
Robust Detection of Radiation Threat

Eric Lei

12 April 2017

Abstract

Background. We consider the problem of radiological source detection in the presence of noisy background environments. The objective is to detect dangerous materials sources by their unique radioactive signatures. The foundations of this problem are presented here.

Related work. There has been much research on statistical and signal processing techniques for source detection. These methods mostly concentrate on anomaly detection when one has little information about the target source. The most popular class of methods involves Principal Components Analysis to learn the expected characteristics of background. There are also many detection methods that do not involve Principal Components Analysis.

Existing methods. We describe PCA anomaly detection, Censored Energy Window, Matched Filter, and Gaussian-Poisson MAP estimation in detail.

Uncertain Censored Energy Windows. A method is proposed for a setting in which information about threat templates is incomplete.

Source independence experiment. An experiment is performed that simulates incomplete knowledge of the source template.

Simultaneous estimation of source strength and background. Another method is proposed that has much lesser dependence on training data.

Generalization experiment. An experiment is conducted by simulating test data that greatly differ from training data.

Conclusion. Both new methods are empirically demonstrated to perform better than competitors.

Keywords: anomaly detection, gamma-ray spectroscopy, radiation monitoring, canonical correlation, Kalman filter

Word Count: 197 words

1 Introduction

In nuclear physics, an important problem is the detection of threatening sources (Runkle et al., 2006). We consider a mobile detector used for search, such as a vehicle that travels through an urban environment to locate dangerous radioactive materials. The primary challenge in this domain is a low signal-to-noise ratio (Tandon, 2016). Urban environments contain large numbers of nuisance (benign) sources—potentially every material present—which can heavily obfuscate the signal from the target source (Aucott et al., 2014). These nuisance sources may include ordinary objects like concrete or bananas as well as shielding materials purposely used to hide the source. Moreover, since the environment changes as the detector moves, the background radiation fluctuates widely, further compounding this problem (Aucott et al., 2014). To address this challenge, the types of research performed in this area have been twofold. The first type concerns the development of physical sensors to detect radioactive particles such as in Ziock et al. (2004), Aucott et al. (2015), and Miller et al. (2016). The second type employs statistical and signal processing techniques to analyze data from sensors. In this study we focus only on the second type.

Several statistical and signal processing methods for spectral data have found success in source detection. However, some of these references only appear in the form of internal technical reports. Others contain unpublished details in their algorithms. Therefore, this paper aims to survey several of these methods, provide detailed algorithmic descriptions, and document the statistical framework used for this problem. The methods discussed are PCA-based anomaly detection, Censored Energy Window, Matched Filter (Tandon, 2016), and Gaussian-Poisson MAP estimation. We describe the algorithms for computing each method. Then we compare and contrast the underlying motivations and informational assumptions. We also illustrate the differences in detection capability on a real radiation dataset.

In addition, this work investigates how two assumptions of the methods could be violated, significantly worsening detection performance. For each of the assumptions, we propose a novel method that performs well when it is violated. The first method addresses the assumption that information about the target source is complete. We simulate incomplete information by removing the source from a known source library that is utilized by the method, thereby modeling a situation in which a previously unknown source is encountered. The method is based on the Censored Energy Window and extends it with Canonical Correlation Analysis (Hotelling, 1936). The second method address the assumption that training data are available and match the test data. The method serves as a direct improvement to the Gaussian-Poisson method by exploiting the time series nature of the data. Specifically, it incorporates local background information through an adaptive Kalman filter. The filter simultaneously estimates local source strength and background shape. Both methods are applied to datasets and demonstrated to improve performance over competitors.



(a) Methods ranked on the source independence scale. (b) Methods ranked on the generalization scale.

Figure 1: Two scales illustrating the importance of different assumptions.

We rank detection methods on two separate theoretical scales. First, methods are sorted by their degree of independence from information about the target source. When this information is inaccurate or incomplete, some methods lose more performance than others. We denote this the source independence scale (Fig. 1a). Second, methods are sorted by their dependence on a training set of background measurements. Specifically, we consider their robustness when the test set comes

from a different distribution. We denote this the generalization scale (Fig. 1b). Each of the new methods ranks highly on one of the scales (shown in boldface).

This paper is structured as follows: §2 introduces necessary background; §3 reviews related work; §4 covers several existing methods in detail; §5 and §6 propose a method with less dependence on source information and describe an experiment with that method; §7 and §8 do the same with a method that depends less on training data; §9 concludes the paper.

2 Background

This section states the foundations of the source detection task. We consider the problem of radiological source detection in the presence of noisy background environments (Aucott et al., 2014). The objective is to detect dangerous materials sources by their unique radioactive signatures. This task can be accomplished by moving a gamma-ray spectrometer through the environment of interest, such as a city, and collecting radiation samples. A gamma-ray spectrometer counts photons with which it makes contact. Each particle has its own frequency, recorded by the spectrometer. Over an interval of time, particles can be counted in each range of frequencies, known as *energy bins*. The energy bins are typically linearly or quadratically spaced and cover the vast majority of usually observed frequencies. The counts in an interval of time over a spectrum of energy bins constitute one measurement, such as in Fig. 2. By recording a sequence of measurements, a spectrometer collects a time series of spectra. The source detection task is to score either an individual measurement or a set of measurements according to how likely a given source is present (Tandon, 2016). The classification is usually performed in real-time. Unlike common online learning scenarios, however, no feedback is observed.

Every material, benign or harmful, has an energy signature that can be defined as the rate over time of particles emitted in each energy bin. The count of particles in an interval of time and an energy bin can be viewed as a Poisson random variable with mean given by the energy signature (Bai et al., 2011). The Poisson counts can be sampled and added to background measurements to form injected (source-containing) measurements, as measurements containing actual threatening sources are otherwise difficult to obtain. An example of a source signature and injection is shown in Fig. 2. Observe how the injected measurement is significantly higher than the background in energy bins where the source is strongest. Furthermore, the strength of a source signature is called the *intensity* and defined as the sum of rates over all energy bins. The intensity is inversely proportional to the square of distance from the source (Miller et al., 2016). When the intensity is higher, a detection method should be more confident that a source is present.

Although the energy signature is known for most materials, this knowledge is not especially useful in a noisy environment because there are multiple materials at varying distances and the environment changes depending on the spectrometer’s position. Consequently, the aggregate signature of the

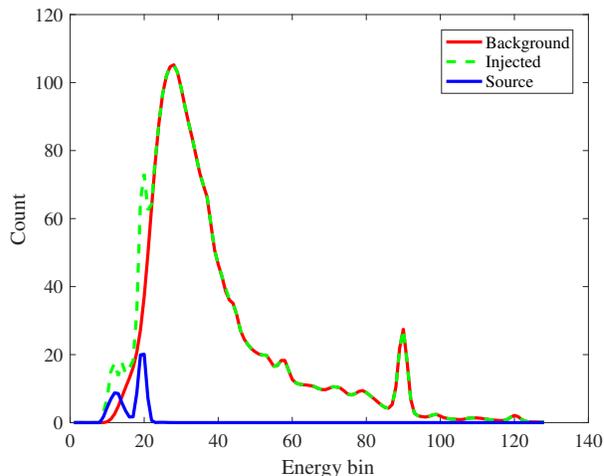


Figure 2: Example of a background sample, injected sample, and template.

background sources is dynamic and unknown, although it is often approximated as the mean of background spectra from a training set. The signature of the target source, on the other hand, can be assumed to be at least partially known by many detection methods (Nelson and Labov, 2012; Tandon, 2016). Source signatures can be considered known because of libraries constructed by domain experts: nuclear physicists have enumerated the templates of a wide array of possible source types with different configurations of fissile material and shielding (Nelson and Labov, 2012). Then detection methods that search for only one source type can be employed on every template in a library (Tandon, 2016). The information about a source template is utilized by specialized methods such as Matched Filter, Censored Energy Window, and Gaussian-Poisson MAP estimation, although the degree of assumed knowledge can still differ between them (Fig. 1a).

