated matrix $\tilde{\mathbf{X}} = [(\mathbf{x}_1 - \bar{\mathbf{x}}), (\mathbf{x}_2 - \bar{\mathbf{x}}), \dots, (\mathbf{x}_n - \bar{\mathbf{x}})]$. Singular Value Decomposition (SVD) provides $\tilde{\mathbf{X}} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$. The columns of \mathbf{U} are the *principal components*: an orthonormal basis with axes in the order of decreasing data variance. We form a low-dimensional basis \mathbf{A} using the first d' columns of \mathbf{U} . One can also compute the matrix \mathbf{A} via classical Principal Component Analysis (PCA), e.g., using the eigenvectors corresponding to the largest eigenvalues of the covariance matrix $\tilde{\mathbf{X}}\tilde{\mathbf{X}}^T$.

We quantify the novelty of observation \mathbf{x}_i using the Euclidean distance to the subspace, equivalent to the L2-norm *reconstruction error* after first transforming \mathbf{x}_i into the low-dimensional basis and then reconstructing an approximation $\hat{\mathbf{x}}_i$. This leads to the following discriminant function which is large for novel data and zero for points on the linear manifold.

$$(2) \quad f(\mathbf{x}_i) = \|\mathbf{x}_i - \hat{\mathbf{x}}_i\| = \|(\mathbf{x}_i - \bar{\mathbf{x}}) - \mathbf{A}\mathbf{A}^T(\mathbf{x}_i - \bar{\mathbf{x}})\|_2$$

The eigenvalue decomposition makes computing \mathbf{A} difficult for large n . However, it is important that our basis accommodate large data sets and long-timescale changes in the background. One solution is to periodically recompute the entire matrix \mathbf{A} in batch mode using a recent subset of the data. In this work we employ the online approach of Ross et al. [14, 19] for efficient online updates to the mean $\bar{\mathbf{x}}$ and eigenbasis \mathbf{A} . This approach updates an SVD decomposition defined by some previous data $\tilde{\mathbf{X}}_p = \mathbf{U}_p\mathbf{\Sigma}_p\mathbf{V}_p^T$. Each block update has a data matrix \mathbf{X}_q with mean $\bar{\mathbf{x}}_q$ and decomposition $\tilde{\mathbf{X}}_q = \mathbf{U}_q\mathbf{\Sigma}_q\mathbf{V}_q^T$. This gives a combined dataset $\mathbf{X}_r = [\mathbf{X}_p|\mathbf{X}_q]$. Fortunately one can compute an updated mean $\bar{\mathbf{x}}_r$ and eigenbasis $\tilde{\mathbf{X}}_r = \mathbf{U}_r\mathbf{\Sigma}_r\mathbf{V}_r^T$ without having to store the old data explicitly. We refer the reader to the original work [14] for details, but summarize their approach in Algorithm 1. It relies on the widely-studied R-SVD procedure [8] which exploits the fact that a low-rank update to the eigenbasis is decomposable into efficient block operations. The method extends R-SVD to the case where the data are not assumed to have zero mean.

Algorithm 1: Ross et al. Algorithm for Sequential Eigenbasis Updates.

Input: Previous mean $\bar{\mathbf{x}}_p$

Previous decomposition $\mathbf{U}_p\mathbf{\Sigma}_p\mathbf{V}_p^T$

Additional data \mathbf{X}_q

Output: Revised mean $\bar{\mathbf{x}}_r$

Revised decomposition $\mathbf{U}_r\mathbf{\Sigma}_r\mathbf{V}_r^T$

Compute $\bar{\mathbf{x}}_r = \frac{n}{n+m}\bar{\mathbf{x}}_p + \frac{m}{n+m}\bar{\mathbf{x}}_q$, where $n = |\mathbf{X}_p|$ and $m = |\mathbf{X}_q|$

Compute $\mathbf{E} = \left(\mathbf{X}_q - \bar{\mathbf{x}}_r\mathbf{1}_{(1 \times m)}\right)\sqrt{\frac{nm}{n+m}}(\bar{\mathbf{x}}_p - \bar{\mathbf{x}}_q)$

Use $\mathbf{U}_p\mathbf{\Sigma}_p\mathbf{V}_p^T$ with \mathbf{E} as input to R-SVD to compute $\mathbf{U}_r\mathbf{\Sigma}_r\mathbf{V}_r^T$

An advantage of the Ross et al. method is that one can downweight the old basis to introduce a forgetting factor that allows the influence of old data to decay gradually as new points are added. This lets the basis shift to track a nonstationary distribution, and it accommodates observations of arbitrary length.

