Effect of signal contamination in matched-filter detection of the signal on a cluttered background

James Theiler and Bernard R. Foy

Abstract— To derive a matched filter for detecting a weak target signal in a hyperspectral image, an estimate of the band-to-band covariance of the target-free background scene is required. We investigate the effects of including some of the target signal in the background scene. Although the *covariance* is contaminated by the presence of target signal (there is increased variance in the direction of the target signature), we find that the *matched filter* is not necessarily affected. In fact, if the variation in plume strength is strictly uncorrelated with the variation in background spectra, the matched filter and its signal-to-clutter ratio (SCR) performance will not be impaired. While there is little *a priori* reason to expect significant correlation between the plume and the background, there usually is some residual correlation, and this correlation leads to a suppressing effect that limits the SCR obtainable even for strong plumes.

These effects are described and quantified analytically, and the crucial role of this correlation is illustrated with some numerical examples using simulated plumes superimposed on real hyperspectral imagery. In one example, we observe an order of magnitude loss in SCR for a matched filter based on the contaminated covariance.

Index Terms— Adaptive signal detection, Clutter, Covariance matrices, Gases, Hyperspectral imagery, Matched filters

I. INTRODUCTION

DETECTING a weak signal on a cluttered background is a problem that arises in a variety of contexts, but we will consider in particular the detection of gaseous plumes in hyperspectral imagery. In the weak plume limit, the effect of the plume can be represented as a superposition of signal and background

$$\mathbf{r} = \epsilon \mathbf{b} + \mathbf{z} \qquad (1)$$

where $\mathbf{r}$ is the observed radiance, $\mathbf{b}$ is the *known* plume signature, and $\mathbf{z}$ is the *unknown* background clutter. Eq. (1) corresponds to the radiance in a single pixel; here, $\mathbf{r}$, $\mathbf{b}$, and $\mathbf{z}$ are d-dimensional vectors corresponding to the d channels of the hyperspectral image, and ϵ is a scalar proportional to the plume strength. In the thermal infrared, ϵ further depends on the thermal contrast and ground emissivity, and it can be positive or negative depending on whether the plume is in emission or absorption. In shorter visible and near-infrared (VNIR) wavelengths, the plume is only absorptive. We take ϵ to be independent of wavelength, which turns out to be a fair approximation in most cases [1]. For the exposition here, we assume that Eq. (1) holds and that ϵ is a scalar, though it varies from pixel to pixel. Since the magnitude of the plume

signature $\mathbf{b}$ can be incorporated into the plume strength ϵ in Eq. (1), we will impose the condition that $\mathbf{b}^T\mathbf{b} = 1$.

A matched filter is a d-dimensional vector $\mathbf{q}$ that is applied to a pixel vector $\mathbf{r}$ to produce a scalar $\mathbf{q}^T\mathbf{r}$ which emphasizes the signal while suppressing the background clutter; that is, $\mathbf{q}^T\mathbf{b}$ is large, but $\mathbf{q}^T\mathbf{z}$ is small.

The adaptive matched filter [2] can be derived by optimizing an estimated signal-to-clutter ratio. We first estimate the background statistics for the scene without the plume, computing the average $\langle\mathbf{z}\rangle = (1/N)\sum_{i=1}^{N}\mathbf{z}_i$, and the covariance matrix

$$K_o = (1/N)\sum_{i=1}^{N}(\mathbf{z}_i - \langle\mathbf{z}\rangle)(\mathbf{z}_i - \langle\mathbf{z}\rangle)^T, \qquad (2)$$

where $\mathbf{z}_i$ is off-plume background radiance in the ith pixel, and N is the number of pixels in the scene. From this covariance matrix, we can express the ability of the matched filter to suppress the background in terms of the variance of the matched filter image over the (plume-free) scene. That is:

$$
\begin{aligned}
\mathrm{Var}(\mathbf{q}^T\mathbf{z}) &= (1/N)\sum_{i=1}^{N}|\mathbf{q}^T(\mathbf{z}_i - \langle\mathbf{z}\rangle)|^2 \\
&= (1/N)\sum_{i=1}^{N}\mathbf{q}^T(\mathbf{z}_i - \langle\mathbf{z}\rangle)(\mathbf{z}_i - \langle\mathbf{z}\rangle)^T\mathbf{q} \\
&= \mathbf{q}^T K_o \mathbf{q}. \qquad (3)
\end{aligned}
$$

We can express the performance of the matched filter in terms of a signal-to-clutter ratio. The signal is the matched filter $\mathbf{q}$ applied to the plume $\epsilon\mathbf{b}$, and the clutter is given by the variance over the background.