In some cases, the source template might be assumed to be unknown to reduce reliance on a source library. A common technique in this scenario is PCA-based anomaly detection (Tandon, 2016). However, the performance is much worse than source-aware methods, which leaves room for a method that utilizes more information but is robust to the completeness of the source library. In addition to unexpected source types, the library may overlook distortions of the source spectrum cause by scattering of occluding objects and intentional shielding. To fill this void, we propose a detection method based on Canonical Correlation Analysis (Hotelling, 1936). Additionally, one requirement shared by all methods listed above is a training set of background measurements. Typically some properties of the training background are learned, and when test measurements deviate from those properties, they are considered more likely to have a source. However, the training set could differ substantially from the test set, creating an opportunity to reduce dependence on the training set by updating the model with the test observations. To accomplish this, we propose an adaptive Kalman filter that simultaneously estimates local source strength and background spectrum.

3 Related work

There has been much research on statistical and signal processing techniques for source detection. These methods mostly concentrate on anomaly detection when one has little information about the target source. The most popular class of methods involves Principal Components Analysis to learn the expected characteristics of background. A study by Aage and Korsbech (2003) proposed noise-adjusted singular value decomposition and a method of “stripping” away common background features. Another technique based on Principal Components Analysis is presented by Runkle et al. (2006). A similar method was suggested by Du et al. (2010) using k -nearest neighbors in the principal components space of the background.

There are also many detection methods that do not involve Principal Components Analysis. In Anderson et al. (2008), an approach was presented that computes a ratio between observed and expected counts in key energy bins. This method may be viewed as a precursor to the Censored Energy Window (Tandon, 2016) method reviewed below. An improved version of the ratio method was given by Pfund et al. (2016). A method was developed by Bai et al. (2011) and Kump et al. (2013) that uses the variable selection capability of LASSO when fitting the observed spectra to their expected values. Other types of statistical methods were surveyed by Fagan et al. (2012), including peak finding and discriminant analysis. The authors point out that Bayesian methods were largely untapped at the time and could improve performance by incorporating all sources of uncertainty. Subsequently, a Bayesian approach by Tandon et al. (2016) demonstrated how to aggregate multiple detection methods to achieve better detection.

4 Existing methods

This section describes several existing methods for threat detection.

4.1 PCA anomaly detection

Unsupervised anomaly detection is useful when information is limited. While libraries of source templates exist, it may be important to limit a method’s dependence on them in case the information is false or unavailable. If no information about the source is available, then an approach based on Principal Components Analysis (PCA) may be used (Tandon, 2016). This approach utilizes a training set of background measurements. It computes the top principal components of the background and then scores new measurements by their reconstruction error. Conceptually, it learns the most important characteristics of the background and then subtracts them from the new measurements. If the remainder is small, then the measurement is adequately explained by background features. If the remainder is large, however, then the measurement is not explained by background, which indicates a separate source.

The algorithm is given in Alg. 1. The number k of principal components used is usually set to $k = 5$. Here, $\text{diag}(\cdot)$ on a square matrix returns the main diagonal as a vector, while the same operator on a vector returns a diagonal matrix with the vector as the main diagonal.

Algorithm 1 PCA anomaly detection

```
1: procedure TRAIN(data  $X \in \mathbb{R}^{n \times d}$ , number  $k$  of components)
2:    $X \leftarrow \text{FILTER}(X, [.3, .5, 1, 1, .5, .3])$ 
3:    $X \leftarrow \text{MOVINGAVERAGE}(X, 10)$ 
4:    $\mu \leftarrow \text{MEAN}(X)$ 
5:    $R \leftarrow \frac{1}{n}X^T X - .99\mu\mu^T$  // preserve 1% of mean
6:    $w \leftarrow \text{diag}(R)$ 
7:    $W \leftarrow \text{diag}(\sqrt{w + 1})$ 
8:    $V_k \leftarrow k$  leading eigenvectors of  $W^{-1}RW^{-1}$ 
9:    $T \leftarrow I - W^{-1}V_kV_k^T W$ 
10:  return  $T$ 
11: procedure TEST(data  $X \in \mathbb{R}^{n \times d}$ , trained  $T \in \mathbb{R}^{d \times d}$ )
12:   $X \leftarrow \text{FILTER}(X, [.3, .5, 1, 1, .5, .3])$ 
13:   $S \leftarrow \text{SUMSQUARES}(XT)/\text{SUM}(X)$  // element-wise division
14:  return  $S$ 
15: procedure FILTER( $X \in \mathbb{R}^{n \times d}$ , weights  $w \in \mathbb{R}^k$ )
16:  for  $j = 1, \dots, d$  do
17:    for  $i = 1, \dots, n$  do
18:       $Y_{ij} \leftarrow \sum_{t=0}^{k-1} w_{t+1}X_{i-t,j}$  where  $X_{i-t,j} = 0$  when  $i \leq t$ 
19:  return  $Y$ 
20: procedure MOVINGAVERAGE( $X \in \mathbb{R}^{n \times d}$ ,  $k \in \mathbb{R}$ )
21:  // uniform moving average with length  $k$  trailing window
22: procedure SUMSQUARES( $X$ )
23:  // computes vector of row-wise sums of squared elements of  $X$ 
24: procedure SUM( $X$ )
25:  // computes vector of row-wise sums of  $X$ 
```

4.2 Censored Energy Window

A benchmark method developed by Nelson and Labov (2012) is the Censored Energy Window (CEW), which assumes partial knowledge of the threat spectrum. It takes as input a set of energy bins in which the threat is expected to be seen most clearly, which is called the signal-to-noise ratio-maximizing energy window. Intuitively, CEW learns how many particles to expect within the energy window when there is no source. If an observation exceeds this expectation, then a source is more likely. The optimal energy window can be precomputed given a source template by a procedure shown in Alg. 2.

Algorithm 2 Energy window computation

- 1: **procedure** OPTIMIZE(source template $s \in \mathbb{R}^d$, mean background spectrum $\mu \in \mathbb{R}^d$)
 - 2: **for** $k = 1, \dots, d$ **do**
 - 3: $r_k \leftarrow s_k / \mu_k$
 - 4: **for** $k = 1, \dots, d$ **do**
 - 5: $\mathcal{C} \leftarrow$ set of bins with k highest r_j values
 - 6: $S_k \leftarrow \sum_{j \in \mathcal{C}} s_j$
 - 7: $B_k \leftarrow \sum_{j \in \mathcal{C}} \mu_j$
 - 8: $k^* \leftarrow \operatorname{argmax} S_k / \sqrt{B_k}$
 - 9: **return** set of bins with k^* highest r_j values
-

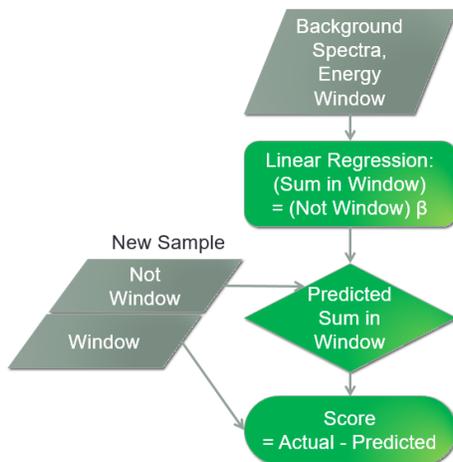


Figure 3: The CEW detection algorithm.

