BAYESIAN DETECTION OF SOLID SUBPIXEL TARGETS

James Theiler,[1] *Stefania Matteoli,*[2] *and Amanda Ziemann*[1]

[1]Los Alamos National Laboratory, Los Alamos, NM 87545 USA
[2]National Research Council (CNR) of Italy, Institute of Electronics,
Information Engineering and Telecommunications, Via G. Caruso 16, 56122, Pisa, Italy

ABSTRACT

We implement and evaluate a Bayesian detector for opaque subpixel hyperspectral targets of unknown abundance. Using both simulated and real hyperspectral backgrounds, we compare this detector to the more conventional generalized likelihood ratio test (GLRT) approach, identifying theoretical differences and observing numerical similarities. Among the theoretical advantages provided by the Bayesian detector is *admissibility*, which means that no detector can be uniformly superior to it. Potential disadvantages include the need to choose a prior distribution, and the computation required to integrate that distribution. For solid subpixel targets, the uniform prior is a natural choice, and we find that adequately-accurate numerical integration can be achieved with only a few evaluations of the likelihood function. We show results for targets implanted in both simulated and real data.

Index Terms— Algorithm, Hyperspectral imagery, Target detection, Likelihood ratio, Composite hypothesis testing, Bayes, GLRT, Multivariate t distribution

You often learn more by being wrong for the right reasons than right for the wrong reasons.
-— Norton Juster, *The Phantom Tollbooth*

1. TARGET DETECTION

For an opaque subpixel hyperspectral target with reflectance spectrum given by $\mathbf{t}$, subpixel abundance a (defined as the fraction of the pixel covered by the target), and atop a background spectrum $\mathbf{z}$, the observed mixed spectrum will be given by the replacement model [1]:

$$\mathbf{x} = (1-a)\mathbf{z} + a\mathbf{t}. \tag{1}$$

Under the null hypothesis, the target is absent and $\mathbf{x} = \mathbf{z}$. If the target abundance a were known in advance, we could derive the optimal detector in terms of the likelihood ratio:

$$L(\mathbf{x}) = \frac{P_{\text{target}}(\mathbf{x})}{P_{\text{bkg}}(\mathbf{x})} = \frac{(1-a)^{-d} P_{\text{bkg}}\left(\dfrac{\mathbf{x}-a\mathbf{t}}{1-a}\right)}{P_{\text{bkg}}(\mathbf{x})}. \tag{2}$$

For the multivariate t distribution with mean $\boldsymbol{\mu}$ and covariance R, we use:

$$P_{\text{bkg}}(\mathbf{x}) = \left[(\nu-2) + (\mathbf{x}-\boldsymbol{\mu})' R^{-1} (\mathbf{x}-\boldsymbol{\mu})\right]^{-\frac{d+\nu}{2}}, \tag{3}$$

while for the non-parameteric model, we use a variable-width kernel density estimator:

$$P_{\text{bkg}}(\mathbf{x}) = \sum_{n=1}^{N} \frac{1}{r_k^d(\mathbf{y}_n)} \kappa\left(\frac{\mathbf{x}-\mathbf{y}_n}{r_k(\mathbf{y}_n)}\right), \tag{4}$$

where $\{\mathbf{y}_n\}$ is a sample of N pixels from the image, and $r_k(\mathbf{y}_n)$ corresponds to the distance from $\mathbf{y}_n$ to its kth nearest neighbor among that sample. Following Ref. [2], we use the Epanechnikov kernel, take $k = 75$, and let N correspond to all pixels in the image except the pixel under test.

1.1. GLRT DETECTOR

Because the target abundance is unknown *a priori*, however, this is a composite hypothesis testing problem, and there is no single optimum solution. The traditional approach is to employ the generalized likelihood ratio test (GLRT). That is: find $\widehat{a}$ that maximizes the likelihood in Eq. (2), and use $L(\widehat{a}, \mathbf{x})$ as the detector. That is:

$$\mathcal{D}_{\text{GLRT}}(\mathbf{x}) = \max_a L(a, \mathbf{x}) = L(\widehat{a}, \mathbf{x}). \tag{5}$$

For the likelihood function in Eqs. (2,3), which corresponds to the replacement model with an elliptically-contoured multivariate t-distributed background, it turns out that it is possible to solve for $\widehat{a}$ analytically and to thus obtain a closed-form solution for the detector in Eq. (5) [3, 4]. For different models and/or more complicated backgrounds, however, a closed-form solution may not be possible.

