PNNL-17873



Prepared for the U.S. Department of Energy under Contract DE-AC05-76RL01830

Development of a Random Field Model for gas Plume Detection in Multiple LWIR Images

Patrick Heasler

September 2008



DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor Battelle Memorial Institute, nor any of their employees, makes **any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or Battelle Memorial Institute. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.**

PACIFIC NORTHWEST NATIONAL LABORATORY operated by BATTELLE for the UNITED STATES DEPARTMENT OF ENERGY under Contract DE-AC05-76RL01830

Printed in the United States of America

Available to DOE and DOE contractors from the Office of Scientific and Technical Information, P.O. Box 62, Oak Ridge, TN 37831-0062; ph: (865) 576-8401 fax: (865) 576-5728 email: reports@adonis.osti.gov

Available to the public from the National Technical Information Service, U.S. Department of Commerce, 5285 Port Royal Rd., Springfield, VA 22161 ph: (800) 553-6847 fax: (703) 605-6900 email: orders@ntis.fedworld.gov online ordering: http://www.ntis.gov/ordering.htm



PNNL-17873

Development of a Random Field Model for gas Plume Detection in Multiple LWIR Images

Patrick Heasler

September 2008

Prepared for the U.S. Department of Energy under Contract DE-AC05-76RL01830

Pacific Northwest National Laboratory Richland, Washington 99352

Executive Summary

This report develops a random field model that describes gas plumes in LWIR remote sensing images. The random field model serves as a prior distribution that can be combined with LWIR data to produce a posterior that determines the probability that a gas plume exists in the scene and also maps the most probable location of any plume. The random field model is intended to work with a single pixel regression estimator–a regression model that estimates gas concentration on an individual pixel basis.

Acknowledgments

Financial support was provided by the U.S. Department of Energy under Contract DE-AC05-76RL01830.

Abbreviations and Acronyms

B: The bayes factors for each pixel, arranged in a vector. B_j represents the Bayes factor calculated for pixel *j*. When there are several images, B_k represents the Bayes factors from image k; $B_k = (B_{1k}, B_{2k}, ..., B_{jk}...)$.

FCP: False Call Probability.

- λ : The apriori rate of occurrence of plumes at the site, measured in plumes/pixel.
- \mathcal{P}_j : An elementary event (bit-map) that represents a plume originating from pixel x_j . \mathcal{P}_0 represents a plume originating from x = 0, so \mathcal{P}_j is a translated version of \mathcal{P}_0 . This is a 2D array and is used to represent an arbitrary plume originating from j.

POD: Probability of Detection.

- site, scene, image: The site is the area on the ground that being imaged. Several images of the site are taken over time (with each image represented as a 2-dimensional matrix. When the images from each time are assembled together into a 3D matrix, this is a scene.
- S_{ji} : A bit-map constructed from a plume shape library; This is the i'th member of the shape library, with the plume origin translated to pixel x_i .
- \vec{X} : The variable X, (a sequence or array), organized as a vector. This notation is sometimes dropped.
- \tilde{X} : The variable X is to be considered a 2D array.
- *n*: Represents the number of pixels in an image or at the site.
- *p*: Number of images (time-slices) recorded.
- *m*: Number of plume shapes in the plume library.

R(event): Odds-ratio associated with the occurrence of the *event*. That is;

$$R(\text{event}) = \frac{Pr(\text{event})}{Pr(Z=\mathbf{0})}$$
(1)

- *I*: Represents an elementary event in the event space of *Z*, a binary sequence of a certain fixed length.
- *Z*: The binary random field that describes plumes at the site.

Contents

Exe	cutiv	e Summ	ary	iii
Ack	nowl	edgmen	ıts	v
Abb	orevia	tions an	nd Acronyms	vii
1.0	Gen	eral De	scription of the Problem	1.1
	1.1	Mathe	matical Assumptions	1.2
	1.2	Mathe	matical Description of Desired Results	1.2
2.0	Ran	dom Fie	eld Model for Plumes	2.1
	2.1	Single	Pixel Regression Model	2.2
	2.2	Posteri	or for Random Field Model	2.3
	2.3	Priors	and Associated Posteriors	2.3
		2.3.1	Single Plume Model	2.4
		2.3.2	Multiple Points of Origin	2.5
		2.3.3	Multiple Shape Model	2.7
		2.3.4	Multiple Image Model	2.7
		2.3.5	General Model	2.8
3.0	Use	fulness	of the Random Field Model: POD Curves for Some Benchmark Scenarios	3.1
	3.1	Compa	arison of Test Statistics	3.2
4.0	App	olication	of the Random Field Model to Data	4.1
	4.1	Descri	ption of the Data	4.2
	4.2	Result	s of Model Fit	4.4
	4.3	Contin	uous Plume Shape Modification	4.6
	4.4	Backg	round Subtraction for Multiple Images	4.9

	4.5 Results for the Improved Model	. 4.9	
5.0	Conclusions	5.1	
6.0	References	6.1	
App	endix A – Calculation of Null and Alternative Distributions for Benchmark Scenarios	A.1	
	A.1 Detection Statistics to be Evaluated	. A.1	
	A.1.1 Null and Alternative Distributions	. A.2	
	A.2 Bayesian Detection Statistic	. A.2	
	A.2.1 Bayes Factor	. A.3	
	A.2.2 Plume Shapes	. A.3	
	A.2.3 Plume Rate of Occurrence	. A.4	
	A.2.4 Strategy for Calculation of Distributions	. A.4	
App	endix B – Plots of Single-Pixel Regression Fits to LWIR Data	B .1	
App	endix C – POP Plots from the Improved Random Field Model	C.1	
Appendix D – Monte Carlo Check of the Random Field Model D.			

Figures

1.1	Plumes estimated from LWIR images taken 15 minutes apart	1.1
3.1	Detection Performance of T_{bayes} , T_{max} , and T_{smooth} for Base Scenario	3.3
3.2	Detection Performance for different SNR	3.4
3.3	Detection Performance for different Plume Lengths	3.5
3.4	Detection Performance for different Numbers of Images (Time-slices)	3.6
4.1	LWIR Spectrum from Site	4.1
4.2	Brightness Map of Site	4.2
4.3	Example of Correlations Caused by Background	4.5
4.4	Gaussian Plumes used in the Random Field model	4.8
A.1	Example of a Plume Library	A.3
D.1	Gaussian Plume used in the Simulation	D.1
D.2	5 Simulated Data Cubes	D.2
D.3	POP's for the 5 Simulated Data Cubes	D.3

1.0 General Description of the Problem

The problem this report considers is plume detection at an industrial site using hyperspectral LWIR images. This report focuses on the development of an algorithm that would improve estimates obtained from existing algorithms by exploiting additional information available concerning the plumes. The specific features of the detection scenario we intend to exploit are;

- The site of interest has been imaged several times with the hyperspectral LWIR system.
- The approximate shape (direction, length) of any prospective plume can be determined from meteorological (MET) data.
- The plumes emanate from point-sources in the site. Consequently, the point of origin of a plume is the same in all registered LWIR images.

The multiple images are assumed to have been taken in the space of minutes to hours; In other words, in a short enough period of time that the background can be assumed to be constant, but any plumes present will show movement from image to image, except at the point of origin. An illustration of plumes conforming to such a scenario is illustrated in Figure (1.1). Three passes over the site were made resulting in a variation of plume direction from pass to pass.

We will not attempt to estimate plumes from raw LWIR hypercubes. We will assume the existence of a good algorithm for detecting a gas (or set of gases) in a single pixel. Examples of such "single-pixel" algorithms include the matched filter algorithms described in Burr and Hengartner (2006), the Bayesian regression model described in Heasler et al. (2007), or the BMA





Plume Origin

model described in Burr et al. (2006).

1.1 Mathematical Assumptions

The model that we develop is formulated on the following assumptions;

- The a-priori probability that a plume originates in the scene at a particular point is constant. That is, the probability of any pixel being the source of a plume is the same as any other.
- The plume shape is described probabilistically by means of a *shape library*. This library will contain bit-maps describing single plumes of varying lengths, directions and widths, each with an associated probability of occurrence. The library will be constructed from MET data with plume direction determined by wind direction, lengths determined by wind speed.
- The plume *point of origin* is assumed to be the same location in each image and represent the extreme up-wind end of the plume.
- The LWIR spectrum associated with an image pixel, *i*, will be summarized by a Bayes Factor, B_i representing the evidence of a plume being present in this particular pixel. This Bayes Factor is computed by application of a single-pixel detection algorithm to the LWIR data.