Given the energy window, CEW is trained on background measurements, fitting a multiple linear regression to predict the sum of counts inside the window from the spectrum of counts outside of it. The threat detection score is a difference between the predicted and the actually observed count in the window. The simple intuition behind this method is that if a source is present, then more counts will probably be observed inside but not outside the energy window. The steps of CEW, which are relatively straightforward, are displayed in Fig. 3. Note that the linear regression step may use ridge regression with a small regularization parameter to alleviate overfitting and collinearity. In addition, two steps are included before calculation of the score. First, after linear regression predicts the in-window sum of counts \hat{y} of a test observation, the prediction is mapped to $\tilde{y} = \sqrt{\hat{y}^2 + 1/4}$. Lastly, since the count can be modeled as a Poisson random variable, the Anscombe

transform (Anscombe, 1948) is useful to change the distribution into an approximately Gaussian one with unit variance. This transform is simply the function $A(x) = 2\sqrt{x + 3/8}$. It is applied to \hat{y} and the observed in-window sum y to stabilize their variance. The final score is $A(\hat{y}) - A(y)$.

4.3 Matched Filter

Algorithm 3 Matched filter

```

1: procedure TRAIN(data  $X \in \mathbb{R}^{n \times d}$ , source template  $s \in R^d$ , number  $k$  of components)
2:    $\mu \leftarrow \text{MEAN}(X)$ 
3:    $W \leftarrow \text{diag}(\mu + 0.001)^{-1}$ 
4:    $\Sigma \leftarrow \text{COV}(X)$ 
5:    $B \leftarrow k$  leading eigenvectors of  $\Sigma$ 
6:    $T \leftarrow s^\top W (I_d - B(B^\top W B)^{-1} B^\top W)$ 
7:   return  $T$ 
8: procedure TEST(learned parameter  $T$ , measurement  $x \in R^d$ )
9:   return  $x^\top T / \|x\|_1$ 

```

If perfect knowledge of the source template is assumed, then more powerful classes of methods can be used such as Matched Filters (MF). In general, an MF achieves the optimal signal-to-noise ratio when there is additive noise by correlating a template with an unknown signal such as background (Turin, 1960). Here we refer to only a specific type of MF, the Orthonormal Subspace Projection Matched Filter (Tandon, 2016). Conceptually, MF is similar to the PCA method. In particular, an orthonormal subspace is found from a training set using PCA. A test measurement is projected onto the basis and back, and the residual between reconstruction and original is computed, thereby eliminating the most discriminative background features. However, whereas PCA computes the norm of this residual, MF computes the inner product with the source template. Thus, a higher score is given when the residual has a similar shape to the source. The algorithm is detailed in Alg. 3. Again, it works well to use $k = 5$ principal components.

4.4 Gaussian-Poisson MAP estimation

Besides the Matched Filter, there is another method available when the source template is known: Gaussian-Poisson (GP) MAP estimation. Uniquely among the methods mentioned so far, it defines a generative model. Through this distribution, it computes maximum likelihoods for the two cases of when a source is present and absent. The background energy template is drawn from a Bayesian prior, a multivariate normal distribution with fixed parameters. Simultaneously, the source intensity is drawn from a univariate normal distribution. Although the background rates and source intensity are not truly normally distributed, the priors capture the potential variability caused by the changing environment and distance to the source. The background counts are drawn from Poisson distributions whose parameters are given by the background energy template. Counts from the source are drawn from Poisson distributions whose parameters are given by the source template scaled by intensity. The observed combined counts are the sum of the background and source counts. The model is depicted in Fig. 4. The parameters of the source intensity are not shown because they are insignificant: the variance is set to a high value to create an uninformative prior.

Next we define the method. Let $z^{(i)}$ be the z^{th} element of vector z . The count $y_t^{(i)}$ of photons in the i^{th} energy bin is a Poisson distributed random variable with rate $x_t^{(i)} + J_t s^{(i)}$, where x_t is the observation-specific vector of background gamma-ray rates, s is the source spectrum, and J_t is the

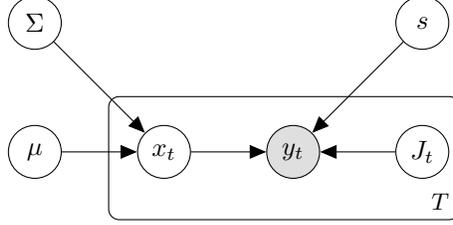


Figure 4: The Gaussian-Poisson model of radiation.

observation specific source intensity. For simplicity of presentation we omit the subscript t below. If x and S were known, and if it were further known that the source is either present (with intensity J) or absent (intensity zero), then the Bayes decision rule for determining the presence or absence of the source would be given by

$$\log(\Pr(y; x + Js)) - \log(\Pr(y; x)) = -\mathbf{1}_d^\top Js + y^\top \log\left(\frac{x + Js}{x}\right) > \tau$$

where τ is the decision threshold, d is the number of energy bins, $\mathbf{1}_d$ is a length d vector of ones, and $\log(\cdot)$ is taken element-wise. In practice s may be known, but x and J vary with t . In this case, we establish a Gaussian prior belief on the background rates and take the MAP estimate for x . We do the same for source intensity, using an uninformative prior with high variance. The resulting score function is given by

$$\begin{aligned} GP(y; \mu, K) = & \\ & \max_{x>0, J>0} \left\{ -\frac{1}{2}(x - \mu)^\top \Sigma^{-1}(x - \mu) - \mathbf{1}_d^\top(x + Js) + y^\top \log(x + Js) \right\} \\ & - \max_{x>0} \left\{ -\frac{1}{2}(x - \mu)^\top \Sigma^{-1}(x - \mu) - \mathbf{1}_d^\top x + y^\top \log(x) \right\} \end{aligned}$$

where μ and Σ are mean and covariance of the background prior. The constituent optimization problems are concave, making $GP(y; \mu, \Sigma)$ efficient to compute using quasi-Newton methods.

5 Uncertain Censored Energy Windows

In this section, a method is presented for a setting in which information about threat templates is incomplete. When methods like CEW and MF are given accurate information about the threats, they usually perform satisfactorily. In practice, however, the design of a threat could vary due to different configurations of material, shielding, and environment (Tandon, 2016). These methods can heavily underperform when their information is incomplete or inaccurate. For example, the computed energy window for CEW could be inaccurate. If so, the observed in-window sum of counts might not drastically change when a source is present. To remedy this shortcoming, we propose a method called Canonical Correlation Analysis (CCA) detection. This method discovers a source in the presence of background noise while tolerating imperfect knowledge of the source template. It could be useful in practical applications when the designs of sources of harmful radiation are not precisely known. On the source information scale, CCA detection can be placed in between PCA anomaly detection, which uses no source-specific information, and source-type-aware methods like Matched Filter.

5.1 Canonical Correlation Analysis

Canonical Correlation Analysis (CCA) is a statistical method for finding correlations between two sets of variables (Hotelling, 1936). Each correlation can be viewed as a latent factor that helps explain the variance in both sets simultaneously. Formally, given random vectors $X \in \mathbb{R}^{d_1}$ and $Y \in \mathbb{R}^{d_2}$, CCA solves the optimization problem $\max_{u,v} \text{Corr}(X^T u, Y^T v)$, in which $u \in \mathbb{R}^{d_1}$ and $v \in \mathbb{R}^{d_2}$ are vectors of coefficients. This procedure is iterated up to $\min\{d_1, d_2\}$ times under the constraint that $X^T u$ is uncorrelated with $X^T u'$ for any previously found u' , with an analogous constraint on $Y^T v$. When $X \in \mathbb{R}^{n \times d_1}$ and $Y \in \mathbb{R}^{n \times d_2}$ are centered data matrices, the solution to CCA has a well known estimator. Let $\Sigma_{XX} = X^T X$, $\Sigma_{YY} = Y^T Y$, and $\Sigma_{XY} = X^T Y$. Then the CCA coefficients U of X are the $\min\{d_1, d_2\}$ leading eigenvectors of

$$\Sigma_{XX}^{-1} \Sigma_{XY} \Sigma_{YY}^{-1} \Sigma_{XY}^T$$

with coefficients V of Y given by an analogous formula.