1.2. BAYESIAN DETECTOR

The Bayesian approach is, first, to posit a prior $p_a(a)$ which corresponds to what we might know about the target abundance *a priori*; that is, *before* we see the data. Then, instead

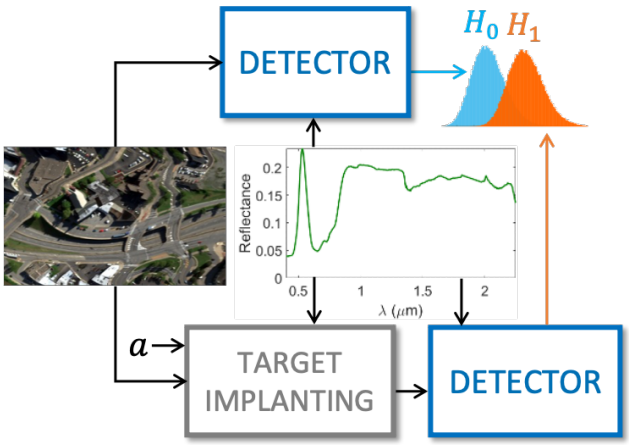

Fig. 1: In the matched-pair method of implanting targets, two images are utilized. The first is treated as background (corresponding to null hypothesis H_0), and the second is a copy of the first but with targets implanted at a specified abundance a into every pixel (corresponding to alternative hypothesis H_1). The detector is applied to both images, forming the two histograms of detector values. In some experiments here, we used the urban highway image as background, and in other experiments, we simulated the background with draws from an EC distribution. Note that this figure also shows the RIT SHARE 2012 image used in the numerical experiments, as well as the target spectrum that was implanted.

of taking the peak value of the likelihood function, we take the weighted average over the range $0 \leq a \leq 1$:

$$\mathcal{D}_{\text{Bayes}}(\mathbf{x}) = \int_0^1 L(a, \mathbf{x}) p(a) da. \tag{6}$$

A natural choice, given that we do not have *a priori* information about target abundance, is to use the uniform prior: $p(a) = 1$. With this choice, we have *not* obtained a closed form solution for $\mathcal{D}_{\text{Bayes}}(\mathbf{x})$. But we have found that an adequate estimate for this integral can be obtained numerically, with relatively few evaluations[1] of $L(a, \mathbf{x})$.

Thus the GLRT and Bayesian approaches *both* consider the likelihood curve. Informally, we might say that a "larger likelihood" indicates a higher confidence in target presence – and while the GLRT looks at the peak of the curve, the Bayesian detector looks at the area under it.

1.3. Admissibility

Following Kay [6] (and more formally, Lehmann and Romano [7]), we say that one detector is more *powerful* than another if it produces an unambiguously better Receiver-Operator Characteristic (ROC) curve – that is to say: the detection rate is better over the full range of false alarm rates. A detector is *uniformly more powerful* if it is more powerful for every value of a. The holy grail of composite hypothesis testing is to find a *uniformly most powerful* (UMP) detector, but most composite hypothesis testing problems – including this one – do not admit a UMP solution. For these problems, no one detector is always best. The goal for the practitioner, then, is to find a detector that is best for the scenario(s) of interest.

One of our motivations for investigating the Bayesian detector is that Bayesian detectors are known to be *admissible* [6, 7]. A detector is admissible if no other detector is uniformly more powerful. This is a good thing – if a detector is not admissible, that means there's an unambiguously better detector out there, somewhere. It does *not* follow, however, that an admissible detector is uniformly more powerful than all other detectors.

As will be shown in the numerical experiments, the admissible Bayesian detector is not always better than the *ad hoc* GLRT detector. It very often achieves performance that is nearly identical to the GLRT, and in some cases, particularly in the low false alarm rate regime, it is the GLRT that is better.