1.2 Mathematical Description of Desired Results

The problem will be treated in a Bayesian fashion, with Bayesian outputs expected. Consequently the principal results will be posterior probabilities associated with the following events;

Single Plume Originates from pixel *i*: Pr(Plume at i | LWIR Data), this is called *POP* or Probability of a Plume.

Plume Exists at Site: *Pr*(Plume LWIR Data), this is POP at a site.

k **Plumes Exist at Site:** *Pr*(k Plumes LWIR Data)

Pr(Plume at i | LWIR Data) can be formed into a site-map of probabilities, with each pixel representing the probability that a plume might originate from that pixel.

2.0 Random Field Model for Plumes

The mathematical model that we will use to describe plumes is called a *random field model*, see Cressie (1991) for a detailed discussion. Let x represent the location of a pixel at the site and t the time that the site was imaged, then Z(x,t) represents the binary random field that describes plumes at the site. The definition is;

$$Z(x,t) = \begin{bmatrix} 0 & \text{pixel } x \text{ is not in a plume at time } t \\ 1 & \text{pixel } x \text{ is in a plume at time } t \end{bmatrix}$$
(2.1)

Since the pixel locations are discrete, as well as the times, the stochastic process is actually a *lattice* process defined only on the space-time points (x_j, t_k) , for j = 1, 2, 3...n and k = 1, 2, 3...p. We will use the notation \vec{Z} to represent a vector containing all elements of $Z(x_j, t_k)$, in lexicographic order and when one needs to refer to space-time pixel i = (j, k), the short-hand notation $Z_i = Z(x_j, t_k)$ will be used.

In order to make probability statements about \vec{Z} , we must define the probabilities associated with each *elementary event* in this process's event space, u. An elementary event for \vec{Z} can be represented by a binary number, I (i.e. a sequence of zeros and ones) of length np. When the digits in I are properly placed in a 3D array, a bit-map image of plumes at the site will result. An elementary event of particular interest will be represented by $\mathbf{0}$, which represents a sequence of np zeros, with the associated bit-map representing no plumes at the site.

The random field model is defined by specifying the probabilities $Pr(\vec{Z} = I)$ for all possible $I \in \mathcal{U}$. While this is conceptually simple, actual specification is impossible for a general random field because there are so many elementary events in \mathcal{U} . For example, two images of a 100 × 100 site would yield np = 20,000 and $2^{20,000}$ elementary events, and probabilities associated with these events could not be computed on any computer.

To produce a solvable problem, our strategy will be to drastically reduce the number of elementary events that are associated with nonzero probabilities. The only elementary events that will be allowed (i.e. have positive probability) will be those corresponding to bit-maps that have "plume-like" shapes. This will reduce the number of entries in the elementary event space (at least non-zero entries) to the order of 100*n*, which can easily be calculated and stored on a computer.

To simplify calculations concerning this random field (specifically the calculation of it's posterior), we will describe the distribution in terms of an *odds-ratio*, R(I), which is defined as;

$$R(I) = \frac{Pr(\vec{Z} = I)}{Pr(\vec{Z} = \mathbf{0})}$$
(2.2)

If the odds ratio can be defined over the elementary events, then their probabilities are given by;

$$Pr(\vec{Z} = I) = \frac{R(I)}{\sum_{J \in \mathcal{U}} R(J)}$$
(2.3)

So one can easily transform odds-ratios into probabilities. One should note that the transforma-

tion is not computationally feasible if the number of non-zero elementary events is too large.

2.1 Single Pixel Regression Model

We assume that the hypercube data, $L_i = L(x_j, t_k)$ (i.e. L_i is an LWIR spectrum) is employed by a Bayesian regression algorithm to calculate the probability that a plume exists in space-time pixel i = (j,k). In other words, the probability;

$$Pr(Z_i = 1|L_i) \tag{2.4}$$

is calculated. If we have this probability, we can calculate the associated *Bayes factor*, which describes the contribution that the data made to this probability, without the effect of any prior information that might have been incorporated into the model. We want to strip such prior information from the data, because it will be combined with new prior information in the form of a random field model. This probability is related to the Bayes factor,

$$B_{i} = \frac{Pr(L_{i}|Z_{i}=1)}{Pr(L_{i}|Z_{i}=0)}$$
(2.5)

via the formula;

$$\frac{Pr(Z_i = 1|L_i)}{Pr(Z_i = 0|L_i)} = \frac{Pr(L_i|Z_i = 1)Pr(Z_i = 1)}{Pr(L_i|Z_i = 0)Pr(Z_i = 0)} = B_i \frac{Pr(Z_i = 1)}{Pr(Z_i = 0)}$$
(2.6)

In other words, the posterior odds-ratio at pixel *i* is the Bayes factor times the a-priori odds-ratio. The Bayes factor resembles the likelihood ratio statistic of classical statistics, except that it may include model parameters that have been integrated out. For example, if the spectrum L_i is described by the distributional model $f_1(L_i|\theta)$ when a plume exists and $f_0(L_i|\theta)$ when it does not, then;

$$B_{i} = \frac{\int f_{1}(L_{i}|\theta)p_{1}(\theta) d\theta}{\int f_{0}(L_{i}|\theta)p_{0}(\theta) d\theta}$$
(2.7)

In practice, it is difficult to perform the integration over the model parameters (as represented by θ , see Kass and Raftery (1994)), so approximations to the Bayes factor have been developed. For example, if maximum likelihood was used to produce an estimate $\hat{\theta}$, then the standard approximation for B_i is;

$$B_i = \frac{f_1(L_i|\hat{\boldsymbol{\theta}})}{f_0(L_i|\hat{\boldsymbol{\theta}})} \tag{2.8}$$

which is the likelihood ratio statistic to test for the existence of a plume.

In the calculations that follow, we will typically produce the Bayes Factor by dividing the singlepixel posterior by the prior odds-ratio. However, one can also consider using the posterior oddsratio in place of the Bayes Factor. When the odds-ratio in Equation (2.6) is used, the resulting elementary event odds-ratios are simply off by a constant factor. Hence the "true" posterior differs from the calculated posterior by this constant factor.

We will make the assumption that the data, L_{i1} , and L_{i2} originating from two separate pixels are independent of each other, so that the multi-dimensional analog of the Bayes factor can be

computed through multiplication. In other words;

$$\frac{Pr(\vec{L}|\vec{Z}=I)}{Pr(\vec{L}|\vec{Z}=\mathbf{0})} = \prod_{i=1}^{np} \frac{Pr(L_i|Z_i=I_i)}{Pr(L_i|Z_i=0)} = \prod_{i=1}^{np} B_i^{I_i} = \prod(B^I)$$
(2.9)

where the notation $\prod(X)$ represents $\prod_i X_i$ for any vector X. This assumption is probably the most questionable assumption made in construction of the model. Because of background radiation, the spectrum of pixel *i* can be highly correlated to its neighbors.

2.2 Posterior for Random Field Model

This section will calculate the posterior distribution for \vec{Z} , assuming a prior specified by Equation (2.2), and hypercube data, \vec{L} . In other words, we want to calculate the posterior distribution, $Pr(\vec{Z} = I | \vec{L})$, or equivalently the posterior odds ratio;

$$R(I|\vec{L}) = \frac{Pr(\vec{Z} = I|\vec{L})}{Pr(\vec{Z} = 0|\vec{L})}$$
(2.10)

Using a single-pixel model, the hypercube data has been distilled into a *Bayes factor* at each pixel location, *i*, which is denoted by B_i and defined by Equation (2.5). The formula for the posterior odds-ratio, R(I|L) is given by;

$$R(I|\vec{L}) = \frac{Pr(\vec{Z} = I|\vec{L})}{Pr(\vec{Z} = \mathbf{0}|\vec{L})}$$

$$= \frac{Pr(\vec{Z} = I)Pr(L|\vec{Z} = I)}{Pr(\vec{Z} = \mathbf{0})Pr(L|\vec{Z} = \mathbf{0})}$$

$$= R(I)\prod_{i}B_{i}^{I_{i}}$$
or
$$= R(I)\prod(B^{I}) \qquad (2.11)$$

This formula shows that the posterior odds-ratio, $R(I|\vec{L})$ is the prior odds-ratio times a product of Bayes factors, a formula that resembles the single pixel formula presented in Equation (2.6). Equation (2.11) is the basic formula we will use to calculate posterior probabilities for elementary events.