5.2 CCA detection

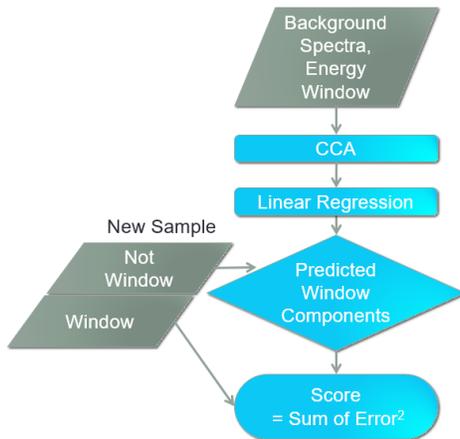


Figure 5: The CCA detection algorithm.

Our method takes as input a collection of background measurements and an energy window. This window usually maximizes signal-to-noise ratio (SNR) for a particular source template of interest, as in Alg. 2, but it can be arbitrary, and we vary its quality in our experiments. The first step is to apply CCA to the training set. Here, X and Y correspond to photon counts inside and outside the energy window respectively. Next, for each pair (u, v) of weights found, we fit a simple linear regression of $X^T u$ on $Y^T v$. We compute the residuals for all samples and fit a univariate normal distribution to them. Then given a new sample to classify, we compute the regression residuals and find their z -scores for each pair (u, v) . The sign of each score is not meaningful, unlike CEW in which an elevated observation suggests a source. Thus, we square the z -scores and compute the sum as the final score. In the same way that CEW finds a single-output linear relationship between photon counts inside and outside the energy window, CCA finds arbitrary multiple-to-multiple linear relationships. A diagram of the method is given in Fig. 5.

6 Source independence experiment

This section describes an experiment that simulated incomplete knowledge of the source template.

6.1 Data

Our dataset was a collection of over 86,000 gamma-ray measurements recorded in one-second intervals by a sodium-iodide detector mounted on a vehicle moving around an urban area in Baltimore, MD, USA. On average there were 2,600 photon counts per second. Each measurement had $d = 116$ linearly spaced energy bins. The measurements were assumed to be background data. In preprocessing they were randomly partitioned into training and test sets. Along with the background, we had a library of 67 source templates corresponding to high-fidelity source simulation applying different configurations of radioactive material and shielding. The templates were normalized to have intensities of 100 counts per second. Each template was sampled as a spectrum of Poisson random variables once for every test observation. These samples were then added to the test observations to form synthetic positive measurements. This process resulted in labeled test data for every source in the library.

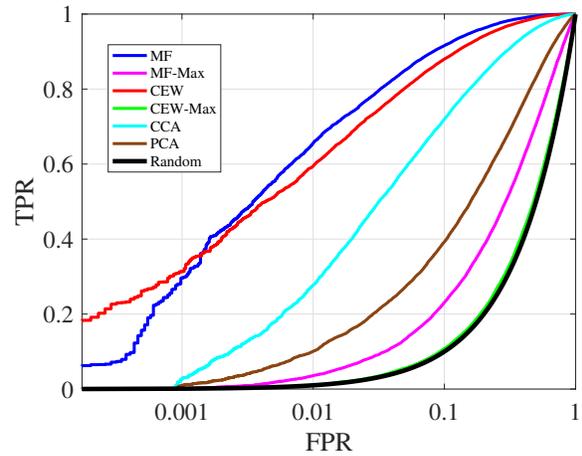
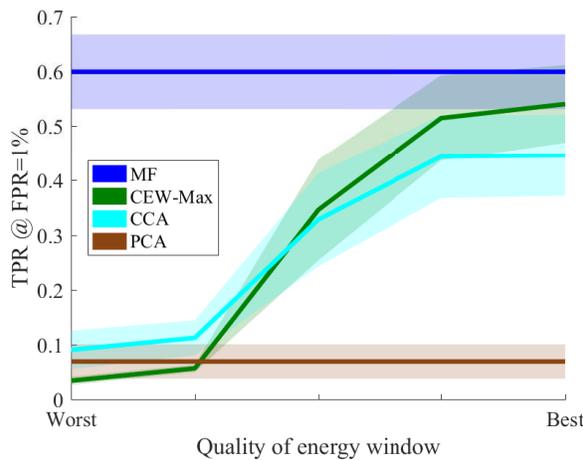
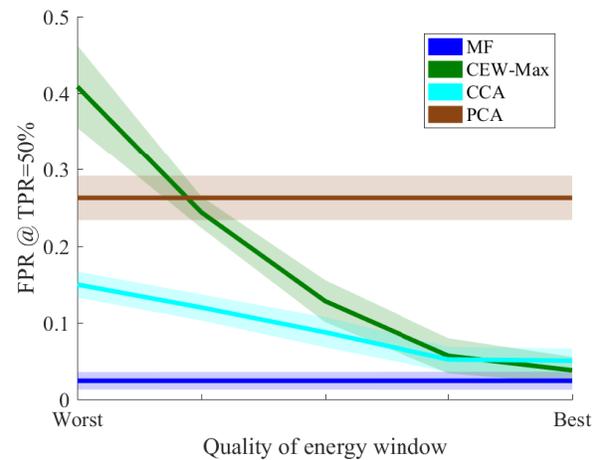


Figure 6: Example of methods' performance on a single source template.

6.2 Results



(a) Comparison between methods' TPR.



(b) Comparison between methods' FPR.

Figure 7: Comparison between methods as energy windows change from common to source-specific.

The methods compared to the CCA method were MF, CEW, and PCA. MF and CEW had two versions: in the first, the methods used the true source template; in the second, they marginalized over every template in the library by using the maximum score (CEW-Max and MF-Max). Now MF uses strictly more information about the source spectrum, so it approximately upper-bounds performance as in Fig. 6, though it could possibly still be worse than CEW in some cases.

In the first part of the experiment, we directly measured the impact of energy window quality. Note that only CEW and CCA use energy windows. We checked how CEW-Max and CCA degraded as the energy window changed from an accurate, source-specific window to a common window that maximizes average SNR across all considered sources but it was not specifically tailored to any of them. A convex combination was taken between each source and the average source with weights of 0, 0.25, 0.5, 0.75, and 1. The optimal energy window was computed for the combined source. We computed the false positive rate (FPR) at the true positive rate (TPR) of 50%, and TPR at the FPR of 1%. The results, averaged across sources, are displayed in Fig. 7b and 7a, along with the source-optimal MF and PCA anomaly detection. With source-specific windows, CCA performed slightly worse on average than CEW. With a common window, CCA enjoyed large, statistically significant advantages in both metrics. TPR was higher by 0.05 and FPR was lower by 0.16.

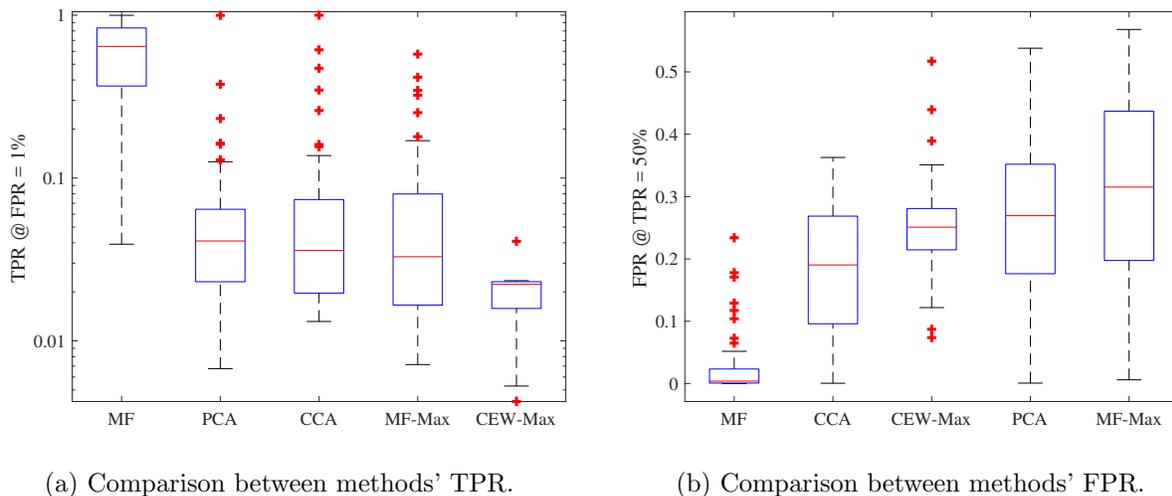


Figure 8: Comparison between methods when target source is missing from the library.