3.2. Semi-supervised eigenbases. Automated novelty detection should exclude rare events that are known to be uninteresting. In particular, one might anticipate specific *false alarms* due to instrument noise, interference, or other mundane but intermittent phenomena. Alternatively, a human user could provide feedback on previous detections that turned out to be uninteresting. We incorporate information about these known false alarms with a second basis trained to model rare or anomalous, but uninteresting, patterns. Our semi-supervised novelty detection method uses a combined subspace with both supervised and unsupervised components. It therefore adapts to long-term background trends while still excluding known false alarms. Algorithm 2 summarizes Semi-Supervised Eigenbasis Novelty Detection (SSEND).

Algorithm 2: Semi-Supervised Eigenbasis Novelty Detection (SSEND)

Input: False alarm data \mathbf{X}_f of size l datapoints,
Context timesteps \mathbf{X}_c of size m datapoints f for each false alarm
Size m block updates of streaming, unsupervised data \mathbf{X}_u
Output: Novelty scores $f(\mathbf{x}_1), f(\mathbf{x}_2), \dots$ for \mathbf{X}_u

Offline

For each false alarm timestep \mathbf{x}_f :

Concatenate m zero-meaned context timesteps to yield $\tilde{\mathbf{X}}_c = [\mathbf{x}_{c1} - \bar{\mathbf{x}}_c, \dots, \mathbf{x}_{cm} - \bar{\mathbf{x}}_c]$
Use PCA with $\tilde{\mathbf{X}}_c \tilde{\mathbf{X}}_c^T$ or SVD with $\tilde{\mathbf{X}}_c = \mathbf{U}_c \Sigma_c \mathbf{V}_c^T$ to compute an orthonormal basis \mathbf{U}_c
Compute the residual $\mathbf{x}_s = \mathbf{x}_f - \bar{\mathbf{x}}_c + \mathbf{U}_c \mathbf{U}_c^T (\mathbf{x}_f - \bar{\mathbf{x}}_c)$
Append the residual \mathbf{x}_s to the supervised training data set \mathbf{X}_s

Compute $\tilde{\mathbf{X}}_s = [\mathbf{x}_{s1} - \bar{\mathbf{x}}_s, \dots, \mathbf{x}_{sl} - \bar{\mathbf{x}}_s]$

Use PCA or SVD with $\tilde{\mathbf{X}}_s$ to compute the supervised, orthonormal basis \mathbf{U}_s

Online

Using the first block \mathbf{X}_{u1} , compute an initial mean $\bar{\mathbf{x}}_p$ and eigenbasis $\mathbf{U}_p \Sigma_p \mathbf{V}_p^T$

For each subsequent \mathbf{X}_{uj} :

Compute a revised mean $\bar{\mathbf{x}}_r$ and a revised decomposition $\mathbf{U}_r \Sigma_r \mathbf{V}_r^T$ using the Ross et al. method

Define a combined basis $\mathbf{U}_c = [\mathbf{U}_r | \mathbf{U}_s]$

Use QR decomposition to find a semi-supervised basis \mathbf{U}_{ss} that makes \mathbf{U}_c orthonormal

For each $\mathbf{x}_i \in \mathbf{X}_{uj}$:

Compute $f(\mathbf{x}_i) = \|(\mathbf{x}_i - \bar{\mathbf{x}}_r) - \mathbf{U}_{ss} \mathbf{U}_{ss}^T (\mathbf{x}_i - \bar{\mathbf{x}}_r)\|_2$

SSEND has both offline and real time elements (shown visually in Figure 2). Offline, we accumulate the false alarm set of momentary events known to be uninteresting. These timesteps will also contain incidental background noise that is unrelated to the event itself, so we must take additional steps to isolate the false alarm pattern from its local context. We construct a data matrix \mathbf{X}_c using the timesteps directly prior to each event, and model them with a context basis \mathbf{U}_c using a procedure like PCA. We then project each false alarm onto its context basis, leaving a residual value that is the unique signal of that event (the part which is distinct from its local background).