$$\mathrm{SCR} = \frac{|\mathbf{q}^T(\epsilon\mathbf{b})|^2}{\mathbf{q}^T K_o \mathbf{q}} = \frac{\epsilon^2|\mathbf{q}^T\mathbf{b}|^2}{\mathbf{q}^T K_o \mathbf{q}} = \frac{\epsilon^2\mathbf{q}^T\mathbf{b}\mathbf{b}^T\mathbf{q}}{\mathbf{q}^T K_o \mathbf{q}}. \qquad (4)$$

The SCR is optimized by a matched filter given by

$$\mathbf{q}_o \propto K_o^{-1}\mathbf{b}, \qquad (5)$$

where '$\propto$' indicates proportionality.[1] The SCR for this matched filter is given by plugging Eq. (5) into Eq. (4):

$$\mathrm{SCR} = \frac{\epsilon^2\mathbf{b}^T K_o^{-1}\mathbf{b}\mathbf{b}^T K_o^{-1}\mathbf{b}}{\mathbf{b}^T K_o^{-1}\mathbf{b}} = \epsilon^2\mathbf{b}^T K_o^{-1}\mathbf{b}. \qquad (6)$$

The expression in Eq. (5) can also be derived in terms of the generalized likelihood ratio test (GLRT) [3], [4], and

---

[1] We say that vector $\mathbf{u}$ is proportional to vector $\mathbf{v}$ (that is, $\mathbf{u} \propto \mathbf{v}$) if there exists a scalar a such that $\mathbf{u} = a\mathbf{v}$. While there has traditionally been some concern about appropriate normalization (for instance, to produce a fixed false alarm rate), we remark that the SCR is unaffected by the magnitude of the matched filter, so proportionality is adequate for our purposes.

for gaussian data can be shown to be optimal [5]. When the background is nongaussian [6], [7], [8], then other approaches may be called for. But for this paper, we will assume that the background is gaussian and that its statistical error is negligible. The general problem of weak signal detection in hyperspectral imagery is reviewed in Refs. [9] and [10].

Gerlach [11] considered the effect of having one of the samples $\mathbf{z}_i$ contaminated by the target signal. In Gerlach's treatment, consideration was given to how this contaminant affected the threshold for constant false alarm rate (CFAR) detection. Our concern here is with how the effect can alter the *direction* of the matched filter $\mathbf{q}$.

In Eq. (5), we derived the classic matched filter result, when the background is uncontaminated by plume. In Section II, we include the plume in the estimate of the covariance, and show that as long as the plume is uncorrelated with its background, it will have no effect on the matched filter. In Section III, we relax the assumption that plume and background are uncorrelated; we introduce the vector ζ to characterize this correlation, and show how ζ modifies the matched filter $\mathbf{q}$. In Section IV, we show the effect of this modified matched filter on the signal-to-clutter ratio. Finally in Section V, we illustrate these effects on real data with simulated plumes.

II. EFFECT OF INCLUDING ON-PLUME PIXELS

In practice, we measure the *total* radiance $\mathbf{r}$ which includes plume as well as background. We compute the mean $\langle \mathbf{r} \rangle = (1/N) \sum_{i=1}^{N} \mathbf{r}_i$, and covariance matrix

$$K = (1/N) \sum_{i=1}^{N} (\mathbf{r}_i - \langle \mathbf{r} \rangle)(\mathbf{r}_i - \langle \mathbf{r} \rangle)^T, \quad (7)$$

where $\mathbf{r}_i = \epsilon_i \mathbf{b} + \mathbf{z}_i$ is the radiance in the ith pixel. Note that $\langle \mathbf{r} \rangle = \langle \epsilon \rangle \mathbf{b} + \langle \mathbf{z} \rangle$, where $\langle \epsilon \rangle = (1/N) \sum_{i=1}^{N} \epsilon_i$, and the expression for covariance is

$$
\begin{aligned}
K = \quad (1/N) \sum_{i=1}^{N} \big[\ & (\epsilon_i - \langle \epsilon \rangle)^2 \mathbf{b}\mathbf{b}^T + \\
& (\epsilon_i - \langle \epsilon \rangle)\mathbf{b}(\mathbf{z}_i - \langle \mathbf{z} \rangle)^T + \\
& (\epsilon_i - \langle \epsilon \rangle)(\mathbf{z}_i - \langle \mathbf{z} \rangle)\mathbf{b}^T + \\
& (\mathbf{z}_i - \langle \mathbf{z} \rangle)(\mathbf{z}_i - \langle \mathbf{z} \rangle)^T \ \big].
\end{aligned} \quad (8)
$$

Now, here comes an important assumption: the plume strength is uncorrelated with the underlying background; *i.e.,*

$$(1/N) \sum_{i=1}^{N} (\epsilon_i - \langle \epsilon \rangle)(\mathbf{z}_i - \langle \mathbf{z} \rangle) = \mathbf{0}. \quad (9)$$