In the second part of the experiment, a lack of information was simulated. To do so, we clustered the threat library using K -means with $K = 10$. For MF-Max, CEW-Max, and CCA, we tested each source by removing its cluster from the data set and marginalizing over the remainder. Fig. 6 shows an example of the Receiver Operating Characteristic (ROC) of the classifiers on a particular source template. The distribution of FPR and TPR over threats is displayed in Fig. 8. By TPR (Fig. 8a), it is difficult to distinguish between PCA, CCA, and MF-Max. By FPR (Fig. 8b), however, the CCA method usually does better than MF-Max, CEW-Max, and PCA.

6.3 Discussion

CCA detection has been demonstrated to be more robust to poor knowledge of the target source template. When the source is unknown to the library, CCA detection usually performs better than PCA detection, CEW-Max, and MF-Max. Likewise, when the energy window is suboptimally computed—corresponding to a situation in which the source template is unknown or inaccurate—CCA detection performs better than CEW-Max by a significant margin. We can hypothesize a reason for why CCA detection is more independent of source information than CEW: whereas CEW finds one relationship between in-window and out-of-window energy bins, CCA detection finds multiple uncorrelated relationships, which reduces the chance that every relationship becomes uninformative when the window is inaccurate. Intuitively, by taking a simple sum in the window,

CEW only utilizes a fraction of the spectral information available. By comparison, CCA detection utilizes much more by finding more than one relationship. Furthermore, a plausible explanation for why CCA is more robust than MF is that CCA simply assumes weaker knowledge of the source—an energy window instead of the exact template. Intriguingly, CCA also performs better than the completely source independent PCA detection even though the knowledge about the source is imperfect. In contrast, MF and CEW, which depend much more on source information, are worse than PCA. Thus, CCA detection may represent a superior trade-off between robustness to source information and capacity to leverage that information.

7 Simultaneous estimation of source strength and background

This section discusses a method that has much lesser dependence on training data. All previous methods that were discussed here in detail rely on training data to learn characteristics of the distribution of background spectra. When the training set resembles the test set, the methods perform well. Nonetheless, in practice, training data may be scarce, or the test observations come from a much different distribution. In these cases the methods become less useful. We suggest a direct improvement to Gaussian-Poisson MAP estimation for the setting of perfect information about the source template. We demonstrate how to limit dependence on training data and adapt to the test data by exploiting the local structure of the time series of observations. In particular, we simultaneously estimate source strength and background spectrum. The method may be categorized as an adaptive Kalman filter, which is a class of methods that dynamically estimate covariances of process and measurement noise when they are unknown (Mehra, 1970). On the generalization scale, this method ranks above all others in this study.

7.1 Notation

We consider discrete stochastic processes indexed by $t \in \{1, \dots, T\}$. We use $a_{i:j}$ to refer to $\{a_i, a_{i+1}, \dots, a_j\}$ in a stochastic process $\{a_t\}$. Let $x^{(i)}$ be the i^{th} element of vector x . Let $x^{(i:j)}$ be a vector containing elements i through j (inclusive) of x . Let $\text{diag}(x)$ denote a diagonal matrix with x as the main diagonal for a vector x .

7.2 Kalman filter for radiation

The Kalman filter (Kalman, 1960) is a widely used method for estimating an unobserved signal underlying an observed noisy discrete time series. In plain terms, it filters out statistical noise to obtain more precise estimates of a signal, such as in Fig. 9. Consequently, it is a natural candidate to model radiation. A Bayesian approach, it tracks a mean and covariance of a prior over the signal at every time step. An example of a common application would be estimating the

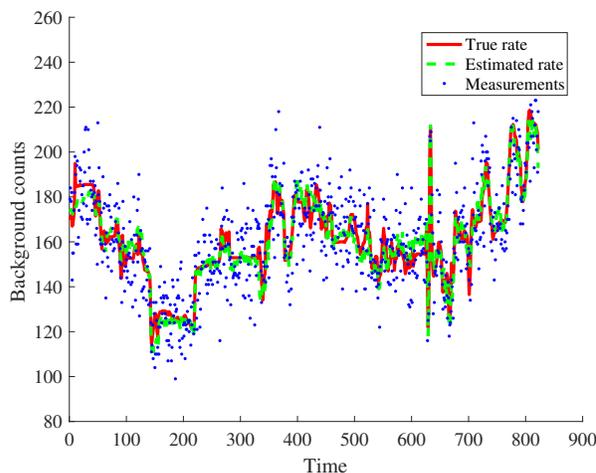


Figure 9: Using the Kalman filter to estimate the background rates from noisy measurements in an energy bin.

position of a vehicle using GPS. Since GPS measurements contain substantial noise, they could be non-smooth. When a Kalman filter is applied, however, one can incorporate information about the velocity and acceleration of the car, inferred from the previous estimates of position, to refine the current estimate.

The Kalman filter estimates the state of a linear dynamical system in discrete time. Let $x_t \in \mathbb{R}^d$ be the state space and $y_t \in \mathbb{R}^p$ be the observation at time t . Let $A \in \mathbb{R}^{d \times d}$ and $C \in \mathbb{R}^{p \times d}$ be the transition and emission matrices respectively. Let $Q_t \in \mathbb{R}^{d \times d}$ and $R_t \in \mathbb{R}^{p \times p}$ be the covariances of process noise and measurement noise respectively. The parameters A , C , $\{Q_t\}$, and $\{R_t\}$ are assumed to be known. In this model we allow the covariances are non-stationary, which is uncommon, but as long as they are non-random all calculations are functionally identical to the stationary case. The dynamical system is defined as

$$x_{t+1} = Ax_t + w_t \quad y_t = Cx_t + v_t$$

where $E(w_t) = 0$, $\text{Var}(w_t) = Q$, $E(v_t) = 0$, and $\text{Var}(v_t) = R$. We assume the $\{w_t\}$ and $\{v_t\}$ are independent. The update equations are well documented and thus omitted here. The state space model can also contain a control input that affects x_t , but we exclude it in this work as it does not fit the application. The Kalman filter can be proved to give the best linear minimum mean squared error estimator of x_t given $y_{1:t}$. Also, if the noise variables are Gaussian, then it is the optimal minimum mean squared error estimator (Shimkin, 2009).

Now we discuss how the Kalman filter can be applied to simultaneously estimate source intensity and background rates. In the GP method, the Gaussian prior on the background rates has a time-invariant mean and covariance. The method can be improved by dynamically estimating the background rates at each step to better estimate the parameters of the prior. To do so we apply the Kalman filter. The state space $x_t \in \mathbb{R}^{d+1}$ contains the background rates in each of d energy bins in the first d elements. Furthermore, the source intensity must be modeled in order to decouple its effect on the measured spectra from the background radiation. Source intensity can vary with distance and obfuscation. As a result, we model source intensity in the last element $x_t^{(d+1)}$. Next, the observation $y_t \in \mathbb{R}^d$ contains the measured counts at time t . The transition matrix A is set to the identity matrix because we have no *a priori* expectation as to how the background rates change. We propose a dynamical system given by

$$x_{t+1} = x_t + w_t \quad y_t = x_t^{(1:d)} + x_t^{(d+1)}s + v_t$$

where $s \in \mathbb{R}^d$ is the target source template. The second equation implies a form for the emission matrix C . The Kalman filter estimates the background rates $x_t^{(1:d)}$, which are used as the mean μ of the Gaussian prior in the GP method. In addition, the covariance K of the prior is estimated as the covariance of the estimated background rates up to time t .