$$(3) \quad \mathbf{x}_s = \mathbf{x}_f - \bar{\mathbf{x}}_c + \mathbf{U}_c \mathbf{U}_c^T (\mathbf{x}_f - \bar{\mathbf{x}}_c)$$

We concatenate the residuals into a dataset and train a supervised basis \mathbf{U}_s . This is our false alarm model. At runtime, we compute an adaptive mean $\bar{\mathbf{x}}_r$ and basis \mathbf{U}_r using the Ross et al. method as before, and define a combined semi-supervised $\mathbf{U}_{ss} = [\mathbf{U}_r | \mathbf{U}_s]$ to span both the supervised and the unsupervised data. We orthogonalize the new basis using QR decomposition with the Gram-Schmidt

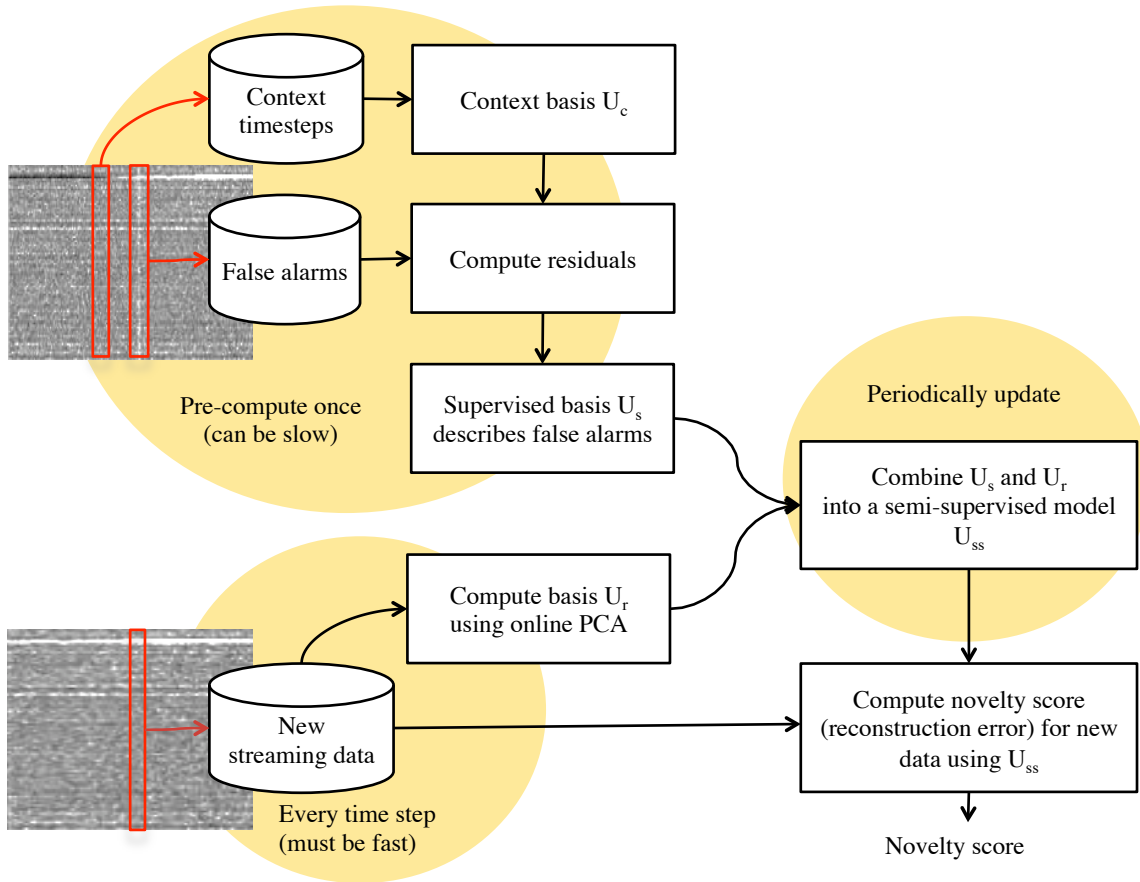


FIGURE 2. Semi-Supervised Adaptive Novelty Detection concept

method. The reconstruction error with respect to the combined model yields a more reliable novelty score.