2.3 **Priors and Associated Posteriors**

In this section, we will develop random field models with increasingly complex elementary events (as described by their priors). The reason for developing several models is because the more complex models can be computed in terms of the less complex models, and this allows us to focus on complicating details one at a time instead of all at once.

The models discussed are;

Single Plume Model: We assume a single plume, of specific shape might exist in one image.

Thus there is at most one plume in the image and there is only one image (time-slice) of the site.

- Multiple Points of Origin Model: We assume that more than one plume might exist in the image. Each plume would have a different point of origin. As in the previous case, there is only one image of the site, and the plumes can have only one shape.
- **Multiple Shape Model:** We assume that the plume shape might be any of the shapes present in the plume library in a specific image. This is the single plume model with an allowance for different plume shapes.
- **Multiple Image Model:** Multiple images (time=slices) have been made of the site, and at a possible point of plume origin the plume must exist in all of the the images. However, it might not have the same shape.
- **General Plume Model:** This assumes multiple shapes and points of origin, and multiple images (time-slices).

2.3.1 Single Plume Model

To obtain a tractable prior, the idea is to restrict the elementary events $I \in \mathcal{U}$ to a set that can be produced on a computer. Except for the elementary event **0**, each elementary event is supposed to describe plumes in the images, which can be indexed by their point of origin.

We will begin by developing a "one-plume" model, show how the posterior can be calculated for this restrictive model, and then extend the model by including multiple plumes in a site, multiple plume shapes, and finally multiple time-slices. Initially, we will assume;

- There is only one LWIR image of the site (i.e. one time-slice).
- The a-priori probability of a plume originating from any pixel j is λ , (with $1 n\lambda$ representing the a-priori probability of no plume being at the site).
- At most, only one plume can exist at the site. Hence, there are either no plumes at the site (with prior probability $1 n\lambda$) or exactly one (with probability $n\lambda$).
- The plume originating at location x_j has the same shape as the plume at location x = 0, consequently the bit-map describing a plume at location j, \mathcal{P}_j , is given by a translation;

$$\mathcal{P}_j(x) = \mathcal{P}_0(x - x_j) \tag{2.12}$$

• This plume can have only one shape. The plume shape is described by a single bit-map, $\mathcal{P}_0(x)$, representing a plume that originates from the origin x = 0. $\mathcal{P}_0(x)$, will always include pixel x = 0, so $\mathcal{P}_0(0) = 1$.

With these assumptions, the only non-zero elementary events for $\vec{Z} = I$ correspond to the bitmaps, I = 0, and $I = \mathcal{P}_j$ for j = 1, 2, 3, ...n. These assumptions concerning plume shape have reduced the elementary event space to a tractable number (n + 1). Furthermore, since \mathcal{P}_j is simply a translation of \mathcal{P}_0 , the posterior odds-ratio can be written as a convolution of the Bayes Factors with the \mathcal{P}_0 bit-map and computed quite quickly using a fast Fourier transform. The convolution analog to Equation (2.11) is;

$$log(R(\mathcal{P}_j|\vec{L})) = log(R(\mathcal{P}_j)) + \sum_k \mathcal{P}_0(x_k - x_j) log(B_k)$$
(2.13)

Or, if we let \tilde{R} , \tilde{P}_0 and \tilde{B} represent 2D arrays, then the above formula can be written as the convolution;

$$\log(\tilde{R}|L) = \log(\tilde{R}) + \tilde{\mathcal{P}}_0 \otimes \log(\tilde{B})$$

=
$$\log(\frac{\lambda}{1 - n\lambda}) + \tilde{\mathcal{P}}_0 \otimes \log(\tilde{B})$$
 (2.14)

The convolution formula Equation (2.14) ignores the difficulties that occur on the image edges. If a plume originates at a pixel near the image edge, it will extend beyond the image and must be truncated to the image pixels. To make the convolutions produce correct results, on needs to pad the Bayes factor array, \tilde{B} , with 1's on the edges; A Bayes factor of 1 indicates no evidence for the existence of a plume or its absence.

The log of an odds-ratio log(R), is also called the *negepotential function* (see Cressie (1991), pg414), so the formula for the posterior could have been developed in terms of a negepotential function instead of the odds-ratio, R. It should also be noted that the odds-ratio for the no-plume event, **0** is also easy to calculate. By definition, $R(\mathbf{0}|L) = 1$.

Equation (2.14) allows us to calculate the posterior probabilities associated with *a plume* at the site, (site POP) and *a plume originating from pixel j*, (POP at j) which is the desired objective.

2.3.2 Multiple Points of Origin

In the previous section, we were able to calculate the posterior probabilities using a prior that assumed *at most one* plume existed at the site and therefore in the image. This assumption will not be reasonable for many sites. In fact, one could argue that it is more likely to have either no plumes at the site (plant shut down), or several (plant operating).

In this section, we extend the "at most one plume" prior into a multiple plume prior and find that the multiple plume model is really based on the at-most-one-plume model. In other words, fitting the one-plume model and then combining the probabilities in appropriate fashion will yield results about multiple plumes.

The most natural assumption for creating a multiple plume random field is to assume that;

• the number of plumes at the site is Poisson, with the probability of *k* plumes at the site given by;

$$Pr(k \text{ Plumes}) = exp(-n\lambda)\frac{(n\lambda)^k}{k!}$$
 (2.15)

• and the points of origin for the k plumes are independent of each other and randomly chosen from the site-map. Hence, the probability that the k plumes originate from pixels $j_1, j_2...j_k$ is $1/n^k$.

The elementary event corresponding to two plumes originating from locations j1 and j2 is represented by the bit-map $\mathcal{P}_{j1} + \mathcal{P}_{j2} - \mathcal{P}_{j1}\mathcal{P}_{j2}$. If the plumes are far enough away from each other so they don't intersect, then the two-plume bit map reduces to $\mathcal{P}_{j1} + \mathcal{P}_{j2}$, and we can consider this to be a reasonable approximation because the probability of two plumes being close to each other is small. In like manner, the bit-map containing k plumes at j1, j2...jk can be approximated by

$$\mathcal{P}_{j1} + \mathcal{P}_{j2} + \ldots + \mathcal{P}_{jk} \tag{2.16}$$

We now explicitly derive the posterior odds-ratio associated with k = 1, 2 and then present the formula for arbitrary k. For one plume, occurring at location j, the prior is;

$$R(\mathcal{P}_{j}) = \frac{Pr(\text{Plume originating at } j)}{Pr(\text{No plumes in site})}$$
$$= \frac{exp(-n\lambda)(n\lambda)\frac{1}{n}}{exp(-n\lambda)}$$
$$= \lambda$$
(2.17)

and the posterior is;

$$R(\mathcal{P}_j|L) = R(\mathcal{P}_j) \prod (B^{\mathcal{P}_j})$$

= $\lambda \prod (B^{\mathcal{P}_j})$ (2.18)

For 2 plumes at locations j1, j2, the prior is;

$$R(\mathcal{P}_{j1} + \mathcal{P}_{j2}) = \frac{Pr(\text{Two plumes at } j1 \text{ and } j2)}{Pr(\text{No plumes in site})}$$
$$= \frac{exp(-n\lambda)\frac{(n\lambda)^2}{2!}\frac{1}{n^2}}{exp(-n\lambda)}$$
$$= \frac{\lambda^2}{2!}$$
$$= \frac{1}{2!}R(\mathcal{P}_{j1})R(\mathcal{P}_{j2})$$
(2.19)

Substituting this expression into Equation (2.11), we find that the posterior odds-ratio is;

$$R(\mathcal{P}_{j1} + \mathcal{P}_{j2}|L) = R(\mathcal{P}_{j1} + \mathcal{P}_{j2}) \prod (B^{\mathcal{P}_{j1} + \mathcal{P}_{j2}})$$

$$= \frac{1}{2!} R(\mathcal{P}_{j1}) R(\mathcal{P}_{j2}) \prod (B^{\mathcal{P}_{j1}}) \prod (B^{\mathcal{P}_{j2}})$$

$$= \frac{1}{2!} R(\mathcal{P}_{j1}|L) R(\mathcal{P}_{j2}|L)$$
(2.20)

Using similar algebra, one can show that for *k* events the formula is;

$$R(\sum_{i=1}^{k} \mathcal{P}_{j_i}|L) = \frac{1}{k!} \prod_{i=1}^{k} R(\mathcal{P}_{j_i}|L)$$
(2.21)

In other words, we can calculate multiple-plume odds-ratios from the single-plume odds-ratios, $R(\mathcal{P}_j|L)$. Because of the nature of equation (2.21), we need not consider all possible multiple plume combinations; We know the only plume combinations with significant probability originate from the locations where $R(\mathcal{P}_j|L) >> 0$. This set of locations will most likely be small, and thus we need only compute all posterior odds-ratios on the subsets of these locations.