Next we explain how to set the remaining parameters, the transition noise covariances $\{Q_t\}$ and emission noise covariances $\{R_t\}$.

7.3 Adaptive filtering

Conventionally, the covariances of process and measurement noise are taken as input to the Kalman filter; in many scenarios, however, their values are unknown (Mehra, 1970). For instance, in the radiation application it is possible that training data are scarce or differ from the test data, so it would be infeasible to estimate the covariances beforehand. When the covariances are estimated in real-time, the resulting method is known as an adaptive Kalman filter. This class of methods has a large body of research. The different approaches include Bayesian, maximum likelihood,

covariance matching, and correlation techniques (Ottersten et al., 1998; Akesson et al., 2008; Bavdekar et al., 2011; Odelson et al., 2016). Most approaches construct estimators with desirable statistical properties such as asymptotic normality, unbiasedness, and consistency. Yet these approaches are often computationally difficult and complicated to implement. Here we propose a simple method for estimating the noise covariances by exploiting knowledge of the radiation domain. We do not show any statistical properties of the estimators but instead showcase their utility empirically. Although performance could be further improved by a more accurate and precise means of covariance estimation, the aim of this work is not to compare this method to other adaptive filters but only to show that it performs well on the task at hand.

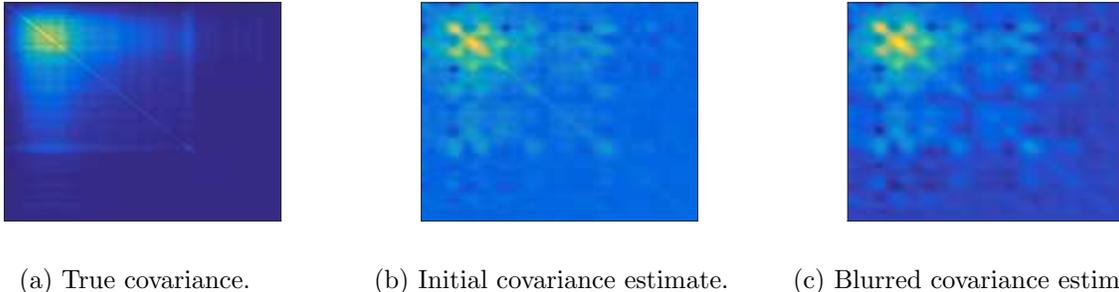


Figure 10: Estimating the process noise covariance.

The method for computing the noise covariances is to use the previous state estimates as training data. In the process noise covariance Q_t , only the upper $d \times d$ block is estimated. For the remainder, we assume the background rates to be independent of source intensity, and the variance of the intensity is set to 10. The exact value is relatively insignificant. For the upper $d \times d$ block, we first compute $\Delta \hat{x}_i = \hat{x}_i - \hat{x}_{i-1}$ for $i = 1, \dots, t - 1$. We let \hat{Q}_0 be the sample covariance of this variable. A heat map of one instance of this matrix is given in Fig. 10b. The values are not locally smooth: nearby elements form a jagged surface. For comparison, Fig. 10a shows the true covariance (whose calculation is explained in the next experiment section); nearby elements form a smooth surface. Therefore, the next and final step is to apply Gaussian smoothing to \hat{Q}_0 . We selected a kernel length of 10 with a unit variance, but the smoothing hyperparameters were not found to be especially important. The result is shown in Fig. 10c. In our experiments, this process produced an effective, if potentially biased, estimator.

The measurement noise R_t can also be easily computed. Given the state x_t , the measurement noise $v_t = y_t - x_t$ is Poisson. In particular, $v_t^{(j)} \sim \text{Poisson}(x_t^{(j)})$. Furthermore, the elements are known to be conditionally independent. Therefore $R_t = \text{diag}(x_t)$. However, since \hat{x}_t is not known at time t , we instead estimate $\hat{R}_t = \text{diag}(\hat{x}_{t-1})$. In practice, this estimate performs better when divided by 2, which causes the estimated state to track the measurements more closely.

7.4 Practicalities

Burn-in period We found it useful to enforce a burn-in period in our experiments. In particular, we assumed that the first $B = 150$ time steps never contained interference, so we forced $\hat{x}_t^{(d+1)} = 0$ when $t \leq B$ regardless of the output of the GP method or Kalman filter. The assumption was justified because a spectrometer can stay in a zone known to have no source during the burn-in period. This approach allowed the filter to calibrate to the correct noise covariances. Furthermore, in the initial 5 time steps, we overrode the estimated rates with the mean of the observations up to the current time step to calibrate to the correct initial state. However, the burn-in period is a

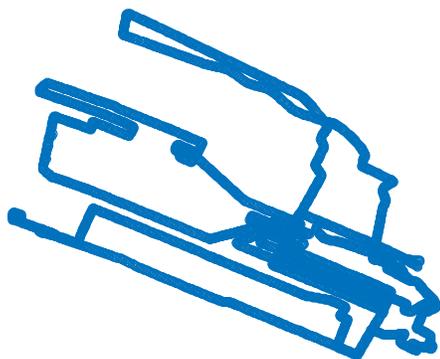


Figure 11: Path of the sensor.

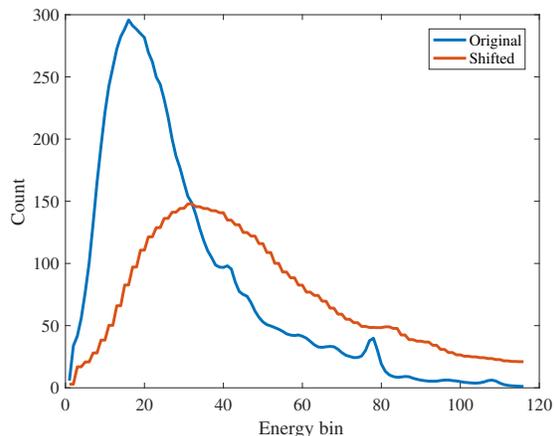


Figure 12: Training observations shifted to decrease correlation with test observations.

much weaker requirement than a training set. If the burn-in period were used for training other methods, they would still be susceptible to changes in the background distribution because they are non-adaptive. This method, on the other hand, adapts to the local distribution and only requires the burn-in period to ensure that the parameters are initialized correctly.

Unquantified uncertainty There is an inconsistency in the assumptions of this method. While the conventional Kalman filter assumes Q_t and R_t are non-random, the estimators in this method depend on the output of the filter itself at the previous time steps, which is subject to random noise. Hence, when the filter estimates the uncertainty of its prediction, it misses the additional uncertainty contributed by the covariance estimators. We leave this issue as future work.

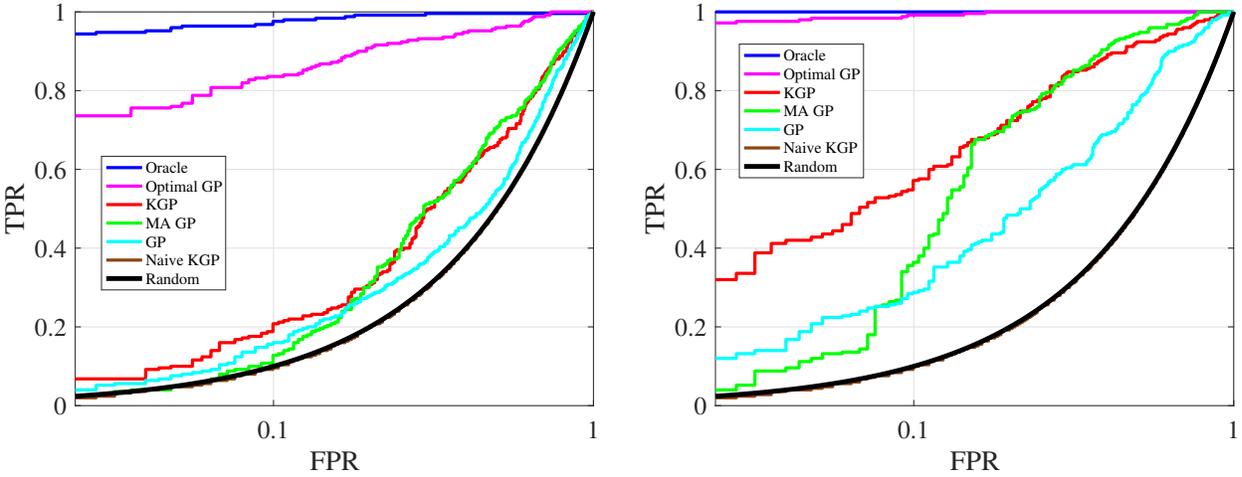
Negative values Intensity and background rates are always nonnegative in reality. However, the Kalman filter has no restriction on the signs of the process. Hence, after every iteration, we enforce nonnegativity by letting $\hat{x}_t \leftarrow \max(0, \hat{x}_t)$.