We will write

$$\varepsilon^2 = (1/N) \sum_{i=1}^{N} (\epsilon_i - \langle \epsilon \rangle)^2 \quad (10)$$

as the variance of the plume concentration ϵ over the scene. From Eq. (9) and Eq. (10), we can write an expression which shows how the covariance matrix is altered by having the plume in the scene:

$$K = K_o + \varepsilon^2 \mathbf{b}\mathbf{b}^T. \quad (11)$$

The "corrupted" matched filter is given by $\mathbf{q} \propto K^{-1}\mathbf{b}$ (compare to Eq. (5)), which is equivalent to $K\mathbf{q} \propto \mathbf{b}$, or

$$K_o\mathbf{q} + \varepsilon^2 \mathbf{b}\mathbf{b}^T\mathbf{q} \ \propto \ \mathbf{b}. \quad (12)$$

Since $\mathbf{b}^T\mathbf{q}$ is a scalar, we have that the term $\varepsilon^2 \mathbf{b}\mathbf{b}^T\mathbf{q} = \varepsilon^2 (\mathbf{b}^T\mathbf{q})\mathbf{b}$ is proportional to $\mathbf{b}$, and so we have that $K_o\mathbf{q} \propto \mathbf{b}$, or $\mathbf{q} \propto K_o^{-1}\mathbf{b} \propto \mathbf{q}_o$. But since the normalization on the matched filter is unimportant, this means that the matched filter is *unaffected* by the presence of plume in the scene.

This is initially counter-intuitive: adding a plume adds signal variation in the $\mathbf{b}$ direction and leads to a covariance matrix that is "fatter" in that direction. Informally, the effect of the K^{-1} in Eq. (5) is to de-emphasize the directions where K is fatter (in particular, the eigendirections associated with the largest eigenvalues). However, as the mathematics shows, this does not apply to the specific direction $\mathbf{b}$.

III. CORRELATION OF BACKGROUND AND PLUME

Although it is reasonable to expect the plume and the underlying scene to be more or less uncorrelated, one should not expect the expression in Eq. (9) to be strictly zero. In practice, there should be some residual correlation, and we will define

$$\zeta = (1/N) \sum_{i=1}^{N} \left(\frac{\epsilon_i - \langle \epsilon \rangle}{\varepsilon} \right) (\mathbf{z}_i - \langle \mathbf{z} \rangle), \quad (13)$$

where the denominator ε is defined in Eq. (10) and is used here for normalization with respect to plume strength. Although we no longer neglect the correlation embodied in ζ, we do still expect ζ to be small, since it corresponds to the extent to which the statistics of on-plume pixels differ from the statistics of off-plume pixels.

The expression for covariance in Eq. (8) is now given by

$$K = K_o + \varepsilon \mathbf{b}\zeta^T + \varepsilon \zeta \mathbf{b}^T + \varepsilon^2 \mathbf{b}\mathbf{b}^T. \quad (14)$$

The matched filter that will be obtained using this covariance will satisfy $K\mathbf{q} \propto \mathbf{b}$, so we can write:

$$(K_o + \varepsilon \mathbf{b}\zeta^T + \varepsilon \zeta \mathbf{b}^T + \varepsilon^2 \mathbf{b}\mathbf{b}^T)\mathbf{q} \propto \mathbf{b}. \quad (15)$$

Now two of the terms, $\varepsilon \mathbf{b}\zeta^T\mathbf{q}$ and $\varepsilon^2 \mathbf{b}\mathbf{b}^T\mathbf{q}$, are proportional to $\mathbf{b}$, so

$$K_o\mathbf{q} + \varepsilon \zeta \mathbf{b}^T\mathbf{q} \ \propto \ \mathbf{b}. \quad (16)$$

It can be shown that Eq. (16) is satisfied by[2]

$$\mathbf{q} \propto K_o^{-1} \left(\mathbf{b} - \frac{\varepsilon \mathbf{b}^T K_o^{-1}\mathbf{b}}{1 + \varepsilon \mathbf{b}^T K_o^{-1}\zeta} \zeta \right). \quad (17)$$

The first term is proportional to $\mathbf{q_o}$, but when there is correlation between background and plume strength (embodied in a nonzero ζ), the direction of the matched filter is altered by the presence of plume in the scene.