[1] Results here are based on only five values: $a \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$. Furthermore, since detections are based on likelihood *ratios*, we do not need *absolute* accuracy. Another way to think of it is that we are computing the exact Bayesian detector for a slightly different prior (sum of five delta functions instead of uniform); it may not exactly agree with the result obtained with a uniform prior, but it is still guaranteed to be admissible. Indeed, reasonable detectors can be derived from Bayesian priors that are comprised of a single well-chosen delta function [5].

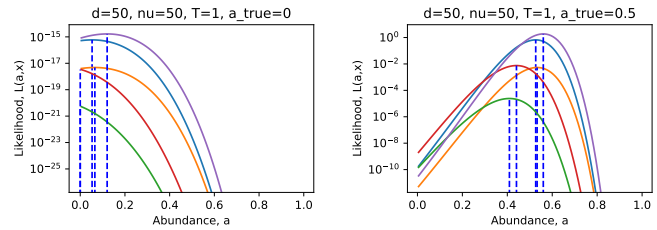

Fig. 2: Likelihood as a function of a, for true abundance values $a = 0$ (left) and $a = 0.5$ (right). We sampled five points at random (the same five points in both plots – one with and one without target present) and show how likelihood varies with a. The vertical dashed lines correspond to GLRT values $\widehat{a}$ for each of the five points. Note that the points with target *present* (in the right panel) have larger likelihoods than the points with target *absent* (in the left panel).

1.4. Other advantages

In addition to the theoretical advantages, there are also some practical advantages to Bayesian detectors.

One advantage is straightforwardness. Once we are outside the realm of trying to find a closed-form solution, the integral in Eq. (6) is straightforward to numerically estimate. (We note that numerical estimates of the peak can also be done straightforwardly.)

Another advantage arises when the background distribution does not admit an analytical model. In that case, matched-pair machine learning [8, 9] can be used to imple-

ment the detector by first implanting targets of abundances drawn from the prior and then by using binary classification to distinguish background from background+target.

2. EXPERIMENT AND RESULTS

2.1. Fully simulated data

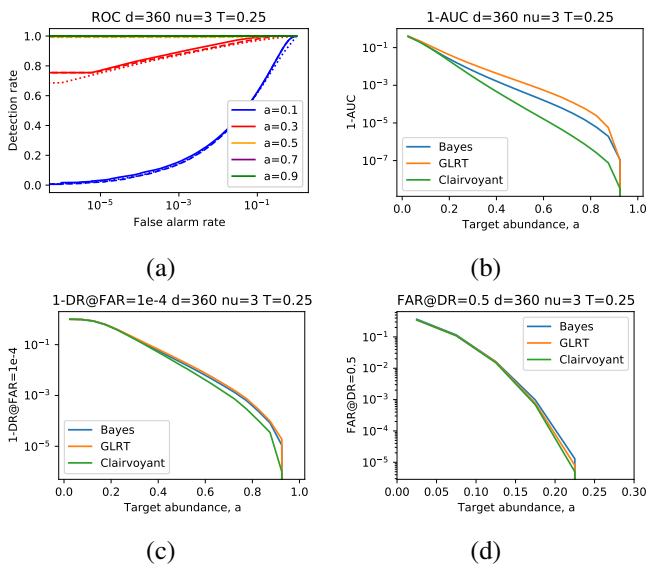

Fig. 3: Comparison of GLRT and Bayesian detectors in detecting implanted plumes on a simulated multivariate t-distributed background, with parameters $d = 360$, $\nu = 3$, and target strength $T = 0.25$. (a) ROC curves ares shown for GLRT (dotted lines), Bayesian (dashed lines), and clairvoyant detectors (solid lines). From these ROC curves, single scalar parameters can be derived. In (b), the AUC (area under the ROC curve; actually one minus AUC) is plotted as a function of a for all three detectors; as expected, the clairvoyant detector has the best performance (highest AUC, or lowest 1-AUC), and we observe that the Bayesian detector slightly outperforms the GLRT detector. In (c), the detection rate is plotted, and again the clairvoyant is best, while the Bayesian and GLRT are nearly identical. In (d), the FAR (false alarm rate) corresponding to a detection rate of 0.5 is plotted. Here the differences are very small; again the clairvoyant is best, and it appears that the GLRT just edges out the Bayesian detector.