Note that the proposed approach does not preserve the mean of the initial false alarm distribution, which is assumed to drift in a similar fashion as the mean of the online distribution. User feedback would permit a more sophisticated system that also updates the false alarm mean and basis online, but we focus here on the simpler case where all training occurs in advance.

The specific semi-supervised approach here is one of a broader family of methods, where the bases to be combined might be learned through many alternative techniques. In particular, there are many options for the supervised portion which occurs offline and does not have a real-time computational constraint. For centered data $\tilde{\mathbf{X}}_s$, classical PCA is tantamount to identifying components \mathbf{z}^* which maximize the magnitude of the projection onto the data covariance matrix:

$$(4) \quad \mathbf{z}^* = \arg \max_{\mathbf{z}} \mathbf{z}^T (\tilde{\mathbf{X}}_s^T \tilde{\mathbf{X}}_s) \mathbf{z} \quad s.t. \quad \mathbf{z}^T \mathbf{z} \leq 1$$

In this paper, we introduce a further innovation designed to improve the interpretability of the learned model. We replace the PCA step with a sparse PCA formulation [11]. Specifically, we incorporate an $L1$ norm penalty on the components.

$$(5) \quad \mathbf{z}^* = \arg \max_{\mathbf{z}} \mathbf{z}^T (\tilde{\mathbf{X}}_s^T \tilde{\mathbf{X}}_s) \mathbf{z} - \lambda \|\mathbf{z}\|_1 \quad s.t. \quad \mathbf{z}^T \mathbf{z} \leq 1$$

This has the effect of driving basis components to zero, with the λ parameter balancing variance maximization and sparsity objectives. These sparse codebooks may benefit interpretability, and can improve generalization performance where the physical processes are themselves known to be sparse.

4. EVALUATION

SSEND was motivated by applications in radio astronomy. We compared performance on a test set of radio array data using five linear and nonlinear novelty detection algorithms: the traditional dedispersion approach, kernel PCA novelty detection [9], one-class SVM novelty detection [20], unsupervised adaptive novelty detection using PCA, and the proposed semi-supervised approach.

4.1. Data set. We use a selected portion of data from the Parkes Multibeam Survey, an extensive search for Pulsars using the Parkes radio telescope of CSIRO [7, 16, 10]. This instrument views the sky simultaneously through 13 receivers, effectively providing 13 independent antennas covering adjacent, and slightly overlapping, areas in the sky. Receiver measurements are recorded at high time resolution and transformed into channelized power measurements corresponding to the squared voltage response at various discrete frequency channels. This specific data sequence contains examples of events known as *peryttons*, first discovered by Burke-Spolar and Bailes in their analysis of Parkes pulsar surveys [3]. Perytons are still poorly understood, and they are scientifically interesting because they vary in frequency and approximate a dispersion curve. However, they do not exactly match a dispersion profile, and their spatial distribution in the sky suggests that they are of terrestrial (possibly atmospheric) origin.

In addition to these features, structured interference is often visible in the form of channel-specific noise and gain fluctuations appearing as horizontal stripes. Such noise is pervasive and typical for highly sensitive, cryogenically cooled receiver feeds. Our tests focus on approximately five minutes of observation time in each of the 13 receivers. This span includes several tens of thousands of timesteps recorded at a cadence of 0.125 milliseconds in each of 96 frequency channels near 1450 MHz. Figure 3 shows three examples of perytons. The red rectangle shows the size of an example data window used to construct \mathbf{x}_i .




FIGURE 3. True anomalies: Peryton events from the Parkes Multibeam Survey.

Figure 4 shows some examples of false alarms that are statistically uncommon but not scientifically interesting. These specific examples are broadband pulses of Radio Frequency Interference (RFI), probably emitted by some local artificial source. Such features are rare enough that they are not well-represented in an unsupervised eigenbasis, but typical enough that they would dominate novelty detection results if not handled explicitly.