---

[2]This can be verified by direct substitution, or it can be derived using the well-known Sherman-Morrison-Woodbury formula: $(A + UV^T)^{-1} = A^{-1} - A^{-1}U(I + V^T A^{-1} U)^{-1} V^T A^{-1}$. This formula could also be used to invert the covariance in Eq. (11) and/or in Eq. (14), but the expressions would include unweidly scalar multipliers that are not important for our purposes.
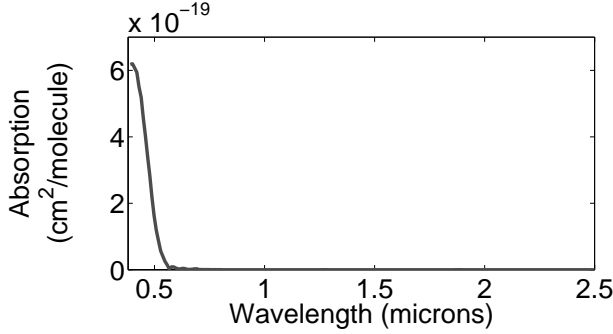
Fig. 1. Absorption spectrum of NO_2 over the AVIRIS wavelength range.
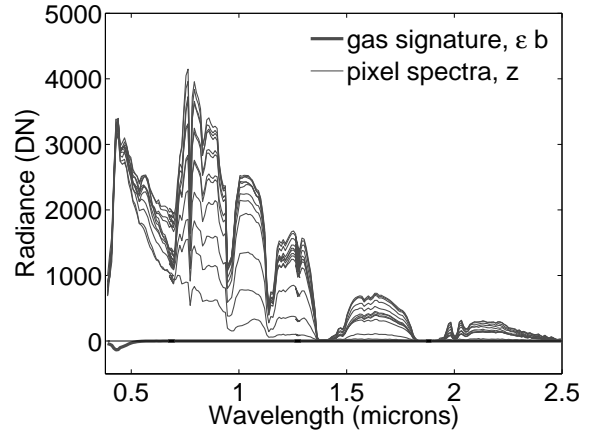

Fig. 2. The positive curves correspond to the spectra z_i of a dozen points randomly selected from the scene. The negative curve is the plume absorption signal εb. The radiance in plotted as digital number (DN), and the ε value shown here corresponds to a plume concentration pathlength of 12.2 ppm-m (parts per million - meters). This is the root mean square plume strength, given by Eq. (10), for our simulated plume.

Since the normalization on $\mathbf{q}$ does not ultimately matter, we will multiply by the scalar $1 + \varepsilon \mathbf{b}^T K_o^{-1} \zeta$ to obtain a form that is strictly linear in ε:

$$\mathbf{q} = \mathbf{q}_o + \varepsilon K_o^{-1} \left([\mathbf{b}^T K_o^{-1} \zeta] \, \mathbf{b} - [\mathbf{b}^T K_o^{-1} \mathbf{b}] \, \zeta \right). \quad (18)$$

IV. EFFECT OF MODIFIED MATCHED FILTER $\mathbf{q}$

We introduce "whitened" [12] vectors

$$\hat{\zeta} = K_o^{-1/2} \zeta, \quad (19)$$
$$\hat{\mathbf{b}} = K_o^{-1/2} \mathbf{b}. \quad (20)$$

to simplify the following scalar expressions:

$$\mathbf{b}^T K_o^{-1} \mathbf{b} = |\hat{\mathbf{b}}|^2, \quad (21)$$
$$\zeta^T K_o^{-1} \zeta = |\hat{\zeta}|^2, \quad (22)$$
$$\mathbf{b}^T K_o^{-1} \zeta = \hat{\mathbf{b}} \cdot \hat{\zeta}. \quad (23)$$

Then, we can rewrite Eq. (18) as

$$\mathbf{q} = \mathbf{q}_o + \varepsilon K_o^{-1/2} \left((\hat{\mathbf{b}} \cdot \hat{\zeta}) \, \hat{\mathbf{b}} - |\hat{\mathbf{b}}|^2 \, \hat{\zeta} \right). \quad (24)$$

From this expression, we have for the signal strength

$$\begin{aligned}
\mathbf{q}^T \mathbf{b} &= \mathbf{q}_o^T \mathbf{b} + \varepsilon \left[K_o^{-1/2} \left((\hat{\mathbf{b}} \cdot \hat{\zeta}) \, \hat{\mathbf{b}} - |\hat{\mathbf{b}}|^2 \, \hat{\zeta} \right) \right]^T K_o^{1/2} \hat{\mathbf{b}} \\
&= [K_o^{-1} \mathbf{b}]^T \mathbf{b} + \varepsilon \left((\hat{\mathbf{b}} \cdot \hat{\zeta}) |\hat{\mathbf{b}}|^2 - |\hat{\mathbf{b}}|^2 (\hat{\mathbf{b}} \cdot \hat{\zeta}) \right) \\
&= \mathbf{b}^T K_o^{-1} \mathbf{b} = |\hat{\mathbf{b}}|^2. \quad (25)
\end{aligned}$$

For the clutter, we have (after some algebraic manipulation)

$$\mathbf{q}^T K_o \mathbf{q} = |\hat{\mathbf{b}}|^2 + \varepsilon^2 |\hat{\mathbf{b}}|^2 \left(|\hat{\mathbf{b}}|^2 |\hat{\zeta}|^2 - |\hat{\mathbf{b}} \cdot \hat{\zeta}|^2 \right), \quad (26)$$

and so the signal-to-clutter ratio is given by

$$\text{SCR} = \frac{\varepsilon^2 |\mathbf{q}^T \mathbf{b}|^2}{\mathbf{q}^T K_o \mathbf{q}} = \frac{\varepsilon^2 |\hat{\mathbf{b}}|^2}{1 + \varepsilon^2 \left(|\hat{\mathbf{b}}|^2 |\hat{\zeta}|^2 - |\hat{\mathbf{b}} \cdot \hat{\zeta}|^2 \right)}. \quad (27)$$