8 Generalization experiment

This section presents an experiment that simulated test data that greatly differed from training data.

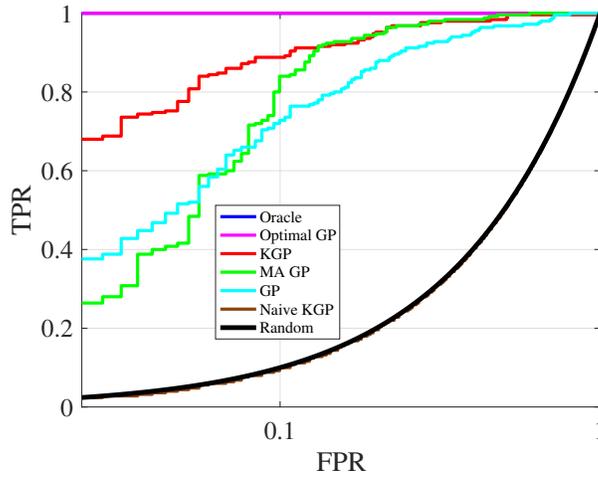
8.1 Data

Our dataset was a collection of over 11,000 gamma-ray measurements recorded in one-second intervals by a sodium-iodide detector mounted on a vehicle moving around an urban area in downtown Berkeley, CA, USA. On average there were 10,000 photon counts per second. Each measurement contained $d = 116$ quadratically spaced energy bins. The measurements were assumed to be background data. Also, the GPS position of the vehicle was recorded. The background rates were estimated from the measurements and viewed as the true rates. The method used was GP estimation over observations within 10 meters of the current measurement being estimated. We also had the same library of 67 source templates used in the previous experiment. Unlike before, only one template was selected: the one with minimum correlation with the mean background spectrum.



(a) ROC with intensity at 75 counts per second.

(b) ROC with intensity at 125 counts per second.



(c) ROC with intensity at 175 counts per second.

Figure 13: Comparison of methods on generalization dataset.

8.2 Procedure

This experiment tested detection of a roadside source in a single pass. In each run of the experiment, a location for the synthetic source was randomly sampled from the set of points at most 20 meters from the path of the sensor (Fig. 11). The test set was taken to be the data points where the sensor was at most 100 meters away from the source, along with all data points in between, forming a contiguous window. The window was extended on either side by 450 observations. The training set was taken to be the remaining data points. To simulate different distributions between the training and test sets, we applied a shifting algorithm to each measurement in the training set (Fig. 12) The counts in any particular energy bin were shifted to higher energy bins. This shift led to Pearson correlation of roughly 50% between the mean training background spectrum and the test background spectra. Every observation in the test set was injected with a sample from the source distribution. The intensity was calculated to be inversely proportional to the square of distance from the source (Miller et al., 2016). The maximum intensity was set to a given value. The methods

were applied to score each test point. Then the maximum score for each method was recorded as the output of the run. Next, all runs were repeated identically except with no source injection. By this procedure we created multiple runs with positive or negative ground truth.

8.3 Results

We compared performance between six methods:

1. Oracle: Poisson likelihood ratio using perfect knowledge of the background rates and source intensity.
2. Optimal GP: GP method using perfect knowledge of the background rates to set the prior mean and covariance.
3. Kalman GP (KGP): GP method using the adaptive Kalman filter to set the prior mean and covariance.
4. Moving Average GP (MA GP): GP method using a 30-point moving average of the background rates to set the prior mean and covariance.
5. GP: GP method using the training data to set the prior mean and covariance.
6. Naive KGP: GP method using the non-adaptive Kalman filter with covariance parameters estimated from the training data in the same manner as KGP.

Fig. 13 shows the ROCs computed over 250 positive and negative runs with maximum intensities of $M \in \{75, 125, 175\}$ counts per second. For every M the best methods are the Oracle and Optimal GP, as expected. The next best is our proposed method, KGP. KGP beats the next best method, MA GP, by a large margin in the low FPR range. GP does poorly, and Naive KGP performs the same as random guessing.

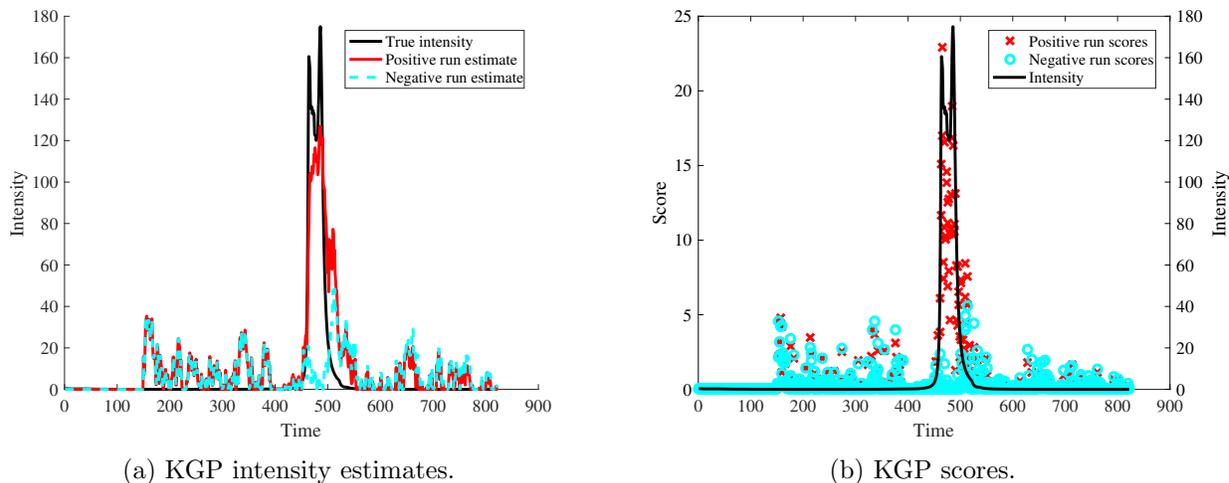


Figure 14: KGP intensity estimates and scores in a positive and negative run.

In Fig. 14a, an example is given of how the adaptive Kalman filter estimates the source intensity in a positive run and the corresponding negative run. Fig. 14b illustrates the scores over time of different methods on the positive run. Fig. 15 shows how the adaptive Kalman filter accurately estimates the

background rate even when the measurements contain a strong source. The measurements greatly exceed the true rates, yet this elevation does not greatly affect the estimated rates.

8.4 Discussion

KGP performs better than MA GP, GP and Naive KGP when the training set is uninformative. These other competitors represent suboptimal ways to adapt to the local data distribution. MA GP attempts to smooth the observed counts to estimate the background rates. However, when the window size is large, it does not track sudden changes in the background. When the window size is small, it mistakenly absorbs the injected source into the estimated background. KGP has an advantage because it allows for sudden changes to be estimated while being robust against a source. Next, GP does not attempt to adapt to local information at all and therefore performs badly, although it is still significantly better than random guessing. Lastly, Naive KGP produces a completely uninformative score (indistinguishable from random guessing in Fig. 13) This method tries to adapt to the test data by estimating background rates and source intensity but does not adapt to the local covariance structure. Evidently, its estimates of rates and intensity are completely inaccurate. In this case, a poor adaptation is worse than no adaptation at all.