4.2. Methodology. We average the data by a factor of 20 down to a temporal resolution of 2.5 ms, and then create a data set from a sequence of short non-overlapping *segments* that cross all 96 vertical frequency channels and 6 horizontal timesteps. This segment width corresponds to a 15ms time interval, found to maximize performance across all methods. We reorder each segment into a single column vector $\mathbf{x} \in \mathbb{R}^{384}$. Finally, we unify data from all beams into one large dataset,




FIGURE 4. False anomalies: Vertical stripes due to broadband RFI that are statistically anomalous but uninteresting.

withholding five beams (38%) for training purposes. These tests consider the proposed SSEND method which combines supervised and unsupervised components and reports reconstruction error $f_{ss}(\mathbf{x}_i)$. Here we trained the subspace \mathbf{U}_s using 30 overlapping segments (\mathbf{X}_s) drawn from three manually-selected broadband RFI pulses. We show results using both the original (dense) solutions and the sparse PCA variation. For comparison, we also report results from the original SSEND version published in earlier work that does subtract the local context of false alarm training events [22].

For comparison, we also consider some alternatives that are broadly representative of popular linear and nonlinear anomaly detectors. First is a purely unsupervised eigenbasis approach based on reconstruction error from a low-dimensional basis $f_u(\mathbf{x}_i)$. It does not explicitly account for RFI. Another popular unsupervised method is the one-class SVM novelty detection of Scholkopf et al. [20]. Here we use a radial basis kernel function selected with a grid search, and treat each test point’s distance to the decision boundary as a real-valued novelty score. A nonlinear, unsupervised version is kernel PCA: a non-linear extension of PCA. Kernel PCA novelty detection first maps the data to a higher (generally infinite) dimensional features space, computes the principal components in this space, projects the transformed data to a lower-dimension manifold, and defines a novelty measure as the reconstruction error in the feature-space. Kernel functions allow the reconstruction error to be calculated without explicitly mapping to the feature space [9]. However, this method never explicitly calculates the principal components so it cannot be used as an adaptive technique in the manner discussed in Algorithm 2. Instead we use the implementation of Hoffmann et al. [9]. We use a radial basis kernel function with parameters selected by a grid search.

Finally, we consider a state-of-the-art radio astronomy solution that uses incoherent dedispersion and summation to search DM values from 0 to 500. We correct each time step separately for each DM, and use the maximum response from all DMs as the novelty score $f_d(\mathbf{x}_i)$. Time averaging did not improve performance for the dedispersion so we report the dedispersion approaches’ single-timestep result. We used the 15ms window for the other methods, which we found to give the best overall performance. We used equivalent preprocessing in all trials.

We identified the precise locations of all peryton events (desired detections) noted in the study by Burke-Spolaor et al. [3]. These appeared to some degree in all antennas, although the signal strength and character varied somewhat even for simultaneous observations. The concatenated dataset provided 88 real novel events for our evaluation. We assigned each peryton an enclosing time interval; any detection in this range counted as having successfully detected the peryton. We take the peryton to be present in all beams even though the actual signal strength varies across receivers. This doesn’t matter for our performance comparison since weak signals penalize all detection methods equally.

We evaluated each method by first computing novelty scores for the entire data set, sorting these scores across all beams, and then counting the result of each trigger in order of decreasing novelty. Each peryton can only be captured once, though multiple triggers within the same event do not

count as false positives. However, any detection falling outside a peryton interval counts as a false positive.