Recall that ϵ is the plume strength at a given pixel, while ε (defined in Eq. (10)) is a global measure of plume strength across the whole image. The ratio ϵ/ε is a dimensionless value that is independent of overall plume strength. In the strong plume (large ε) limit, the SCR in Eq. (27) saturates at

$$\lim_{\varepsilon \to \infty} \text{SCR} = \frac{(\epsilon/\varepsilon)^2}{|\hat{\zeta}|^2} \left(1 - \frac{|\hat{\mathbf{b}} \cdot \hat{\zeta}|^2}{|\hat{\mathbf{b}}|^2 |\hat{\zeta}|^2} \right)^{-1}. \quad (28)$$

This is an upper bound on the SCR that will be obtained for a given plume configuration. Note that as the magnitude of the dimensionless $|\hat{\zeta}|$ increases, the limiting signal-to-clutter ratio drops. From the definitions in Eq. (13) and Eq. (19), we can write a more compact expression:

$$\hat{\zeta} = (1/N) \sum_{i=1}^{N} (\epsilon_i/\varepsilon) \, \hat{\mathbf{z}}_i \quad (29)$$

where $\hat{\mathbf{z}}_i = K_o^{-1/2} (\mathbf{z}_i - \langle \mathbf{z} \rangle)$ is the whitened and mean-subtracted background radiance.

V. NUMERICAL EXAMPLES

We illustrate the effect of including the target in the background with some numerical examples. For the clutter, we use hyperspectral radiance data from the AVIRIS sensor [13]. On this scene, we simulate a plume of NO_2 [14] (see Fig. 1).

Since this data is in the visible and near-infrared (VNIR), the plume is necessarily in absorption. If the absorption as a function of wavelength λ is given by $\tilde{b}(\lambda)$ (plotted in Fig. 1), then the measured radiance in the ith pixel as a function of λ is given by Beer's law:

$$\begin{aligned}
r_i(\lambda) &= z_i(\lambda) \exp(-\epsilon_i \tilde{b}(\lambda)) \quad (30) \\
&= z_i(\lambda) - \epsilon_i \langle z(\lambda) \rangle \tilde{b}(\lambda) \\
&\quad - \epsilon_i [z_i(\lambda) - \langle z(\lambda) \rangle] \tilde{b}(\lambda) + O(\epsilon^2) \quad (31) \\
&\approx z_i(\lambda) - \epsilon_i \langle z(\lambda) \rangle \tilde{b}(\lambda), \quad (32)
\end{aligned}$$

