Robust Bayesian Aggregation for Radioactive Source Detection

Jack H. Good, Ian Fawaz, Kyle Miller, and Artur Dubrawski

The Auton Lab, The Robotics Institute, Carnegie Mellon University

Introduction

Source detection is the problem of determining the presence or absence of certain radioactive materials, usually with gamma-ray spectroscopic data. It has applications in

- medicine
- industrial safety
- national security

The source detection problem, especially in busy and dynamic urban environments, a likely target for radiological dispersion devices, presents challenges:

- variation in background radiation
- uncertain sensor position due to GPS error
- dynamic gamma-ray occlusions
- moving sources
- nuisance sources, such as medical isotopes
- anisotropic shielding

This work focuses on adapting the Bayesian Aggregation (BA) framework [1] to **overcome uncertainty in background, position, and occlusion**. Future work will focus on moving sources, nuisance sources, and anisotropic shielding.

Gamma-ray Spectroscopic Data

A gamma-ray sensor measurement is a vector of photon counts at several energy bins over a fixed time interval. These observations are the input to our model.

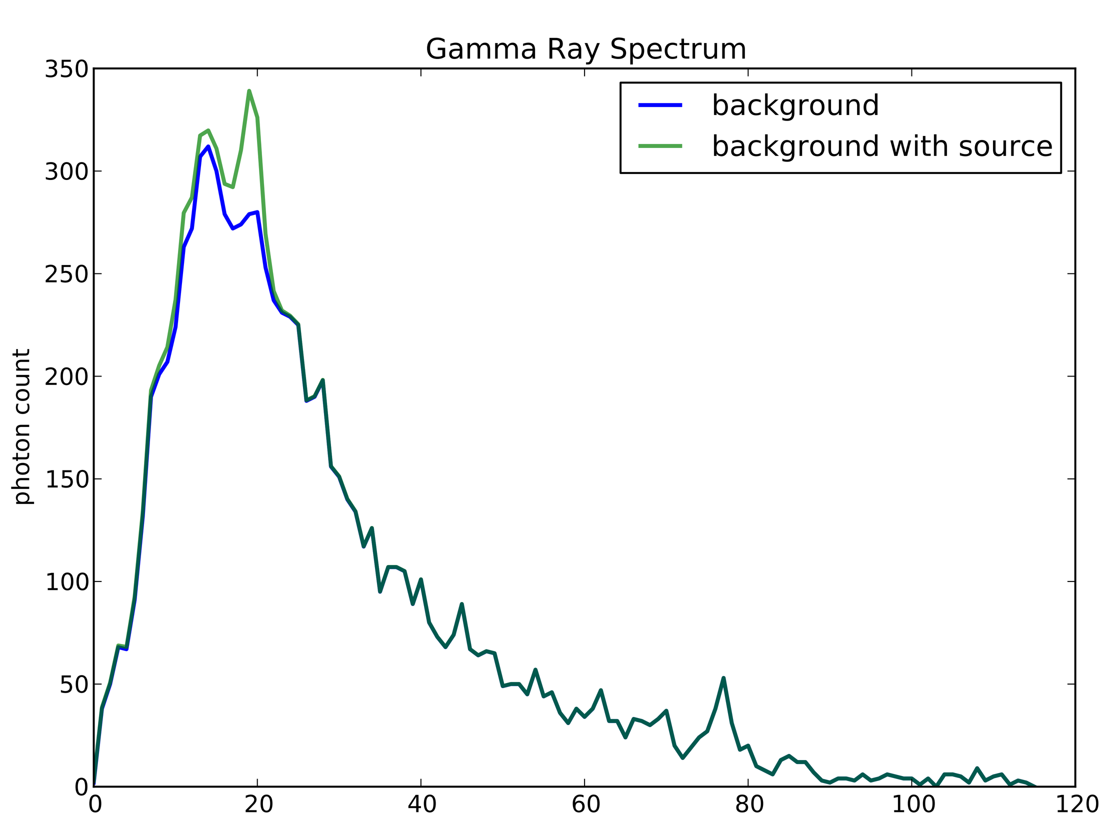

A typical spectroscopic measurement, with and without source injected.