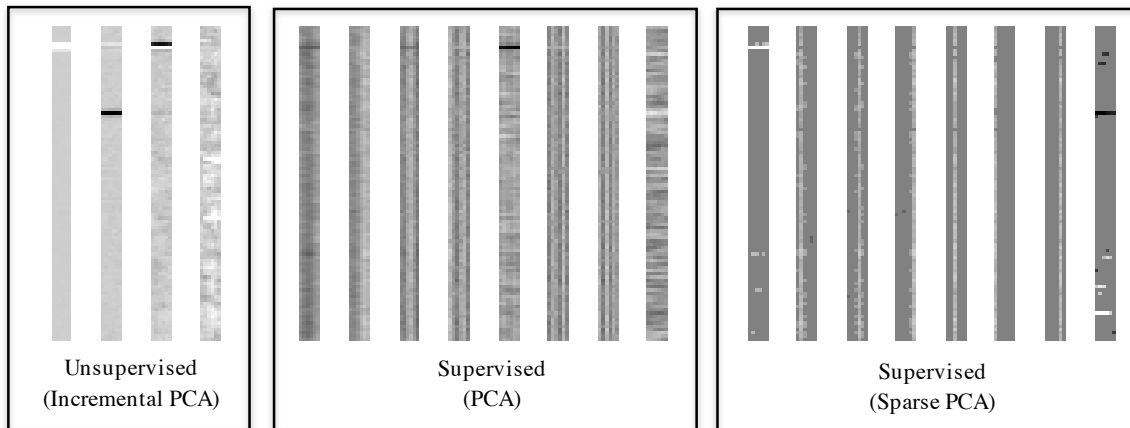


FIGURE 5. Orthonormal principal components used to construct \mathbf{U}_{ss} from \mathbf{U}_r and \mathbf{U}_s (dense or sparse). The unsupervised portion (left) models channelized interference, while the vertical structures in the supervised portion represent momentary broadband RFI.

4.3. Results. A visualization of the unsupervised and supervised bases learned by our method appears in Figure 5. Here we use the top 4 principal components as an unsupervised basis with online updates from the data stream. These *eigensignals* (Figure 5, left) show high magnitude in the most variable channels; at the time this eigensignal snapshot was captured, such channels comprised the major axis of variance for the data set. A supervised basis of 8 dimensions models the known broadband RFI. We show the top eigensignals for both classical and sparse methods in Figure 5 center and right, respectively. Both models capture the vertical profile of momentary RFI pulses at different locations. The sparse basis (right) is clearly interpretable as a combination of short additive broadband components. Paired with either supervised option, the online PCA basis can accurately reconstruct a slow shift in channelized RFI conditions along with any additive RFI pulse. Note that this image shows the orthonormal segments after QR factorization.

Figure 6 compares novelty detection scores for the entire observation sequence of the first test beam, computed with a purely unsupervised basis (standard PCA), f_u , and the semi-supervised approach, f_{ss} . Interesting peryton events are noted by black triangles; the other signal spikes correspond to various kinds of RFI. Five peryton signals were barely visible in the reconstruction error of either method, due possibly to the alignment of non-overlapping segments or the inherently weak visibility of those events in this beam. We exclude these five from the diagram for clarity. In general SSEND responds to the novelty of peryton events while filtering most of the RFI. In contrast, broadband RFI contaminates the purely unsupervised approach; it accounts for the three strongest responses by f_u for this sequence.

Figure 7 shows a Receiver Operating Characteristic (ROC) curve describing the tradeoff in precision and recall rates. We report the number of perytons captured for a variety of false positive budgets, considering the semi-supervised approach as well as the sparse semi-supervised variant which uses sparse PCA for supervised learning stage. False positive budgets beyond 10 are excessive since this would represent greater than one detection event for every 5 seconds of observations (an unrealistic burden on manual post-analysis). Future commensal campaigns with constant observations and higher data volumes will demand even stricter limits. For this low error budget, SSEND considerably outperforms the competing methods: the top 13 signals detected via f_{ss} are due to




FIGURE 6. Semi-supervised learning filters out RFI events that would otherwise dominate the detection results. This time series plot shows per-timestep novelty evaluated for the first beam in the test set. Not all perytons are clearly distinguishable in this beam.

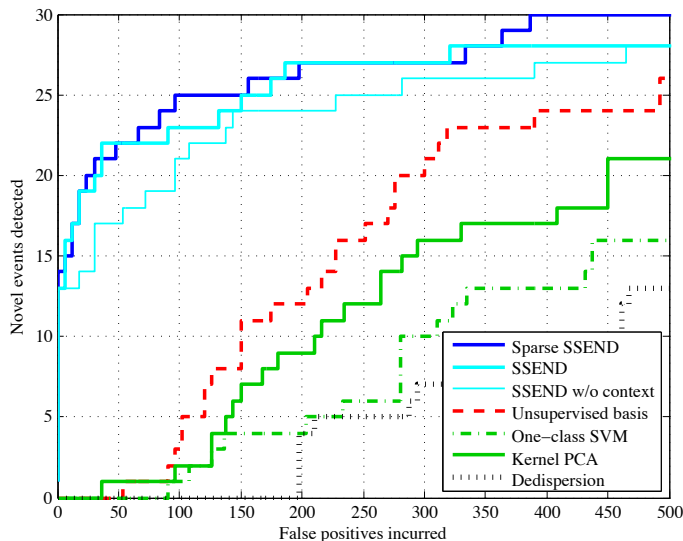


FIGURE 7. ROC curves comparing eigenbasis novelty detection approaches with the traditional dedispersion search.

perlytons, while the kernel PCA technique detects ~ 30 false-positives before the first peryton, and the unsupervised method reports more than 50 false positives before finding the first real peryton. These runner-up methods require 250 and 200 false positives respectively before they match the error-free retrieval rate of the semi-supervised approaches.