where $\langle z(\lambda) \rangle = (1/N) \sum_{i=1}^{N} z_i(\lambda)$. The approximation is accurate if $\epsilon_i \tilde{b}(\lambda) \ll 1$, and $z_i(\lambda) - \langle z(\lambda) \rangle$ is relatively small at wavelengths where $\tilde{b}(\lambda)$ is large. In the simulation, we used Eq. (30) directly; the approximation in Eq. (32) was used to obtain the signature vector $\mathbf{b}$ used for the matched filters. In particular, we take $b'(\lambda) = \langle z(\lambda) \rangle \tilde{b}(\lambda)$ for each wavelength, and then normalize to obtain $\mathbf{b} = \mathbf{b}'/|\mathbf{b}'|$. Fig. 2 shows the spectra from a dozen randomly chosen pixels in this image, and for comparison, the spectrum $\varepsilon \mathbf{b}$.
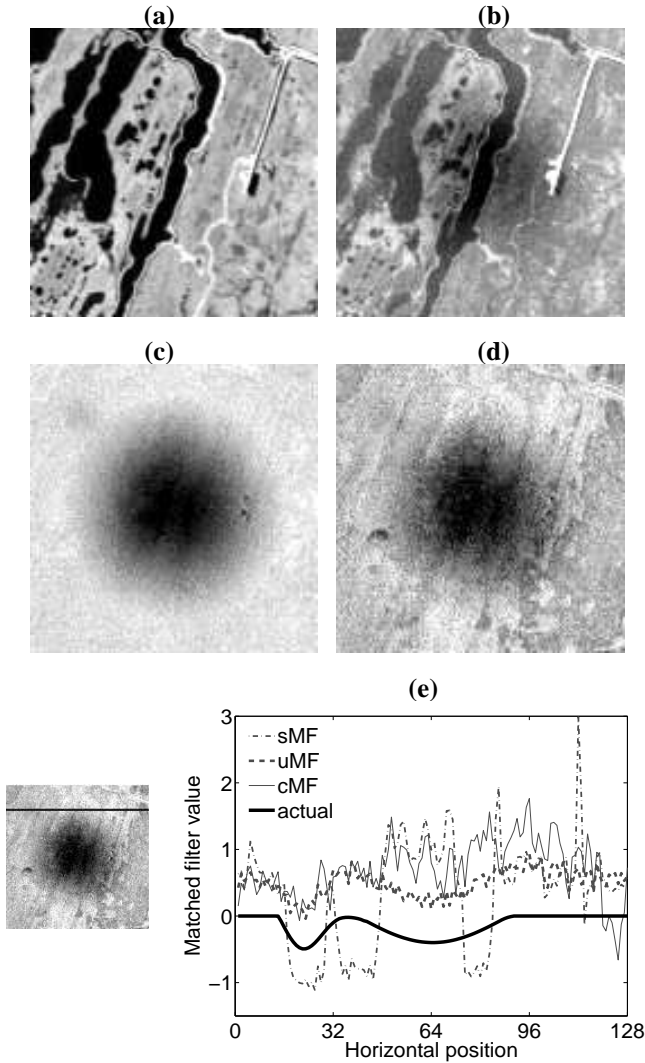
Fig. 3. **(a)** Broadband image of the hyperspectral dataset, including the simulated (and virtually invisible) plume. **(b)** Simple matched filter (sMF) image using the plume signature $\mathbf{q} = \mathbf{b}$. **(c)** Uncontaminated matched filter (uMF), $\mathbf{q} = \mathbf{q_o} = K_o^{-1}\mathbf{b}$, is based on the covariance of the underlying (plume-free) scene. **(d)** Contaminated matched filter (cMF), $\mathbf{q} = K^{-1}\mathbf{b}$, is based on the covariance of the scene with the plume included. **(e)** Profile of the plume strength (along the line indicated by the stripe across the small image to the left of the plot). The horizontal scale corresponds to position, in pixels, across the image; the vertical scale is based on matched filter values which have been normalized across the whole image to have zero mean and variance equal to unity. Shown are profiles for the simple (sMF), the uncontaminated (uMF), and the contaminated (cMF) matched filters. Also shown is the actual plume profile, scaled so that 1 corresponds to $\varepsilon = 12.2$ ppm-m.

Fig. 4. **(a)** Broadband image of the shuffled scene; these are the same pixels as in Fig. 3, but with a randomized spatial layout. **(b)** Simple matched filter (sMF). **(c)** Uncontaminated matched filter (uMF). **(d)** Contaminated matched filter (cMF); unlike Fig. 3(d), the plume has very little effect on the detectability of the plume. **(e)** Profile of the matched filter values; unlike Fig. 3(e), the profiles for the uncontaminated (uMF) and contaminated (cMF) matched filters are nearly the same.

Fig. 3 shows the results of the matched filters applied to the simulated plume imposed on a 128×128 chip from the $d = 224$ channel AVIRIS image. In fact, there are two plumes; a large plume with a gaussian profile in the center of the image, and a smaller and weaker plume in the upper left corner. The correlation between the spatial structure of the image with the spatial structure of the plume concentration leads, in this case, to $|\hat{\zeta}| = 0.43$. Fig. 3(a) shows a broadband image (sum of all channels) of the data with plume, and in this image, the plume is virtually invisible. The simple matched filter $\mathbf{q} = \mathbf{b}$ in Fig. 3(b) enhances the plume signature but does not suppress the background. The optimal matched filter $\mathbf{q} = K_o^{-1}\mathbf{b}$ in
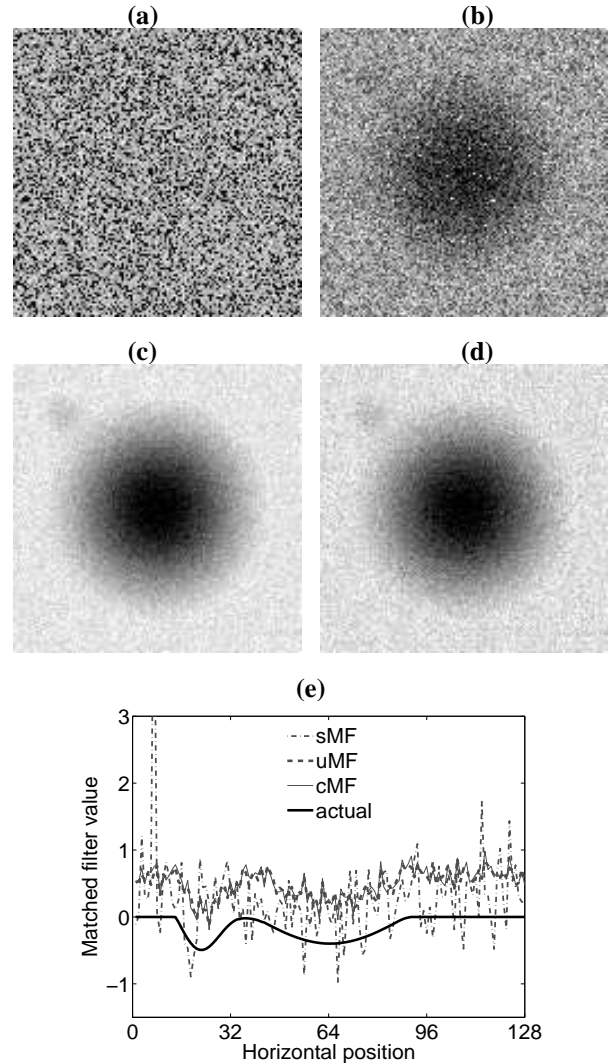
Fig. 3(c) suppresses the background considerably, and in that image even the weaker plume is evident in the image. But when the matched filter is based on a covariance computed from data that is contaminated by the plume, as seen in Fig. 3(d), that matched filter is not as effective at suppressing the background; remnants of that background are evident in the image, and the weaker plume is hidden in that background variation.

