Robust Detection of Radiation Threat Using Uncertain Censored Energy Windows

E. Lei*, K. Miller**, P. Huggins**, A. Dubrawski**
* Machine Learning Department, Carnegie Mellon University ** Auton Lab, Carnegie Mellon University

Introduction. Given gamma-ray spectrometer data collected from noisy background environments, a popular method for detecting threats of known types is the Censored Energy Window (CEW) [1]. CEW aggregates features of the target energy window where we expect to see the source signatures most clearly, potentially losing information in the process. In addition, CEW is sensitive to the quality of selected target energy window, which may be an issue in practice when the optimal energy window is not perfectly known *a priori*.

Objective. We aim to develop a classifier of measured photon count spectra for detecting a source signature that exploits the individual energy bins in the target window and is robust to the choice of this window.

Background. A gamma-ray spectrum measurement is typically represented as a vector of photon counts aggregated over a fixed number of consecutive, non-overlapping energy bins. CEW is trained from measurements of source-free (background) data inside and outside of the source-type-specific energy window [1]. It fits a multiple linear regression to predict the sum of photon counts in the window from the vector of binned photon counts outside of it. The threat detection score is a difference between the predicted and the actually observed sum of photons in the window, normalized by the variance of this sum. However, to reduce the loss of potentially useful information, we instead apply Canonical Correlation Analysis (CCA), a statistical method for finding structured correlations [2], to avoid summing up photons in individual bins of the target energy window. Given two sets of variables $X \in \mathbb{R}^{d_1}$ and $Y \in \mathbb{R}^{d_2}$, CCA solves the optimization problem $\max_{u,v} \operatorname{corr}(X^\top u, Y^\top v)$, 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 $X^\top u$ is orthogonal to $X^\top u'$ for any previously found u', and analogously for $Y^\top v$. In this work, X and Y correspond to photon counts inside and outside the energy window.

Method. Our method takes as training input a collection of background spectra and an energy window. This window usually maximizes signal-to-noise ratio (SNR) for a particular source template of interest, but it can be arbitrary, and we vary its quality in our experiments. The output for a new measured spectrum is a score indicating confidence that the sought after source is manifested in it. The first step of our method is to apply CCA to find correlations between photon counts inside and outside the window in background data. Next, for each pair (u, v) of weights found, we fit a linear regression of $X^\top u$ on $Y^\top v$. We compute the residuals for all samples and fit a univariate normal distribution to them. Then given a new sample to classify, we again 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. 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, so we hypothesize that these correlations are more robust at the threat detection task when the choice of energy window is suboptimal.

To compute the optimal energy window for each source, we rely on the following algorithm, which takes the source signature as input, represented as a vector of Poisson rates over energy bins. We compute the ratio of the rates to the mean background sample and sort the energy bins in descending order of this ratio. For each integer k from 1 to the number of bins, we obtain the sum S_k of the top k bins and the sum B_k of k corresponding mean background bins. We find $k^* = \arg\max_k S_k/\sqrt{B_k}$ and output as the energy window the top k^* bins according to the ratio of source rates to mean background.

Data. Our dataset is a collection of over 86,000 gamma-ray measurements recorded in one-second intervals by a sodium-iodide detector mounted on a truck moving around an urban area. These measurements were assumed to be background data, and a random subset of 35,000 was injected with source signatures to create synthetic positive samples. Each source signature was a vector of Poisson rates over energy bins from which we took i.i.d. samples, which were then added to the background measurements. Each vector was normalized to have a sum, or intensity, of 100. This process resulted in labeled testing data for each of the 67 sources in our library. We used high-fidelity source simulation applying different configurations of radioactive material and shielding.

Results. We compared our method to CEW and match filter (MF). MF uses strictly more information about the source spectrum, so it upper-bounds performance [3]. We checked how the methods degraded as the energy window changed from an accurate, source-specific window to a window that maximizes average SNR across all considered sources but it is 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, and the optimal energy window was computed for the combined source. We computed the average area under the Receiver Operating Characteristic (AUC), false positive rate (FPR) at the true positive rate (TPR) of 50%, and TPR at the FPR of 1%. The results are displayed in Fig. 1, 2, and 3. With source-specific windows, CCA performed slightly worse on average than CEW. With a common window, CCA enjoyed large, statistically significant advantages in all metrics. AUC was higher by 0.15 and TPR by 0.05, and FPR was lower by 0.16.

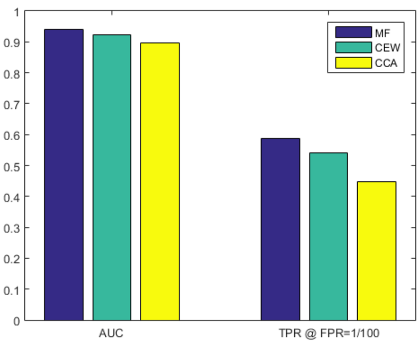

Figure 1: Comparison between methods' AUC as energy windows change from common to source-specific.

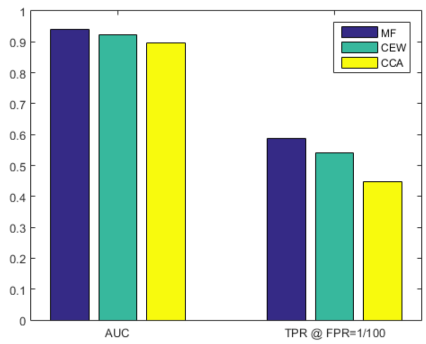

Figure 2: Comparison between methods' TPR as energy windows change from common to source-specific.

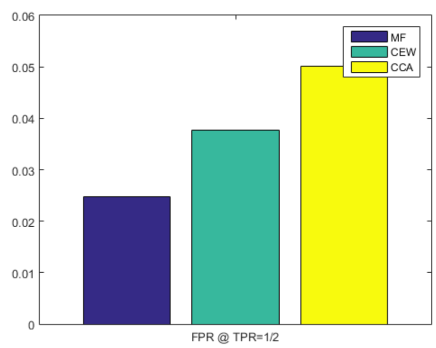

Figure 3: Comparison between methods' FPR as energy windows change from common to source-specific.

Conclusion. We characterized photon counts distributions in gamma-ray spectra using Canonical Correlation Analysis to replace Censored Energy Window with a multiple-to-multiple bin regression. This approach produced a classifier that was substantially more sensitive and threat-specific than CEW when the energy window was not exactly optimal for the particular source template, which is often the case in practice when characteristics of targeted threats are only approximately known.

References.

[1] Nelson, Karl, and Simon Labov. "Aggregation of Mobile Data." Lawrence Livermore National Lab Technical Report 1.2.2 (2012): 2-3.

[2] Hotelling, Harold. "Relations between two sets of variates." Biometrika 28.3/4 (1936): 321-377.

[3] Tandon, Prateek. "Bayesian Aggregation of Evidence for Detection and Characterization of Patterns in Multiple Noisy Observations." No. CMU-RI-TR-15-23. Robotics Institute, Carnegie Mellon University (2015).

Acknowledgements. This work has been partially supported by the U.S. Department of Homeland Security, Domestic Nuclear Detection Office, under competitively awarded grant 2014-DN-077-ARI087-01, by the U.S. Department of Defense, Defense Threat Reduction Agency under award HDTRA1-13-1-0026, and by the National Science Foundation under award 1320347. This support does not constitute an express or implied endorsement on the part of the U.S. Government.