Notably, both dense and sparse SSEND offer comparable performance. For completeness we also report performance of the original SSEND algorithm first reported in [22] which does not consider

the local context around false alarms. Moving to the version reported here produces a slight, but perceptible, improvement. Additional hand-tuned RFI excision rules, such as a ban on zero-DM signals that are likely to be terrestrial, might improve performance further. Naturally, such rules are less general than a purely learning-based approach and might filter other unanticipated anomalies.




FIGURE 8. ROC curves to assess the sensitivity to data segment sizes.

The preceding results form \mathbf{x}_i with a data segment of 15 ms (6 time steps). We evaluated sensitivity to segment length (see Figure 8). Segments of duration 10 – 15 ms performed best for this data set. It is possible that smaller segments are susceptible to noise while larger sizes dilute the perytons. It might improve performance for large segments to use a higher-dimensional basis for the unsupervised component. Such models might do a better job of modeling temporal structure (such as switching interference) that begins to appear at these scales.

We also assessed the runtime of each method to determine whether they could be used in a realistic real-time setting. Using a single core of a modern desktop processor, the runtime of the dedispersion search method averaged 0.16 seconds per DM for the entire subsampled sequence, or ≈ 80 seconds for a typical search over 500 DM values. This could be divided easily among multiple processors to provide faster processing for multiple beams. The eigenbasis approaches’ runtimes depend strongly on the size of the block updates to the eigenbasis. For a single desktop processor core performing block updates of size $m = 100$, the entire observation from a single beam was processed at $5\times$ real time (≈ 10 seconds/beam for the entire dataset). The time required was slightly larger (up to ≈ 20 seconds/beam) for smaller block updates where constant-time overhead costs had a larger impact. The accuracy of these techniques was nearly indistinguishable for all block update sizes we tried. Varying the segment sizes also affected run time by up to a factor of two. Kernel PCA and one-class SVM performed considerably slower than the dedispersion and eigenbasis approaches as all computations were performed with an RBF kernel representation of the data: a representation of size $|\mathbf{x}|^2 = 384^2$ for this work. In our experiments we found these techniques required $\approx 200 - 400$ seconds/beam with block updates of $m = 400$. Furthermore, unlike the dedispersion and eigenbasis techniques, the Kernel PCA and one-class SVM computation times scale quadratically with the size of m . This reduces the generality of these methods, and in combination with their large computational run-times, makes them infeasible as real-time techniques. On the other hand, we found the dedispersion and eigenbases approaches to be readily employable for real-time use on general purpose computing hardware.

5. DISCUSSION

Semi-Supervised Eigenbasis Novelty Detection applies to anomaly detection in domains with real-time requirements, high-dimensional input, and prior knowledge about false alarm events. Of course, it is not necessary to incorporate false alarm information directly into the novelty detection model as we have done here. Instead a pre-classification could filter these events prior to a purely unsupervised novelty detection stage. Nevertheless, there may be other advantages to the combined approach of SSEND. It is simple and easy to implement. The projection shifts to reflect any underlying drift in the mean signal levels, so that a basis trained on previous false alarms remains relevant. Further work will investigate ways to combine multi-scale models when the temporal extent of the interesting events is not known in advance. Finally, application to the broader Parkes survey catalogue will increase practical experience with the technique, and may even reveal additional classes of RFI and astronomical transient events.

6. ACKNOWLEDGEMENTS

This work was made possible through the assistance of Sarah Burke-Spolaor, formerly of CSIRO and now affiliated with the Jet Propulsion Laboratory. More generally we would also thank Swinburne University and CSIRO for a generous data access policy. We also thank J-P Macquart (CSIRO), and Dayton Jones, Robert Preston and Joseph Lazio (Jet Propulsion Laboratory) whose advice and ideas have influenced our investigation. This research was performed at the Jet Propulsion Laboratory, California Institute of Technology, under the Research and Technology Development Strategic Initiative Program, under a contract with the National Aeronautics and Space Administration. Copyright 2012 California Institute of Technology. All Rights Reserved. U.S. Government Support Acknowledged.