Fig. 3(e) shows a profile across the top of the scene that cuts through the center of the weaker plume. This plot recapitulates the results shown in the images in Fig. 3. One can see the dip in the curve (centroid at horizontal position of 21 pixels) due to the weak plume, but only for the uncontaminated matched filter (uMF) is the dip larger than the fluctuations in the matched filter value due to the background clutter. The dip is also seen in the contaminated matched filter (cMF) profile, but its magnitude is on the same order as the fluctuations due

to the background. For the simple matched filter (sMF), the profile is completely dominated by the background clutter.

As a control experiment, we shuffle the pixels in the image. This way, the mean and covariance of the hyperspectral image cube is strictly preserved, but the plume-scene correlation is much smaller. For this image, we have $|\hat{\zeta}| = 0.11$. Fig. 4 shows the shuffled image, and the application of the various matched filters to that data. In Fig. 4(d), in contrast to Fig. 3(d), the effect of including the plume in the covariance estimation is much smaller. The plume strength and the spectral statistics of the background clutter (mean and covariance) are the same in both cases. The difference is that the correlation is much smaller for the shuffled image.
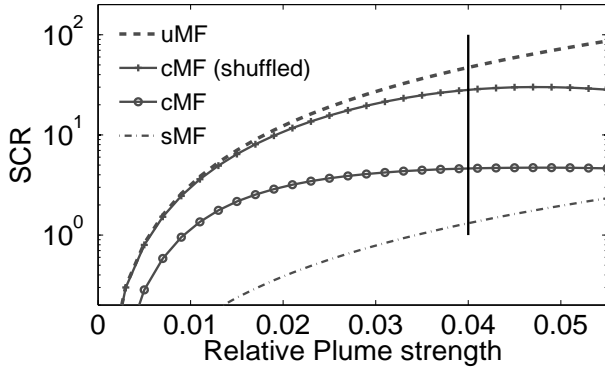

Fig. 5. Average SCR is plotted against relative plume strength for three different matched filters: simple (sMF), uncontaminated (uCF), and contaminated (cMF). The contaminated matched filter result is plotted for both the original (Fig. 3) and the shuffled (Fig. 4) image. The relative plume strength is the signal $|\varepsilon \mathbf{b}| = \varepsilon$ divided by the root mean square of the plume-free image $\langle |\mathbf{z} - \langle \mathbf{z} \rangle|^2 \rangle^{1/2} = \text{trace}(K_o)^{1/2}$. The vertical line corresponds to the plume strength in Figs. 2-4. The SCR is computed from the matched filter images by first distinguishing the on-plume from off-plume pixels. We define S_{on} as the average value of $\mathbf{q}^T \mathbf{r}$ over the on-plume pixels, and S_{off} is the average over the off-plume pixels. With V_{off} is the variance of the matched filter values over the off-plume pixels, then we compute $\text{SCR} = (S_{\text{on}} - S_{\text{off}})^2 / V_{\text{off}}$.

In Fig. 5, we plot the average SCR for matched filter plume detection as a function of relative plume strength. We see that SCR increases with increasing plume strength, but the contaminated matched filters saturate[3] at a maximum value. That maximum value is larger for the shuffled image, where the correlation (in particular, the value of $|\hat{\zeta}|$) is smaller.