Bayesian Aggregation Framework

Bayesian Aggregation is a framework for leveraging evidence from multiple observations $\{x_i\}_{i=1}^n$. For source detection, it is defined as

$$BA(\{x_i\}_{i=1}^n) = \prod_{i=1}^n \frac{P(x_i|\text{source})}{P(x_i|\text{no source})}.$$

We define a vector ζ that captures uncertain physical factors of interest such as source intensity, distance to source, and gamma-ray attenuation due to occlusion. With the assumption that the components of ζ are independent given source intensity I, **we propose two novel variations of the BA score**: one that maximizes over ζ and one that marginalizes over ζ.

$$BA_{\text{max}}(X) = \max_{I\in\mathbf{R}} P(I) \prod_{i=1}^n \max_{\zeta_i\in\Omega} \frac{P(\zeta_i|I)P(x_i|\text{source},\zeta_i)}{P(x_i|\text{no source},\zeta_i)}$$

$$BA_{\text{marg}}(X) = \int_{\mathbf{R}} P(I) \prod_{i=1}^n \int_{\Omega} \frac{P(\zeta_i|I)P(x_i|\text{source},\zeta_i)}{P(x_i|\text{no source},\zeta_i)} d\zeta_i \, dI$$

Matched Filter Likelihood Model

We use matched filter [2], the linear model that maximizes signal to noise ratio for known source templates, as the basis of our likelihood model. It is defined as $MF(x) = s^T\Sigma^{-1}x$, where s is the source spectrum and Σ the background covariance from training data. Since this is a weighted sum of Poisson random variables, it is approximately normally distributed; we thus define a likelihood model

$$\frac{P(x|\text{source},\zeta)}{P(x|\text{no source},\zeta)} = \sqrt{\frac{\sigma_b^2}{\sigma_b^2+\zeta\sigma_s^2}} \exp\left(\frac{(x-\mu_b-\zeta\mu_s)^2}{2(\sigma_b^2+\zeta\sigma_s^2)} - \frac{(x-\mu_b)^2}{2\sigma_b^2}\right),$$

where

$$\mu_s = w\cdot s \qquad \sigma_s^2 = (w\circ w)\cdot s$$
$$\mu_b = w\cdot B \qquad \sigma_b^2 = (w\circ w)\cdot B$$

with $w = \Sigma^{-1}s$ and B the background rate vector.

Priors

We let $P(I)$ be a very weak L1 regularization centered at 0, and we let $P(\zeta_i|I)$ be a L1 regularization centered at some estimate $\hat{\zeta}$ based on a noisy estimate of sensor position, e.g. a GPS reading. There is much flexibility in altering $P(\zeta_i|I)$ to suit the available auxiliary data.

Boundary of Improvement

We derive a first order asymptotic approximation of the maximum difference between the true positive rates of the oracle (fully informed model) and baseline BA at a false positive rate F_0. It is given as

$$\max T - \hat{T} = \text{erf}\left(\text{erf}^{-1}(1-2F_0)\frac{1-\cos\theta}{1+\cos\theta}\right),$$

where θ is a measure of the difference between the true and estimated parameters of the scene and is approximately the angle between $\hat{\zeta}$ and the true ζ.

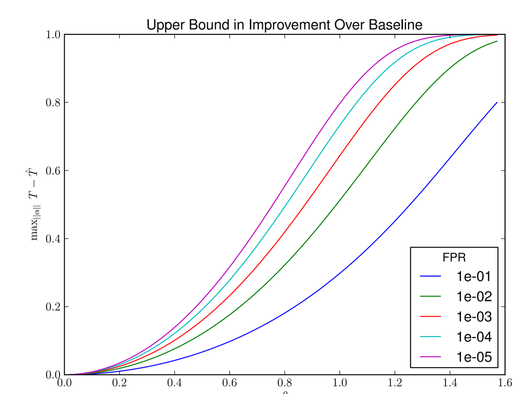

Results

The following plots show improvement over the baseline at a false positive rate of 10^{-3} for various scenes and source and background specifications.