2.3.3 Multiple Shape Model

Assume that we have compiled a library of plume shapes S_{0i} , i = 1, 2, 3...m, with the probability that plume shape *i* might occur being h_i . Each shape S_{0i} , represents a bit-map with the plume origin centered on location x = 0, while S_{ji} represents S_{0i} translated to pixel *j*. The formula for calculating the posterior odds-ratio for one of these shapes is really the same as that given in Equation (2.18), except that the prior probability is modified by h_i . Let us define;

$$R(S_{ji}) = \frac{Pr(\text{Plume originating at } j \text{ with shape } i)}{Pr(\text{No plume})}$$
(2.22)

so that,

$$R(S_{ji}|L) = R(S_{ji}) \prod (B^{S_{ji}})$$

= $h_i \lambda \prod (B^{S_{ji}})$ (2.23)

The odds-ratio for a plume originating at pixel *j* is therefore given by;

$$R(j|L) = R(\text{Plume at } j|L)$$

= $\sum_{i} R(S_{ji}|L)$
= $\lambda \sum_{i} h_{i} \prod (B^{S_{ji}})$ (2.24)

2.3.4 Multiple Image Model

We assume that p images (time-slices) have been recorded, and that any plume in the image must have a common point of origin, but otherwise the shape can vary from image to image. The Bayes-factor data from image k is identified by B_k (Realize that B_k is not a single value as in the previous sections, but a 2D array). Posterior odds-ratios associated with image k will be identified by the notation $R(\cdot|L_k)$. Also let $I = i_1, i_2, ..., i_p$ represent a set of p plume shapes chosen from the shape library, with the elementary event corresponding to S_{jI} defined by;

$$S_{j,i_k}$$
 present in time-slice $k = 1, 2...p$ (2.25)

Then the prior probability of occurrence is given by the odds-ratio;

$$R(\mathcal{S}_{jI}) = \lambda \prod_{k} h_{i_k} \tag{2.26}$$

In other words, the shape in time-slice k is independent of the other time-slices.

With this prior placed upon the elementary events, the posterior becomes;

$$R(S_{jI}|L) = R(S_{jI}) \prod_{k} \prod (B_{k}^{S_{j,i_{k}}})$$

$$= \lambda \prod_{k} h_{i_{k}} \prod (B_{k}^{S_{j,i_{k}}})$$
or
$$= \frac{1}{\lambda^{p-1}} \prod_{k} R(S_{j,i_{k}}|L_{k})$$
(2.27)

The last step follows by substitution of Equation (2.23). Hence the posterior for a multiple image data-cube can be computed by first computing the posterior for each time-slice and then multiplying them together. Furthermore, one need not calculate the odds-ratios for all elementary events but only for those in which the single time-slice odds-ratios, $R(S_{j,ik}|L_k)$ are large.

If one attempted to calculate the odds-ratio for a plume originating at pixel *j* using Equation (2.27), one would have to sum over *I*, resulting in a sum containing m^p terms. It is possible to reduce arithmetic operations by applying the distributive law;

$$R(j|L) = R(\text{Plume Originates at } j|L)$$

$$= \frac{1}{\lambda^{p-1}} \sum_{I} \prod_{k} R(S_{j,i_{k}}|L_{k})$$

$$= \frac{1}{\lambda^{p-1}} \prod_{k} [\sum_{i} R(S_{ji}|L_{k})]$$

$$= \frac{1}{\lambda^{p-1}} \prod_{k} R(j|L_{k}) \qquad (2.28)$$

with $R(j|L_k)$ representing the odds-ratio of a plume being at pixel *j* in time-slice *k*, and computed with Equation (2.24).

2.3.5 General Model

One can combine the modifications discussed in the previous sections to produce a posterior calculation that includes (1) multiple plume shapes, (2) multiple plumes at a site, and (3) multiple time-slices. First note that the odds-ratios of all elementary events associated with such a model cannot be calculated. For example, if the site contains $n = 10^4$ pixels, p = 4 time-slices, and m = 100 plume shapes, then the number of elementary events would be;

$$nm^p = 10^4 \times 100^4 = 10^{12} \tag{2.29}$$

a trillion events, too large to calculate on present computers.

To produce a successful calculation, we exploit Equation (2.28), which can be calculated, and also use the fact that most of the elementary events should be nearly zero. Here is an outline for the calculation;

- 1. Calculate all posterior odds-ratios $R(S_{ji}|L_k)$ j = 1, 2, 3...n, i = 1, 2, 3...m, and k = 1, 2, 3...p. There are *mnp* such odds-ratios, a tractable number.
- 2. Calculate the odds-ratios, R(j|L), associated with a plume originating from pixel *j* using Equation (2.28).
- 3. Identify locations, *j*, where R(j|L) >> 0 and form these locations into the set A_1 . This is the set of possible plume origination sites.
- 4. To calculate probabilities of plume pairs, form the set \mathcal{A}_2 , the set of pairs $i1, i2 \in \mathcal{A}_1$ where (1/2)R(i1|L)R(i2|L) >> 0, and use Equation (2.19) to calculation the pair-wise odds-ratio. Calculate similar odds-ratios for triples, 4-tuples, etc. and form \mathcal{A}_3 , \mathcal{A}_4 , etc. It is assumed that these sets will be small.
- 5. Convert posterior odds-ratios into posterior probabilities using Equation (2.3). These probabilities are the desired output of the calculation.

3.0 Usefulness of the Random Field Model: POD Curves for Some Benchmark Scenarios

In this section, we calculate the false call probability (FCP) and probability of detection (POD) performance for a plume detection statistic derived from the Bayesian Plume Model developed in Section (2). This Bayesian detection statistic is compared to two alternative detection statistics to gauge its performance. The alternative detection statistics represent two of the most commonly used plume detection techniques.

Our objective for this calculation is to choose a simple scenario that can be easily calculated but will show how important plume information might be. If the random field model does not offer significantly improved detection under ideal conditions, then it is not worth pursuing.

The scenarios we base these calculations upon assume that the random field model is correct; That is, it correctly describes the actual shape of the plume present in the data and the distributional properties of the data. To calculate FCP and POD, one must describe a null and alternative distribution for the data. The null distribution describes the resulting image data when no plume exists, while the alternative distribution describes the image date when a plume (or many plumes) exist. With a little reflection, the reader can see that for a problem as complicated as this, there is no single null or alternative distribution. For example, the null distribution is influenced by the site background; Each different site background produces a different null distribution. In a similar manner, the alternative distribution is influenced by temperature contrast, dimensions of the plume, gas concentration, etc.; Different assumptions concerning these variables produce different alternative distributions.

A complete description of the scenario assumptions is given in Appendix A, along with a derivation of the null and alternative distributions for three detection statistics. The three detection statistics evaluated consist of;

Site POP, T_{bayes} : This is the detection statistic produced by the random field model.

- **Maximum Statistic,** T_{max} : The maximum of all the single pixel regressions results. This is equivalent to applying a detection threshold on the single-pixel regression results and is the most direct way to use the single-pixel results for detection.
- **Smoothed Statistic**, T_{smooth} : This smooths the single-pixel results before thresholding. This detection methodology would be the most obvious way to incorporate the plume shape into a detection algorithm.

We calculate the alternate distributions for the following scenarios;

Base Case: The SNR for the plume is 2, the plume length is 5 pixels, and two images (timeslices) have been taken of the site.

SNR 1: Same as base case, but SNR has been lowered to 1.

SNR 5: Same as base case, but SNR has been raised to 5.

Plume Length 2: Same as base case, but the plume length has been lowered to 2.

Plume Length 10: Same as base case, but the plume length has been raised to 10.

Number Images 1: Same as base case, but the number of time-slices has been lowered to 1.

Number Images 3: Same as base case, but the number of time-slices has been raised to 3.

3.1 Comparison of Test Statistics

Before presenting the simulation results in detail, it might be informative to mention some theoretical results. T_{bayes} represents the standard Bayesian statistic used in decision theory. A well-known result (Ferguson (1967), pg59) from decision theory shows that Bayes decision procedure is *admissible*; In other words, the ROC curve associated with T_{bayes} should always "be equal to or above" the ROC curve associated with any other test statistic for plume detection (at least image-wide plume detection). So, theory would lead us to expect the ROC for T_{bayes} to be "better than" that of T_{max} or T_{smooth} , the real question is by how much.