Although the ultimate *proof of the pudding* is performance in real target detection scenarios, direct comparisons of algorithms in such scenarios are inevitably anecdotal, and it is important to understand the behavior of those algorithms as conditions change. In this section, we consider an elliptically-contoured (EC) multivariate t-distributed background, with targets implanted according to Eq. (1). In this idealized setting, we can consider whitened data, which means that we can assume the background data has zero mean and unit covariance.

To assess performance of different algorithms (and our main interest here is the GLRT and Bayesian variants of the EC-FTMF algorithm), the parameters we need to consider are: spectral dimension d, EC non-Gaussianity parameter ν, target abundance a, and a measure of target strength, $T = (\mathbf{t} - \boldsymbol{\mu})' R^{-1} (\mathbf{t} - \boldsymbol{\mu})/d$.

In Fig. 2, we illustrate some typical likelihood functions, $L(a, \mathbf{x})$ plotted as a function of a. The GLRT detector is based on peaks of these curves, while the Bayesian detector is based on the areas of these curves.

In Fig. 3, we compare the GLRT and the Bayesian detectors for plausible parameters $d = 360$, $\nu = 3$, and $T = 0.25$. In this case we see different performance when using the AUC statistic; the Bayesian detector does a better job over all values of a. The FAR@DR=0.5 however shows nearly identical performance, with the GLRT marginally better. We remark that for practical target detection, low false alarm rates are required, and the AUC does a poor job of characterizing the low false alarm rate regime.

2.2. Real data with real targets artificially implanted

The dataset employed in this study uses an image from the RIT SHARE 2012 campaign acquired by a SpecTIR sensor over downtown Rochester, NY [10, 11], shown in Fig. 1. In order to draw ROC curves, we employed the target implantation method [8, 12, 13], sketched in Fig. 1. Basically, we fractionally implanted a selected target spectrum in each pixel of the image with a given target abundance, according to the replacement model of Eq. (1). The detectors were then applied to both the original (target-absent) image and the treated (target-present) image using the implanted target spectrum. The corresponding detection statistics were thresholded for ROC curve evaluation. The target spectrum was acquired over a green wooden panel during the RIT SHARE 2012 campaign. In order to make the target more difficult to detect, here we chose to implant an attenuated target: $\mathbf{t}' = (1-f)\boldsymbol{\mu} + f\mathbf{t}$, with $f = 0.05$. We observe in Fig. 4 that the Bayesian detector does outperform the GLRT detector, but the more important observation is that the difference is miniscule. Indeed, both detectors also outperform the clairvoyant detector (again, by a very small amount). Because the clairvoyant detector's optimality depends on the background distribution being multivariate t, this optimality is no longer guaranteed when applied to real data. This example demonstrates that GLRT vs Bayes is not the whole story; modeling the background distribution is always important [14].

In Fig. 4(f), we also compute ROC curves using both Bayesian and GLRT-based non-parametric detectors derived from the Variable-bandwidth Kernel Density Estimator (VKDE) [2]. We observe that the non-parametric detector achieves higher detection rates when the false alarm rate is also moderately large, but that the EC appears to do a better job at the very lowest false alarm rates.

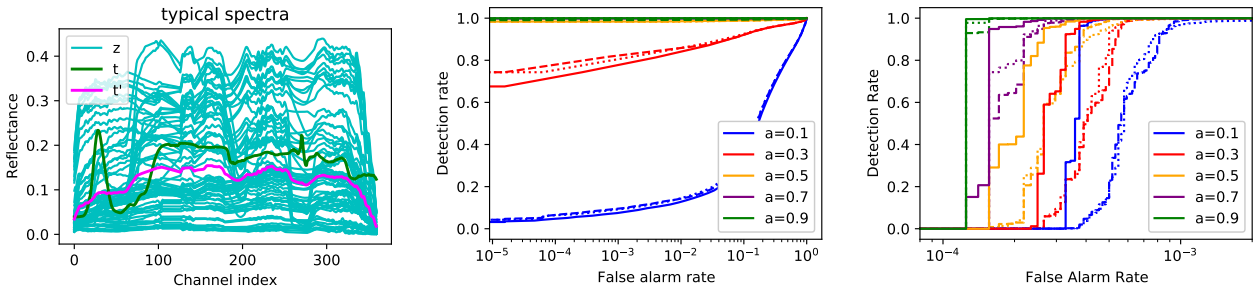

(a) (b) (c)