REFERENCES

- [1] D. Bhattacharya. Detection of radio emission from pulsars. *NATO ASIC Proc. 515: The Many Faces of Neutron Stars*, 1998.
- [2] G. Blanchard, G. Lee, and C. Scott. Semi-supervised novelty detection. *Journal of Machine Learning Research*, 11:2973–3009, 2010.
- [3] S. Burke-Spolaor, M. Bailes, R. Ekers, J. Macquart, and F. Crawford III. Radio Bursts with Extragalactic Spectral Characteristics Show Terrestrial Origins. *The Astrophysical Journal*, 727:18, 2011.
- [4] J. Chennamangalam et al. Software data-processing pipeline for transient detection. *The Low-Frequency Radio Universe, ASP. Conf. Series*, LFRU, 2009.
- [5] J. M. Cordes et al. The dynamic radio sky. *New Astronomy Reviews*, 48:1459–1472, 2004.
- [6] J. M. Cordes and M. A. McLaughlin. Searches for fast radio transients. *The Astrophysical Journal*, 596:1142–1154, October 2003.
- [7] R. T. Edwards, M. Bailes, W. van Straten, and M. Britton. The Swinburne Intermediate Latitude Pulsar Survey. *MNRAS*, 326:358–374, 2001.
- [8] G. Golub and C. Van Loan. *Matrix Computations*. Johns Hopkins Univ Pr, 1996.
- [9] H. Hoffmann. Kernel PCA for novelty detection. *Pattern Recognition*, 40(3):863–874, 2007.
- [10] B. A. Jacoby, M. Bailes, S. M. Ord, R. T. Edwards, and S. R. Kulkarni. A large-area survey for radio pulsars at high galactic latitudes. *The Astrophysical Journal*, 699:401–411, 2009.
- [11] M. Journé, Y. Nesterov, P. Richtárik, and R. Sepulchre. Generalized power method for sparse principal component analysis. *J. Mach. Learn. Res.*, 11:517–553, March 2010.
- [12] E. Keane. The search for nearby RRATs and other transient radio bursts. In *Third Estrela Workshop*, 2008.
- [13] J. Lazio, J. S. Bloom, G. C. Bower, J. Cordes, S. Croft, S. Hyman, C. Law, and M. McLaughlin. The dynamic radio sky: An opportunity for discovery. *Astro2010: The Astronomy and Astrophysics Decadal Survey White Paper no. 176*, 2009.
- [14] J. Lim, D. Ross, R. Lin, and M. Yang. Incremental learning for visual tracking. *Advances in Neural Information Processing Systems*, 1:793–800, 2005.
- [15] A. G. Lyne and F. G. Graham-Smith. *Pulsar Astronomy*. Cambridge University Press, 1998.

- [16] R. N. Manchester et al. Parkes multibeam pulsar survey: I. observing and data analysis systems, discovery and timing of 100 pulsars. *MNRAS*, 328:17–35, 2001.
- [17] M. Markou and S. Singh. Novelty detection: a review. Part 1: statistical approaches. *Signal Processing*, 83(12):2481–2497, 2003.
- [18] M. Markou and S. Singh. Novelty detection: a review. Part 2: neural network based approaches. *Signal Processing*, 83(12):2499–2521, 2003.
- [19] D. Ross, J. Lim, R.-S. Lin, and M.-H. Yang. Incremental learning for robust visual tracking. *International Journal of Computer Vision*, 77:125–141, 2008. 10.1007/s11263-007-0075-7.
- [20] B. Schölkopf, J. Platt, J. Shawe-Taylor, A. Smola, and R. Williamson. Estimating the support of a high-dimensional distribution. *Neural Computation*, 13(7):1443–1472, 2001.
- [21] S. O. Song, D. Shin, and E. S. Yoon. Analysis of novelty detection properties of auto-associators. In *Proceedings of COMADEM*, pages 577–584, 2001.
- [22] D. R. Thompson, W. A. Majid, C. J. Reed, and K. L. Wagstaff. Semi-supervised novelty detection with adaptive eigenbases, and application to radio transients. In *Proceedings of the NASA Conference on Intelligent Data Understanding*, 2011.