Figure (3.1) displays the ROC curves^(a) for the three test statistics using the scenario outlined in the previous sections. Note that *FCP* has been logged, so that we can examine performance in FCP regions of interest; In a typical problem, one would not be interested in running the detection algorithm with a *FCP* much greater than 1%. The defining scenario variables (i.e. SNR, plume length, # of time slices) are listed in the figure headings and represent those of the *Base Case* scenario described previously.

The most obvious conclusion one can draw from Figure (3.1) is that the plume information utilized by both T_{bayes} and T_{smooth} dramatically improves detection performance, as compared to T_{max} . Also, T_{bayes} is substantially better than T_{smooth} , with its POD approximately double that of T_{smooth} (when FCP < 1%). From these results, we can conclude that incorporation of spatial/temporal plume information into the detection algorithm can make a huge difference in POD, if correctly done.

In the following series of ROC plots, we have varied the scenario variables SNR, plume length, and number of time slices, with the objective of obtaining a better understanding of when plume information is really important in a detection scenario. For example, if *SNR* is large, then even T_{max} will work, but how large does *SNR* have to be?

In Figure (3.2), SNR is varied from 1 to 5. As one can see from the ROC plots, this range does seem to bound acceptable performance for all three statistics; With SNR = 1, none of the three statistics can detect plumes acceptably, while with SNR = 5, all three do an almost perfect job. So plume information is important for "weak plumes", with a weak-plume defined by $SNR \approx 2$.

Figure (3.3) shows the effect of plume length on POD. With a plume length of 2 pixels, T_{max} has

⁽a) A ROC curve plots POD versus FCP

Figure 3.1. Detection Performance of T_{bayes} , T_{max} , and T_{smooth} for Base Scenario



snr= 2 plume.len= 5 No.images= 2



Figure 3.2. Detection Performance for different SNR

low POD (but T_{bayes} is much better than the others, producing a POD = 20% at FCP = 1%). On the other hand, a plume length of 10 pixels, makes detection a virtual certainty for T_{bayes} and T_{smooth} but not for T_{max} . It is surprising that T_{max} does so badly with such a large plume. This supplies further evidence of the value of plume information in detection.

Figure (3.4) shows the effect of different numbers of time-slices. In the first case, the number of time slices has been restricted to 1, and in this case we see that T_{bayes} and T_{smooth} deliver almost identical performance. The formulas for T_{bayes} and T_{smooth} are algebraically distinct, so it is a surprise that they deliver similar numerical results. In this case, for example, T_{smooth} has the form;

$$T_{smooth} = \max_{S, pixels} S \otimes Y \tag{3.1}$$

with the maximum being taken over all pixels in the smoothed image, and over all shapes in the library. In contrast, T_{bayes} involves the transformed sum;

$$T_{bayes} = \sum_{\mathcal{S}, pixels} \exp(\gamma(\mathcal{S} \otimes Y - \frac{n_{\ell}\gamma}{2}))$$
(3.2)



Figure 3.3. Detection Performance for different Plume Lengths

If the number of shapes in the shape library were increased (these scenarios use a shape library with only three members), one might expect the performance T_{smooth} and T_{bayes} to diverge, with T_{bayes} being better. On the other hand, with real data, which does not have the exact distributional properties described by the model, one might expect T_{smooth} to behave better than T_{bayes} . In other words, T_{smooth} might be more robust to distributional deviations than T_{bayes} .

When we increase the number of time-slices, the performance of T_{bayes} becomes dramatically better than T_{smooth} , indicating that the temporal information T_{bayes} incorporated into the random field model is very important. This temporal information is that *the plume origin pixel is the same in all time-slices*. With three time slices, T_{bayes} can detect a weak plume with near certainty (i.e. $POD \approx 1$), even at the lowest FCP rates.

These ROC curves show that the Bayesian plume model can deliver dramatically improved detection performance when the plume is small and weak. Such a model is worth further consideration.



Figure 3.4. Detection Performance for different Numbers of Images (Time-slices)

4.0 Application of the Random Field Model to Data

The results in the last section show that the random field model would produce dramatically better performance than single-pixel regression if the data and detection scenario exactly conformed to model assumptions. Of course, this is never the case so an important question to be answered is; How well does the model perform when the assumptions are only approximately satisfied? We would like the detection procedure to work well on real data. Specifically, we would like the random field model to produce realistic Probability of Plume (POP) estimates, and correctly locate any existing plume.

In this section, these questions are examined by applying the model to a real LWIR data set. In this exercise, certain problems with the model became apparent, and were rectified. (see Section 4.3)





Pixel 170 25 Spectrum and Plume Radiance

Black: Background, BG: Observed, Red: Plume Spectrum

4.1 Description of the Data

The data set employed for the test contained 10 images of an industrial site, with the time between images being approximately 10-20 minutes. The site was imaged with an LWIR instrument that expressed the spectrum in 128 bands. An example of a typical LWIR spectrum from this site is presented in Figure (4.1). In this Figure, the blue-green line represents the actual spectrum, while the red and black lines represent results from the regression fits. The black spectral line (which is almost identical to the blue-green spectra) represents the background spectra (i.e. spectra without any plume gases), while the red spectra represents the spectra we would see from a plume that is optically opaque. The space between the black and red spectra is caused by the "emissivity contrast" between the atmosphere and background. If the emissivity contrast was zero, it would not be possible to see any plume over the pixel. From a regression perspective, a small emissivity contrast would result in a large standard deviation for the gas concentration estimates.



Site, Expressed in Brightness (K)



actually originates from a pixel containing a strong gas plume, so one can see that we are trying to find effects that are not obvious in the spectrum; The plume gases have created a small depression in the blue-green spectra at wavenumber 950.

Figure (4.2) presents a brightness map of the site. As one can see from the map, the average temperature at the site is about 300K. One can also see the road grid in the map, building outlines, and plant piping. These features all have different emissivities, and different temperatures, which cause a different capability to estimate gas concentrations over the different features.

The true-state of the site is not known, so we cannot be certain of the plumes that actually exist in the images. We have used the LWIR data to produce a best guess of the true state for the site. In Appendix B, the concentration SNR's (i.e. gas concentration estimate over its standard deviation) are plotted from the regression fits. From the plots, two gas plumes are obvious, but we are really interested in how well the random field model can find weak plumes (with a low false call rate). We have attempted to identify weak plumes by visually examining a movie of each gas SNR and identifying plume-like features. A plume-like feature is a shape that has a definite point of origin but otherwise moves as time progresses.

Table (4.1) presents our best guess as to the plumes present (i.e. under the column labeled "True State". In other words, we will be comparing the random field results to that produced by human evaluation.

	No. Pi	xels With	Random Field	True
Gas	SNR > 2	SNR > 4.5	Model POP	State
1	701	50	1	no plume
2	3089	126	1	no plume
3	654	152	1	strong plume
4	579	183	1	strong plume
5	10	1	0	no plume
6	214	45	1	weak plume
7	64	14	1	weak plume
8	45	11	1	no plume
9	103	27	1	no plume
10	27	11	1	no plume
11	6293	87	1	no plume
12	2221	66	1	no plume
13	20	5	1	no plume

Table 4.1.	Best Estimate of Plumes in Image
------------	----------------------------------

True State determined by visual examination of Data.

4.2 Results of Model Fit

The random field model was fit to the data and these results are presented in Table (4.1), (in column 4) for each gas. To allow comparison to alternate test statistics, we also compiled the number of pixels whose SNR was greater than 2 and 4.5. If there were no plumes and the SNR statistic was behaving correctly, we would expect 2.5% of the pixels (i.e. approximately 450) to be above 2, and there would be a 5% chance that 1 or more pixels would be above 4.5. In other words, SNR > 4.5 produces a plume detection test for a 5% level of significance. From the results in this table, one can see that use of the single pixel SNR produces a test with a high false call rate; All of the 13 gases in the table are judged to be present when using a detection threshold of 4.5.

We had hoped that the random field model would do better than simple thresholding, but it also suffers from a high false call rate. The probability of plume (POP) results for the random field model are presented in column 4 of Table (4.1). Only gas 5 is ruled to be absent; all other gases have a POP of 1, indicating the definite presence of a plume. The random field model also suffers from a high false call rate and does not produce correct POPs. It is obvious that the distributional assumptions that the random field model rests upon are not correct for this data.