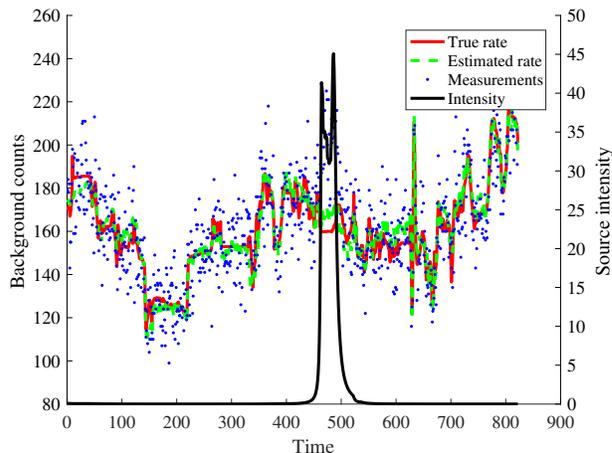


Figure 15: Using the adaptive Kalman filter to estimate the background rates in an energy bin in the presence of a source.

9 Conclusion

This study made three primary contributions. First, several existing methods for threat detection were documented in detail. Most of these methods assumed knowledge of the target source, and all assumed training data. Second, a method was proposed for the setting in which the assumption of knowledge about the source is violated. Third, a method was proposed for the setting in which the assumption of informative training data does not hold. Both new methods were empirically demonstrated to perform better than competitors. In the future, it would be interesting to review these methods when the spectra contain less than 100 total counts per second, which is a major challenge facing many methods (Tandon, 2016).

Acknowledgements

This work has been partially supported by the U.S. Department of Homeland Security, Domestic Nuclear Detection Office’s Academic Research Initiative under awards 2014-DN-077-ARI087-01 and 2016-DN-077-ARI108, and by the U.S Department of Defense, Defense Threat Reduction Agency under award HDTRA1-13-1-0026. This work was the result of collaboration with Kyle Miller¹, Peter

¹Auton Lab, Carnegie Mellon University

Huggins¹, Artur Dubrawski¹, Karl Nelson², Simon Labov², and Jay Jin³, although they may not agree with all the conclusions of this paper.

References

- Aage, H. K. and Korsbech, U. (2003). Search for lost or orphan radioactive sources based on NaI gamma spectrometry. *Applied Radiation and Isotopes*, 58(1):103–113.
- Akesson, B. M., Jorgensen, J. B., Poulsen, N. K., and Jorgensen, S. B. (2008). A generalized autocovariance least-squares method for Kalman filter tuning. *Journal of Process Control*, 18(7-8):769–779.
- Anderson, K. K., Jarman, K. D., Mann, M. L., Pfund, D. M., and Runkle, R. C. (2008). Discriminating nuclear threats from benign sources in gamma-ray spectra using a spectral comparison ratio method. *Journal of Radioanalytical and Nuclear Chemistry*, 276(3):713–718.
- Ancombe, F. J. (1948). The Transformation of Poisson, Binomial and Negative-Binomial Data. *Biometrika Trust*, 35:246–254.
- Aucott, T. J., Bandstra, M. S., Negut, V., Curtis, J. C., Chivers, D. H., and Vetter, K. (2014). Effects of background on gamma-ray detection for mobile spectroscopy and imaging systems. *IEEE Transactions on Nuclear Science*, 61(2):985–991.
- Aucott, T. J., Bandstra, M. S., Negut, V., Curtis, J. C., Meyer, R. E., Chivers, D. H., and Vetter, K. (2015). Impact of detector efficiency and energy resolution on gamma-ray background rejection in mobile spectroscopy and imaging systems. *Nuclear Instruments and Methods in Physics Research, Section A*, 789:128–133.
- Bai, E. W., Chan, K. S., Eichinger, W., and Kump, P. (2011). Detection of radionuclides from weak and poorly resolved spectra using Lasso and subsampling techniques. *Radiation Measurements*, 46(10):1138–1146.
- Bavdekar, V. A., Deshpande, A. P., and Patwardhan, S. C. (2011). Identification of process and measurement noise covariance for state and parameter estimation using extended Kalman filter. *Journal of Process Control*, 21(4):585–601.
- Du, Q., Wei, W., May, D., and Younan, N. H. (2010). Noise-adjusted principal component analysis for buried radioactive target detection and classification. *IEEE Transactions on Nuclear Science*, 57(6 PART 2):3760–3767.
- Fagan, D. K., Robinson, S. M., and Runkle, R. C. (2012). Statistical methods applied to gamma-ray spectroscopy algorithms in nuclear security missions. *Applied Radiation and Isotopes*, 70(10):2428–2439.
- Hotelling, H. (1936). Relations between two sets of variates. *Biometrika*, 28(3/4):321–377.
- Kalman, R. E. (1960). A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering*, 82(1):35.

²Lawrence Livermore National Laboratory

³Google Inc.

-
- Kump, P., Bai, E. W., Chan, K. S., and Eichinger, W. (2013). Detection of shielded radionuclides from weak and poorly resolved spectra using group positive RIVAL. *Radiation Measurements*, 48(1):18–28.
- Mehra, R. (1970). On the identification of variances and adaptive Kalman filtering. *IEEE Transactions on Automatic Control*, AC-15(2):175–184.
- Miller, K., Huggins, P., Labov, S., Nelson, K., and Dubrawski, A. (2016). Evaluation of coded aperture radiation detectors using a Bayesian approach. *Nuclear Instruments and Methods in Physics Research, Section A*, 839:29–38.
- Nelson, K. and Labov, S. (2012). Aggregation of mobile data. *Lawrence Livermore National Lab Technical Report*, 2.2(1):2–3.
- Odelson, B. J., Rajamani, M. R., and Rawlings, J. B. (2016). A new autocovariance least-squares method for estimating noise covariances. In *Texas-Wisconsin Modeling and Control Consortium*, volume 28, pages 391–397.
- Ottersten, B., Stoica, P., and Roy, R. (1998). Covariance matching estimation techniques for array signal processing applications. *Digital Signal Processing*, 8(3):185–210.
- Pfund, D., Anderson, K., Detwiler, R., Jarman, K., McDonald, B., Milbrath, B., Myjak, M., Paradis, N., Robinson, S., and Woodring, M. (2016). Improvements in the method of radiation anomaly detection by spectral comparison ratios. *Applied Radiation and Isotopes*, 110:174–182.
- Runkle, R. C., Tardiff, M. F., Anderson, K. K., Carlson, D. K., and Smith, L. E. (2006). Analysis of spectroscopic radiation portal monitor data using principal components analysis. *IEEE Transactions on Nuclear Science*, 53(3):1418–1423.
- Shimkin, N. (2009). Derivations of the discrete-time Kalman filter: The stochastic state-space model. Technical report.
- Tandon, P. (2016). Bayesian aggregation of evidence for detection and characterization of patterns in multiple noisy observations. *AI Matters*, 2(3):25–26.
- Tandon, P., Huggins, P., Maclachlan, R., Dubrawski, A., Nelson, K., and Labov, S. (2016). Detection of radioactive sources in urban scenes using Bayesian aggregation of data from mobile spectrometers. *Information Systems*, 57:195–206.
- Turin, G. L. (1960). An introduction to matched filters. *IRE Transactions on Information Theory*, 6(3):311–329.
- Ziock, K. P., Craig, W. W., Fabris, L., Lanza, R. C., Gallagher, S., Horn, B. K. P., and Madden, N. W. (2004). Large area imaging detector for long-range, passive detection of fissile material. *IEEE Transactions on Nuclear Science*, 51(5 I):2238–2244.