Each model with several regularization penalties:

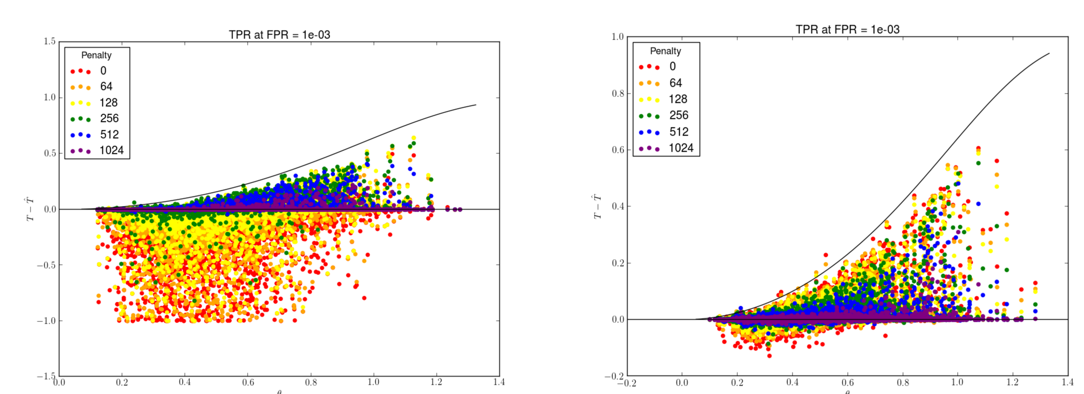

The oracle, the maximization model with penalty 512, and the marginalization model with penalty 256:

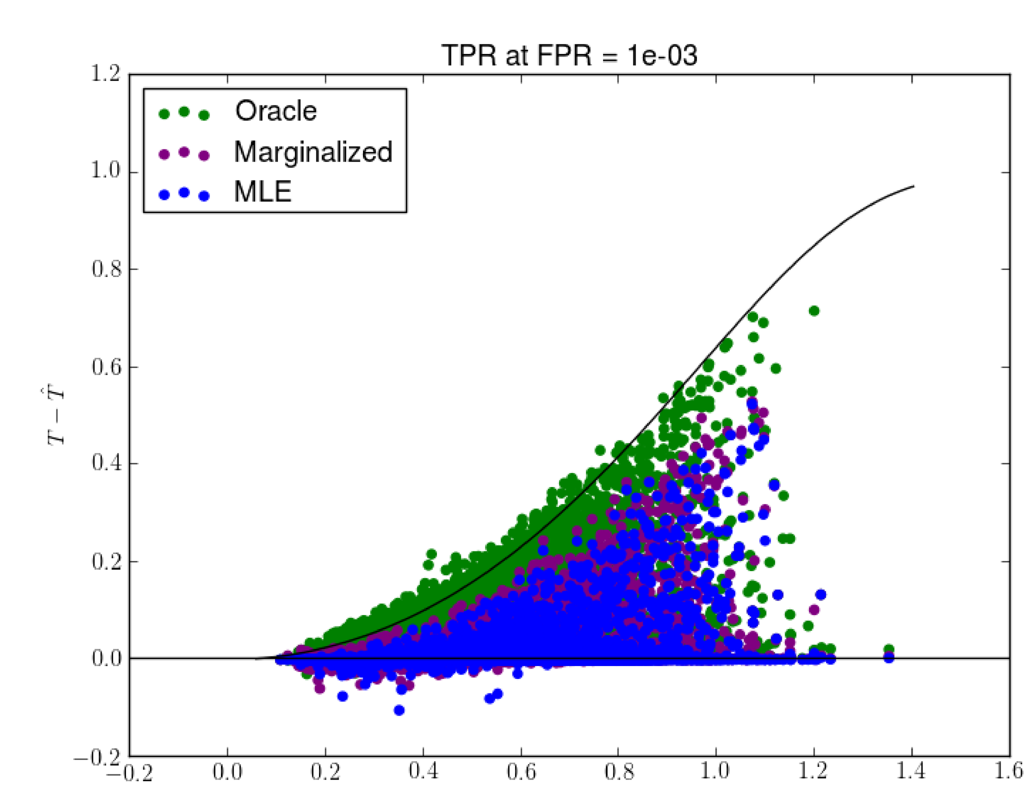

With a tuned regularization, both models have increased performance over the baseline under various levels of uncertainty, and even demonstrate true positive rates close to the oracle.