A careful examination of the data revealed several problems with the model assumptions. The major problem had to do with the use of SNR to define the Bayes factor. Our model assumes that the SNR are independently distributed from pixel to pixel. However, the SNR plots in Appendix B show that this is not the case; These plots display background shapes which cause correlations between adjacent pixels. One might ascribe this problem to a malfunctioning regression fit that is not capable of removing all the background effects, but the problem really is caused by use of SNR to define the Bayes factor used in the random field model. Since the signal noise ratio is the gas concentration estimate divided by its standard deviation (i.e. $SNR = \hat{C}/sd(\hat{C})$, background features will appear in the SNR if they are present in the denominator, $sd(\hat{C})$, but not the numerator \hat{C} . So even if the pixel regression eliminates the background effects from \hat{C} , they may still be present in $sd(\hat{C})$ and cause pixel-to-pixel correlations in the SNR.

A little reflection concerning $sd(\hat{C})$ leads us to conclude that this standard deviation will always be strongly correlated to background because it is inversely related to temperature contrast; A plot of $sd(\hat{C})$ roughly displays the temperature difference between ground and air. Figure (4.3) illustrates this effect, by plotting \hat{C} , $sd(\hat{C})$, and SNR = C/sd(C) for gas 4. One can see that sd(C) basically provides an image of the background, with the road network displaying the smallest temperature contrast, and the buildings the greatest.

Another problem is related to the use of binary plume shape in the random field model. (i.e. the plume is either present or absent in a pixel). In reality the plume should have a shape described by gas concentration. Use of a binary shape forces one to choose an arbitrary threshold that defines when a plume ceases to exist. For example, the binary plume model we have adopted defined a plume to be present when the SNR > 2 and absent when SNR < 2. Of course, the use of 2 to define the presence of a plume is an arbitrary selection; One could have used 1 or 0.5 as the threshold. One would like to employ a definition of a plume that did not rely on an arbitrary threshold.


Finally, an examination of the diagnostics produced by the pixel regression fits indicated that there were problems with about 0.5% of the fits. In these fits, the RMSE was about twice as large as it should be, indicating that either the instrument was malfunctioning, or the background model was inadequate for the data. These fits frequently produced unrealistic temperature contrasts (of up to 100K). Visual indications of such data problems are present in the SNR plots of Appendix B. Some of the plots display horizontal stripes, which are probably caused by a malfunctioning or mis-calibrated detector in the instrument.

To rectify the problems discussed above the following modifications were made to the random field model;

- 1. Pixels with RMSE > 1.5 error were eliminated.
- 2. Binary plume shapes were replaced by shapes defined in terms of gas-concentration. The Bayes Factor was also defined in terms of concentration, C, and its standard deviation, sd(C), so that the correlations in sd(C) could be accounted for.
- 3. Static Background shapes are removed from multiple images using ANOVA.

These modifications are presented in detail in the following sections.

4.3 Continuous Plume Shape Modification

This development assumes that the plume is expressed in terms of concentration, such as a Gaussian plume. A plume originating at pixel *i*, denoted by $\mathcal{P}_i = (\mathcal{P}_{i1} \dots \mathcal{P}_{ik} \dots)$, is a 2-dimensional array with \mathcal{P}_{ik} representing the relative concentration of the gas in pixel *k*. When $\mathcal{P}_{ik} = 0$, the plume is not present in that pixel. We assume that \mathcal{P}_i describes the shape of the actual plume well, but not its magnitude. The actual plume is given by the term $\alpha_i \mathcal{P}_i$, where α_i represents the unknown source emission from pixel *i*, in units of ppmM.

The single-pixel regression model produces a concentration estimate for pixel k of \hat{C}_k which should be estimating the value $\alpha_i \mathcal{P}_{ik}$ with uncertainty (SD) of σ_k (or a precision of $W_k = \sigma_k^{-2}$). In other words, the single-pixel regression supplies us with values of \hat{C}_k and uncertainties σ_k . Given this framework, one can use regression to estimate the prospective source emission, α_i and construct a Bayes Factor from this estimate.

The regression estimate for α_i is;

$$\hat{\alpha} = \frac{\sum_{k} \mathcal{P}_{k} W_{k} \hat{C}_{k}}{\sum_{k} \mathcal{P}_{k}^{2} W_{k}}$$
(4.1)

and with variance;

$$Var(\hat{\alpha}) = \left(\sum_{k} \mathcal{P}_{k}^{2} W_{k}\right)^{-1}$$
(4.2)

Note that $\hat{\alpha}_i$ and $var(\hat{\alpha}_i)$ can be computed at each pixel by convolutions. The convolution formula for these two statistics are;

$$Var(\hat{\alpha}) = (\mathcal{P}^2 \otimes W)^{-1} \tag{4.3}$$

and

$$\hat{\alpha} = (\mathscr{P} \otimes (W\hat{C})) Var(\hat{\alpha}) \tag{4.4}$$

In other words, these convolution formulas provide us with a computationally efficient way to calculate $\hat{\alpha}$ for an entire image.

 $\hat{\alpha}_i$ is approximately normally distributed, so the Bayes Factor for the *origination* of a plume at pixel *i* is;

$$B_{plume,i} = \frac{Pr(\alpha_i > T | \hat{\alpha}_i, Var(\hat{\alpha}_i))}{Pr(\alpha_i \le 0 | \hat{\alpha}_i, Var(\hat{\alpha}_i))}$$

=
$$\frac{\Phi((\hat{\alpha}_i - T) / sd(\hat{\alpha}_i))}{\Phi(-\hat{\alpha}_i / sd(\hat{\alpha}_i))}$$
(4.5)

We have included a threshold of T on the plume source for generality. If nothing is known about the source, one can set T = 0, and this is the threshold we will use.

This Bayes factor, $B_{plume,i}$, replaces the expression B^{I} in Equations (2.11) and (2.24) so that the posterior odds ratio now becomes;

$$R(\mathcal{P}_i|\hat{C}) = R(\mathcal{P}_i)B_{plume,i} \tag{4.6}$$

and

$$R(\text{Plume at } i|\hat{C}) = \lambda \sum_{k} h_k B_{\mathcal{S}_{ik}}$$
(4.7)

with shapes in the plume library indexed by k.

The plume shapes we will employ are derived from a Gaussian plume model. A Gaussian plume models is defined in terms of the following parameters;

x: Distance (i.e. pixel coordinates) are measured in terms of Meters.

Plume width at source: Is given by $\sigma_0 M$.

Plume Diffusion Constant: The formula for plume width x_2 pixels downwind is given by

$$\sqrt{\sigma_0^2 + \sigma_1^2 \frac{x_1}{\nu}} \tag{4.8}$$

Hence σ_1^2 is given in units of M^2/sec There are several formulas for plume "diffusion," depending on the nature of the flow such as laminar or turbulent flow.

Wind Speed: The wind speed is *v* M/sec.

Emission: Emission at source is $E_0 M^3 / sec$ of gas.

Concentration: Concentration of target gas at point is measured as M^3 of gas per M^2 of area. This is equivalent to the instrument measurement of *ppmM*.

So the plume concentration at point (x_1, x_2) , where x_1 is the downwind direction and x_2 is the cross-wind direction, and the plume origin is (0,0), is;

$$\frac{E_0}{v}\phi(x_2;0,\sigma(x_1)) \tag{4.9}$$

with

$$\sigma(x_1)^2 = \sigma_0^2 + \sigma_1^2 \frac{x_1}{v}$$
(4.10)

The regression parameter, α defined above is therefore equal to;

$$\alpha = \frac{E_0}{\nu} \tag{4.11}$$

In other words, the plume regression fits can also supply an estimate of gas emission.



Figure (4.4) presents a library of 9 Gaussian plumes used in the improved random field model. These plumes are produced by using 3 values for the diffusion constant σ_2 to produce wide, intermediate and thin plumes. Three plume directions are also represented in the plume library

Plume direction: 25,40,55°, centered on the Gas 4 plume in Appendix 1.

 σ_0 : 0.75 pixels (A pixel has a dimension of approximately 2M).

 $\sigma_1/\sqrt{(v)}$: 0.58, 0.33, and 0.28 pixels.