Fig. 4: (a) Typical spectra from the image of downtown Rochester, along with the original target spectrum $\mathbf{t}$, and the attenuated spectrum $\mathbf{t}'$ that was implanted into the image. Performance of Bayesian (dashed lines) and GLRT (dotted lines) detectors are compared to the Clairvoyant detector (solid lines): (b) using the EC-model of background in Eq. (3), and (c) using the non-parameteric model of background in Eq. (4). We see that Bayesian and GLRT detectors have very similar performance, and that that performance is not much different from that of the clairvoyant detector. On the other hand, we see a considerable difference in the two background models. The EC-based model appears to be better at very low false alarms, while the NP-based detector is better for false alarm rates larger than 0.001.

3. DISCUSSION AND CONCLUSIONS

Based on our numerical experiments, we found that the admissibility advantage nominally provided by Bayesian detectors did not always lead to substantially better performance in parameter regimes that correspond to practical target detection. Indeed, in the low false alarm rate regime, we often observed the GLRT very competitive with (and in some cases, even outperforming) the Bayesian detector. Even here, however, we found that the Bayesian and GLRT performances were not only very similar to each other, but were very nearly as good as the (formally) unattainable optimum provided by the clairvoyant detector.

We remark that the experiments here were based on the uniform prior. Future experiments with priors that put more weight on smaller values of a may prove to be more competitive. The potential flexibility of the prior may enable us to engineer detectors that optimize properties that we can specify in advance.

4. REFERENCES

[1] D. Manolakis, "Taxonomy of detection algorithms for hyperspectral imaging applications," *Optical Engineering*, vol. 44, pp. 066403, 2005.

[2] S. Matteoli, M. Diani, and G. Corsini, "Closed-form non-parametric GLRT detector for subpixel targets in hyperspectral images," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 56, no. 2, pp. 1568–1581, 2020.

[3] A. Schaum and A. Stocker, "Spectrally selective target detection," *Proc. ISSSR (International Symposium on Spectral Sensing Research)*, p. 23, 1997.

[4] J. Theiler, B. Zimmer, and A. Ziemann, "Closed-form detector for solid sub-pixel targets in multivariate t-distributed background clutter," in *Proc. IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, 2018, pp. 2773–2776.

[5] J. Theiler, "Veritas: an admissible detector for targets of unknown strength," *Proc. SPIE*, vol. 11727, pp. 117270B, 2021.

[6] S. M. Kay, *Fundamentals of Statistical Signal Processing: Detection Theory*, vol. II, Prentice Hall, New Jersey, 1998.

[7] E. L. Lehmann and J. P. Romano, *Testing Statistical Hypotheses*, Springer, New York, 2005.

[8] J. Theiler, "Matched-pair machine learning," *Technometrics*, vol. 55, pp. 536–547, 2013.

[9] A. Ziemann, M. Kucer, and J. Theiler, "A machine learning approach to hyperspectral detection of solid targets," *Proc. SPIE*, vol. 10644, pp. 1064404, 2018.

[10] A. Giannandrea, N. Raqueno, D. W. Messinger, J. Faulring, J. P. Kerekes, J. van Aardt, K. Canham, S. Hagstrom, E. Ontiveros, A. Gerace, J. Kaufman, K. Vongsy, H. Griffith, and B. D. Bartlett, "The SHARE 2012 data campaign," *Proc. SPIE*, vol. 8743, May 2013.

[11] "SpecTIR Hyperspectral Airborne Experiment 2012 (SHARE 2012)," `https://www.rit.edu/cos/share2012/`.

[12] M. Bar-Tal and S. R. Rotman, "Performance measurement in point target detection," *Infrared Physics & Technology*, vol. 37, pp. 231–238, 1996.

[13] M.S. Stefanou and J.P. Kerekes, "A method for assessing spectral image utility," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 47, no. 6, pp. 1698–1706, 2009.

[14] S. Matteoli, M. Diani, and J. Theiler, "An overview of background modeling for detection of targets and anomalies in hyperspectral remotely sensed imagery," *J. Selected Topics in Applied Earth Observations and Remote Sensing (JSTARS)*, vol. 7, pp. 2317–2336, 2014.