VI. CONCLUSION

The band-to-band covariance matrix K is altered when the background is linearly contaminated by signal, but the effect of this contamination on the matched filter $\mathbf{q} = K^{-1}\mathbf{b}$ depends on the correlation between the variation in plume strength over the scene with the variation in the ground spectrum over the scene. The dimensionless vector $\hat{\zeta}$ characterizes this correlation. When $|\hat{\zeta}| = O(1)$, the performance of the matched filter is substantially compromised. We have not investigated typical values of $|\hat{\zeta}|$ in practical plume-detection scenarios, but we emphasize that $|\hat{\zeta}|$ depends on the plume geometry and profile, not on overall plume strength. Along these lines,

[3] In this simulation, as the plume strength is further increased, we begin to see a decrease in the SCR, due to nonlinearity of the Beer's Law absorption.

we also remark that Eq. (29) indicates that $\hat{\zeta}$ will be relatively small as long as the plume is relatively small compared to the size of the scene. In our example, the plume is relatively large, and $|\hat{\zeta}| = 0.43$. In that scene, with a relative plume strength of four percent, the SCR for the contaminated matched filter was reduced by a factor of 12 compared to the uncontaminated matched filter. A second example, based on shuffling the pixels in the first example, had reduced correlation ($|\hat{\zeta}| = 0.11$) but in other respects was identical to the first example, and in that case the SCR loss was only about 30%.

REFERENCES

[1] B. R. Foy, R. R. Petrin, C. R. Quick, T. Shimada, and J. J. Tiee, 'Comparisons between hyperspectral passive and multispectral active sensor measurements." *Proc. SPIE*, vol. 4722, pp. 98–109, 2002.
[2] I. S. Reed, J. D. Mallett, and L. E. Brennan, 'Rapid convergence rate in adaptive arrays," *IEEE Trans. Aerospace and Electronic Systems*, vol. 10, pp. 853–863, 1974.
[3] E. J. Kelly, 'An adaptive detection algorithm," *IEEE Trans. Aerospace and Electronic Systems*, vol. 22, pp. 115–127, 1986.
[4] F. C. Robey, D. R. Fuhrmann, E. J. Kelly, and R. Nitzberg, 'A CFAR adaptive matched filter detector," *IEEE Trans. Aerospace and Electronic Systems*, vol. 28, pp. 208–216, 1992.
[5] L. L. Scharf and B. Friedlander, 'Matched subspace detectors," *IEEE Trans. Signal Processing*, vol. 42, pp. 2146–2156, 1994.
[6] D. Manolakis, D. Marden, J. Kerekes, and G. Shaw, 'On the statistics of hyperspectral imaging data," *Proc. SPIE*, vol. 4381, pp. 308–316, 2001.
[7] C. C. Funk, J. Theiler, D. A. Roberts, and C. C. Borel, 'Clustering to improve matched filter detection of weak gas plumes in hyperspectral imagery," *IEEE Trans. Geoscience and Remote Sensing*, vol. 39, pp. 1410–1420, 2001.
[8] J. Theiler, B. R. Foy, and A. M. Fraser, 'Characterizing non-Gaussian clutter and detecting weak gaseous plumes in hyperspectral imagery," *Proc. SPIE*, vol. 5806, pp. 182–193, 2005.
[9] D. Manolakis and G. Shaw, 'Detection algorithms for hyperspectral imaging applications," *IEEE Signal Processing Magazine*, vol. 19, no. 1, pp. 29–43, 2002.
[10] C.-I. Chang, *Hyperspectral Imaging: Techniques for Spectral Detection and Classification*. New York: Kluwer Academic, 2003.
[11] K. Gerlach, 'The effects of signal contamination on two adaptive detectors," *IEEE Trans. Aerospace and Electronic Systems*, vol. 31, pp. 297–309, 1995.
[12] L. L. Scharf, *Statistical Signal Processing*. Reading, MA: Addison-Wesley, 1990, p. 137.
[13] AVIRIS (Airborne Visible/Infrared Imaging Spectrometer) Free Standard Data Products, Jet Propulsion Laboratory (JPL), National Aeronautics and Space Administration (NASA). http://aviris.jpl.nasa.gov/html/aviris.freedata.html.
[14] A. C. Vandaele, C. Hermans, P. C. Simon, M. Carleer, R. Colin, S. Fally, M. F. Merienne, A. Jenouvrier, and B. Coquart, 'Measurements of the NO_2 absorption cross-section from 42000 cm^{-1} to 10000 cm^{-1} (238-1000 nm) at 220 K and 294 K," *Journal of Quantitative Spectroscopy and Radiative Transfer*, vol. 59, pp. 171–184, 1998.

James Theiler received a Ph.D. in Physics from Caltech in 1987, and subsequently held appointments at UCSD, MIT Lincoln Laboratory, Los Alamos National Laboratory, and the Santa Fe Institute. His interests in statistical data analysis and in having a real job were combined in 1994, when he joined the Space and Remote Sensing Sciences Group at Los Alamos. His professional interests include image processing, remote sensing, and machine learning.

Bernard R. Foy received a Ph.D. in Physical Chemistry from M.I.T. in 1988 and joined the Chemistry Division at Los Alamos in 1990, where he has worked in lidar and passive infrared chemical sensing, and remote sensing data analysis.