The plumes represented in the library have also been normalized so that their maximum value is 1, and so that distance is measured in pixels. If L_p represents the dimension of a pixel, then the formula for converting an α estimate from the regression to an emission is;

$$E_0 = \alpha v \sqrt{2\pi} \sigma_0 \tag{4.12}$$

4.4 Background Subtraction for Multiple Images

The concentration data at pixel location *i* and time *t* is represented by \hat{C}_{it} and an associated standard deviation of σ_{it} . Then the estimate for the background is;

$$\mu_i = \sum_t W_{it} \hat{C}_{it} \tag{4.13}$$

$$W_{it} = \frac{1/\sigma_{it}}{\sum_{i} 1/\sigma_{it}}$$
(4.14)

And the background-subtracted data (BSC) is;

$$BSC_{it} = \hat{C}_{it} - \mu_i \tag{4.15}$$

4.5 Results for the Improved Model

The improved random field model was fitted to the 13 gas concentrations given in Appendix B, with the results presented in Table (4.2). Column 4 in this table presents the POP from this model, while Column 6 indicates whether or not the plume (if it exists) was correctly located by the model. As one can see, the calculated POP (probability of a plume) is now much more reasonable. The strong gas plumes were assigned a POP of 1, indicating a high degree of confidence in a plume. Also, the strong plumes were correctly located by the procedure (see Appendix C).

The false call performance was dramatically improved; With the improved model, all "blank" (i.e. no plume) images were assigned a POP of 0.

However, the performance for the weak plumes is more disappointing; Only one of the three weak plumes (gas 6) was detected with a POP of 95%. One could argue that gas 6 shouldn't be scored as a detection, because it was not correctly located. The improved random field model isn't detecting the weak plumes identified by the human eye.

A review of the weak plume fit diagnostics indicates that one effect is causing the poor detection results; The plumes labeled as "weak" plumes in Table 4.2 could more reasonably be categorized as *intermittent* plumes. In other words, the plumes disappear in several of the time slices. For example, the three weak plumes in Table (4.2) only exist in about 40% of the time-slices. This feature could either be caused by variations in the speed of the wind, or by variations in the stack emission. At any rate, the random field model expects to find the plume in all time-slices and when the plume is absent in even one time-slice, it is assigned low probability of existence. Therefore, an assignment of a low POP to the weak plumes in the table is not unreasonable.

Appendix C displays plots of the Improved Random Field Model fits. If there is a total POP > 0.5%, The pixels with POP greater than 0.5% are marked with a plus. Appendix C displays log of the Bayes factor (equivalent to POP) and also the raw signal noise ratio (SNR) for the purposes of visual comparison.

	No. Pixels With		Random Field	True	Correct
Gas	SNR > 2	SNR > 4.5	Model POP	State	Location
1	701	50	0	no plume	-
2	3089	126	0	no plume	-
3	654	152	1	strong plume	Yes
4	579	183	1	strong plume	Yes
5	10	1	0	no plume	_
6	214	45	.92	weak plume	No
7	64	14	0	weak plume	No
8	45	11	0	no plume	-
9	103	27	0	no plume	_
10	27	11	0	weak plume	No
11	6293	87	0	no plume	-
12	2221	66	0	no plume	-
13	20	5	0	no plume	_

Table 4.2. Results from Improved Random Field Model

5.0 Conclusions

Simulations show that the random field model can produce dramatically improved results when the data meets the distributional assumptions of the data. However, the presence of a varying background in real data causes problems for the random field model as originally formulated.

An improved version of the random field model produces a detection statistic (POP) that gives reasonable results without the use of an empirical calibration, an indication that the improved distributional model fits the data reasonably well. This is in contrast to the use of SNR, which produces unrealistically high false call rates.

It is surprising that random field model was able to produce reasonable results using a plume library of only 9 shapes. A comparison of the plumes in the library (Figure (4.4) and those in the data Appendix B) show that real plumes are much more complicated than the idealized Gaussian Plumes in the library. Also, the behavior of the plume with time is more complicated than that assumed by the model; The real plumes are intermittent and consist of puffs that travel across the image. Use of a more extensive plume library might significantly increase probability of detection for weak plumes.

The test presented in this report shows that the random field model has promise; A more extensive test is required to show that it would work correctly on a broad spectrum of sites/gases. A more extensive test would concentrate on performance on weak plumes, and assure the "true state" is known by artificially injecting plumes in the data set. In a more serious test, a larger plume library should be used and the plume library should be empirically determined.

6.0 References

Berger JO. 1980. Statistical Decision Theory and Bayesian Analysis. Springer Verlag.

Burr T, B Foy, H Fry, and B McVey. 2006. "Characterizing Clutter in the Context of Detecting Weak Gaseous Plumes in Hyperspectral Imagery." *Sensors ISSN1424-8220*.

Burr T, H Fry, B McVey, and E Sander. 2002. "Chemical identification using bayesian model selection." *Spring Research Conference of Statistics in Industry and Technology*.

Burr T and N Hengartner. 2006. "Overview of Physical Models and Statistical Approaches for Weak Gaseous Plume Detection using Passive Infrared Hyperspectral Imagery." *Sensors ISSN1424-8220*.

Cressie N. 1991. Statistics for Spatial Data. Wiley.

Ferguson TE. 1967. Mathematical Statistics, A Decision Theoretic Approach. Academic Press.

Heasler P, C Posse, J Hylden, and K Anderson. 2007. "Nonlinear Bayesian Algorithms for Gas Plume Detection and Estimation from Hyper-spectral Thermal Image Data." *Sensors* 7(ISSN 1424-8220):905–920.

Kass RE and AE Raftery. 1994. Bayes Factors. 254, University of Washington.

Kil DH and FB Shin. 1996. *Pattern Recognition and Prediction with Applications to Signal Characterization*. AIP Press.

McLachlan GJ. 1992. Discriminant Analysis and statistical Pattern Recognition. Wiley.

Sander E. 2003. *The Application of Bayesian Model Averaging to the Detection, Identification and Quantification of Plume Gases*. DOE.

Winkler G. 2006. *Image Analysis, Random Fields and Markov Chain Monte Carlo Methods.* Springer Verlag.

Young S. 2002. Detection and quantification of gases in industrial-stack plumes using thermalinfrared hyperspectral imaging. ATR-2002(8407)-1, Aerospace Corp.

Appendix A

Calculation of Null and Alternative Distributions for Benchmark Scenarios

Appendix A – Calculation of Null and Alternative Distributions for Benchmark Scenarios

The following conditions have been selected to define the null and alternative distributions to determine detection performance of the random field model;

- Dimensions: There are $n_p = 64^2$ pixels in an image. The plume is of length $n_\ell = 2, 5, 10$. The site has been imaged $n_t = 1, 2, 3$ times. A space-time pixel in the data is indexed by *jk*, representing pixel *j* in time-slice *k*.
- The LWIR data from pixel *jk* has been summarized into a normalized signal, Y_{jk} , (for example, by application of a matched filter). When no plume is present, Y_{jk} is normally distributed with mean 0 and standard deviation 1. $(Y_{jk} \approx N(0, 1))$
- There is at most one plume at the site, located for convenience in the middle of the site. The plume is a single line of pixels of fixed length n_{ℓ} . In all time-slices of the site, the plume direction is the same.
- When a plume is present at jk, $Y_{jk} \approx N(\gamma, 1)$, where γ is the signal-to-noise ratio associated with the gases in the plume. The signal/noise ratio in an actual plume would vary, but for simplicity, we will assume it is fixed at *gamma*. Since we are mainly interested in the detection of weak plume, γ will be small, with $\gamma = 1, 2, 5$.

A.1 Detection Statistics to be Evaluated

The most obvious way to determine whether or not there is any evidence of a gas release from the images (without considering spatial/temporal information) is to look for the largest signals in Y_{jk} . If these signals are appropriately large, then one concludes the gas is present at the site and the plume location conforms to those jk where Y_{jk} is above a specified threshold.

This method of detection is equivalent to using the maximum detection statistic;

$$T_{max} = \max_{jk} Y_{jk} \tag{A.1}$$

where the index runs over the $n_{tot} = n_p n_t$ pixels in the data. We choose a threshold t, and call a detection at the site whenever $T_{max} > t$. We call this detection methodology the *Max detection* methodology.