The difference between the two models:

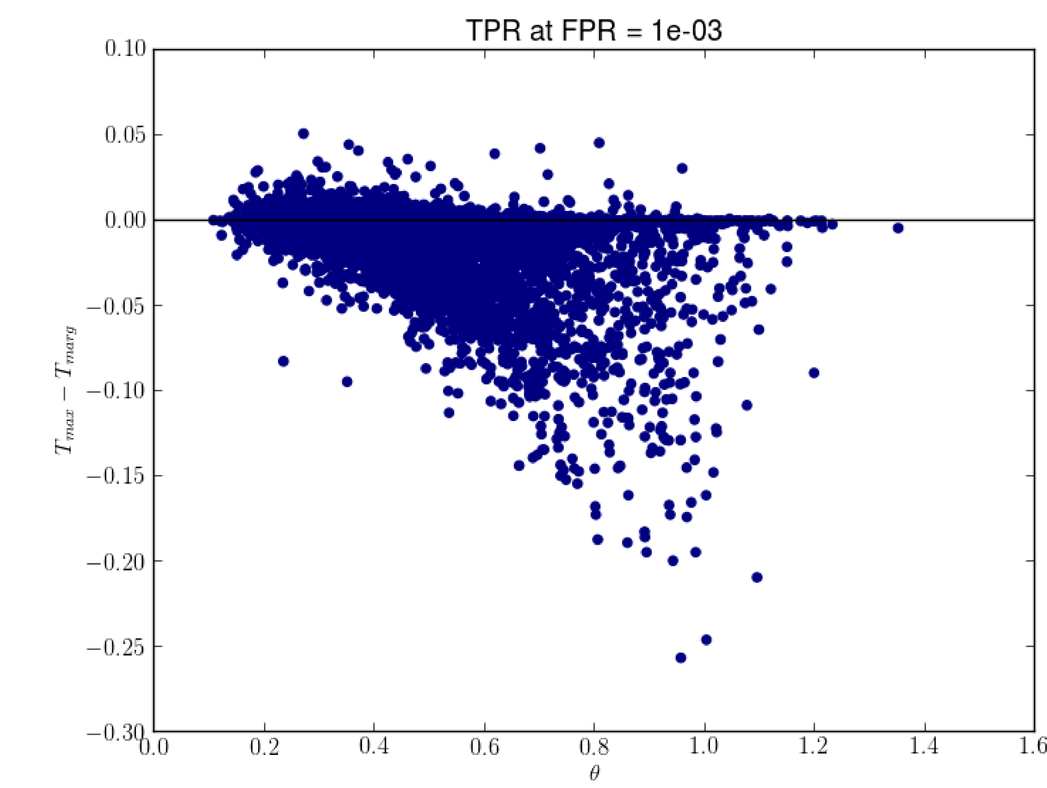

Conclusion

We find that our models are very effective in overcoming targeted sources of uncertainty; both show improvement over baseline BA and come close to the performance of the oracle. **We recommend the marginalization approach, since it shows comparatively better performance and is less sensitive to regularization tuning**, with the caveat that it is more computationally expensive than maximization. Future work will integrate with video and computer vision techniques to reduce uncertainty and provide a narrower hypothesis space.

References

[1] P. Tandon, P. Huggins, R. Maclachlan, A. Dubrawski, K. Nelson, and S. Labov. Detection of radioactive sources in urban scenes using bayesian aggregation of data from mobile spectrometers. *Inf. Syst.*, 57(C):195–206, Apr. 2016.

[2] G. Turin. An introduction to matched filters. *IRE Transactions on Information Theory*, 6(3):311–329, June 1960.

Acknowledgements

Jack Good is funded by the Carnegie Mellon University Robotics Institute REU Site (NSF Grant #1659774). He thanks Kyle Miller and Artur Dubrawski for their mentorship.

Carnegie Mellon University

THE ROBOTICS INSTITUTE

NSF