Another popular variant of the detection strategy described above is to first smooth the signals Y_{jk} using a spatial filter that resembles the plume. If one did not know the shape of the plume, one would smooth using several different shapes and look for a "hit" in the resulting smoothed images. In other words, for this detection strategy, one smooths using a plume shape S, and then takes the maximum, so the detection statistic is;

$$T_{smooth} = \max_{jk} \mathcal{S} \otimes Y \tag{A.2}$$

Since a portion of the Bayesian algorithm involves smoothing, it is of interest to compare T_{smooth} against T_{bayes} ; It is possible that the good performance of T_{bayes} relative to T_{max} may be due to smoothing and if this were to be the case T_{smooth} would be a viable alternative to T_{bayes} . We call this detection methodology the *smoothed detection* methodology.

A.1.1 Null and Alternative Distributions

The null and alternative distributions can be explicitly calculated for T_{max} , since the data Y_{jk} is independent;

$$Pr(T_{max} < t | \text{Null}) = \Phi^{n_{tot}}(t)$$
(A.3)

and

$$Pr(T_{max} < t | \text{Alternative}) = \Phi^{n_{tot} - p}(t) \cdot \Phi^{p}(t - \gamma)$$
(A.4)

where *p* represents the number of pixels in the plume; $p = n_{\ell}n_t$. Using these distributions, we can calculate an ROC curve for T_{max} .

However, such formulas can't be produced for a smoothed version of Y, because the smoothed data points are no longer independent. To calculate the distribution of T_{smooth} we use simulation.

A.2 Bayesian Detection Statistic

The random field model developed in Section (2) does not "detect" plumes in the image, rather it provides a probability that plumes exist (i.e. site POP). However this probability can be easily used for detection by thresholding. In other words, we would say a plume exists at the site if Pr(plume exists|L) is larger than a specified threshold. An equivalent procedure is to threshold on the odds-ratio associated with this event. Since this happens to produce a simpler detection statistic, we define our Bayes detection statistic as follows;

$$T_{bayes} = \frac{Pr(\text{Plume})}{Pr(\text{No Plume})}$$

= $\sum_{j} R(j|L)$ (A.5)

with

$$R(j|L) = R(j|Y) = \frac{1}{\lambda^{n_t - 1}} \prod_k R(j|Y_k)$$
(A.6)

and

$$R(j|Y_k) = \sum_{i} R(S_{ji}|Y_k)$$
(A.7)

and

$$log(R(\mathcal{S}_{ji}|Y_k)) = log(h_i\lambda) + \sum_{j2} \mathcal{S}_{ji}(x_{j2})log(B_{jk})$$
(A.8)

For these equations to completely define T_{bayes} , we need to specify (1) the Bayes factor in terms of the data Y_{jk} , (2) the plume shapes S_{ji} , and h_i , and finally λ .

A.2.1 Bayes Factor

The Bayes factor for pixel jk, B_{jk} , is simply the alternative distribution, divided by the null distribution. Hence $log(B_{jk})$ is;

$$log(B_{jk}) = log\left(\frac{\frac{1}{\sqrt{2\pi}}\exp(-0.5(Y_{jk} - \gamma)^2)}{\frac{1}{\sqrt{2\pi}}\exp(-0.5Y_{jk}^2)}\right)$$

$$log(B_{jk}) = \gamma(Y_{jk} - 0.5\gamma)$$
(A.9)
(A.10)

A.2.2 Plume Shapes

We will assume that the plume library consists of n_d straight-line plumes of length n_ℓ , uniformly spaced between 0 and 90 degrees. As an example, Figure (A.1) illustrates the plume shapes for $n_{\ell} = 10 \text{ and } n_d = 5$

For this problem, S_{Oi} represents a bit-map in direction *i*. Under the alternative hypothesis, a plume does exist in the data, and it conforms to the "middle" shape (i.e. 45° shape) in the library. All directions will be assumed to be equally likely, hence $h_i = 1/n_d$. For this model, we will use $n_d = 3$.





A.2.3 Plume Rate of Occurrence

The plume rate of occurrence, λ occurs as a multiplicative factor in the formulas defining T_{bayes} and therefore does not influence the FCP/POD performance of this statistic. To simplify the formulas, we will set $\lambda = 1$.

A.2.4 Strategy for Calculation of Distributions

The null and alternative distributions are too difficult for analytic calculation, therefore, we produce them by simulation. To produce the alternative distribution of T_{bayes} ;

- 1. Simulate Y_{jk} from a N(0, 1) distribution
- 2. Calculate T_{bayes} from the simulated data
- 3. Repeat the two previous steps about a 1000 times and form the distribution $Pr(T_{bayes} < t | Null)$ from the population of T_{bayes} .

The alternative is produced in the same general fashion, except at the pixel locations in the plume, the Y_{ik} are chosen from a $N(\gamma, 1)$ distribution.

Appendix B

Plots of Single-Pixel Regression Fits to LWIR Data

Appendix B – Plots of Single-Pixel Regression Fits to LWIR Data

The images in this appendix plot $\hat{C}/sd(\hat{C})$ where \hat{C} represents a gas concentration estimate obtained by single-pixel regression of the LWIR data. For a single pixel, the LWIR data is an LWIR spectrum and the regression model attempts to describe radiation from the ground and the effect of atmosphere and plume on this radiation.

The plots are on *SNR* of gas concentration, and if the standard deviations in the denominators are correct, a value greater than about 4.5 in an image would be unusual when no plumes were present (i.e. would only occur about 1% of the time). For each gas, we have presented the SNR for 4 of the 10 time slices.







B.3



B.4









Appendix C

POP Plots from the Improved Random Field Model

Appendix C – POP Plots from the Improved Random Field Model

This appendix displays two plots for each random field model fit to the data (there is one fit for each gas). The first plot presents the log of the plume odds-ratio. This odds-ratio describes the probability that a plume exists at a particular location in the site. Low values indicate low probabilities, while high values indicate high probabilities of a plume (POP). These odds ratios are added up and normalized to produce the POP for the entire site, which is output in the figure heading.

The second plot presents the SNR $(\hat{C}/sd(\hat{C}))$ from the regression data (i.e. the input data to the random field model). The actual regression data consists of a SNR image for each of the 10 time-slices, which are too many images to present. The plot presented represents the maximum SNR over the 10 time-slices. Any locations having a *POP* > .001 are identified in both plots with a "+".







C.3








Appendix D

Monte Carlo Check of the Random Field Model

Appendix D – Monte Carlo Check of the Random Field Model

In order to verify that the Random Field Model was free of any programming bug, we constructed a simple Monte-Carlo simulation to see if it computed reasonable posterior probabilities (i.e. POP's). We found that it did seem to be calculating the correct probabilities, and furthermore was *extremely effective* at detecting weak plumes.

The Monte-Carlo consists of 5 test cubes containing gas concentrations contaminated with Gaussian error (mean 0, standard deviation 1) and each test cube consists of 3 time slices which contain a Gaussian plume in the lower-left hand corner as illustrated in Figure D.1. The plumes originate from the plume library used by the random field model, so the simulated and assumed plumes will match exactly. However, the random field model must select the correct plume shape out of nine possible.

The 5 test cubes contain plumes with increasing stack emissions of 0, 0, 0.5, 1, and finally 2. This implies that the first two test cubes contain no plumes and should produce a low POP, corresponding to a false call probability. In fact, it can be shown that the expected value of a POP associated with a "blank" test cube should be equal to the prior POP used in the Random Field Model. In this case, we have set the prior POP to 0.01, so we would expect a posterior POP of about 0.01 for the first two test cubes.



Because the plume shapes in Figure D.1 all have a maximum of 1 at the origin (which is also the location of the stack), using emissions of 0.5, 1, and 2 results in maximum SNR's of 0.5, 1, and 2 respectively for the test cubes that contain plumes. These are weak plumes, and one can visually see this by examining simulated versions of the data cubes presented in Figure D.2. The images in the Figure display the test cube with the strongest plume (that with a stack emission of 2); It would be difficult to visually detect the plume in these images. Only if an investigator performed the right transformation would he find the plume. For example an average of the three images, as illustrated by the 4'th image, makes the plume apparent.

A fit of the random field model to the test cubes produces the results present in Figure D.3. The images display the log of the plume odds-ratio, and the POP for the site is given in the image headings. As one can see, the POP associated with blank test cubes is low; In fact lower than the expected value of 0.01. Also, the random field model produces strong detection for the two



strongest plumes producing a POP of 1. From the odds-ratio plots, we also see that the plumes are correctly located.

The weakest plume (emission=0.5) does not show a strong detection; Its POP is comparable to the value one would expect from a blank image. However, it does seem to have located the plume correctly.





902 Battelle Boulevard P.O. Box 999 Richland, WA 99352 1-888-375-PNNL (7665) www.pnl.gov

