Algorithms for the Categorization and Identification of IR Military Signatures

Simon Turbide, Tracy Smithson, Daniel St-Germain and Pierre Fournier
Defence R & D Canada – Valcartier, 2459 Boul. Pie XI-Nord, Quebec (Qc), Canada, G3J 1X5

ABSTRACT

Military targets such as aircrafts and flares do not exhibit unique infrared signatures; their emissions are dominated by combustion products (mostly water vapor, carbon dioxide and hydrogen chloride) and hot metal greybody emissions. An algorithm has thus been developed to categorize target signatures based on their emission source components. The signatures are then partitioned, based on their emission components, into groups of similar emission characteristics. Using previous trial data, seven unique flare categories were defined. A second algorithm was finally developed to exploit this signature description and interrogate individual field measurements for target detection and categorization.

Keywords: Categorization, identification, infrared signatures, decoy flares.

1. INTRODUCTION

The DRDC Valcartier Spectral Imagery Laboratory currently supports two infrared ground based hyperspectral imagers designed for field measurements of military targets. The Passive Infrared Ranging And Target Evaluation System (PIRATES) and Baby-PIRATES have been used over the years for the infrared signature characterization of a variety of target types, resulting in a rich library of military related target signatures. For a previous international trial, Canada was tasked to collect the spectral signatures of the infrared (IR) decoy flares deployed by various air platforms. During the three-week measurement period of this trial, 24 different flare types were used, involving 13 aircraft (fighters, transport and helicopters) and resulting in a total of 1000 flare deployment events and a large number of descriptor sets. Each unique set of descriptor is defined by a Flare type, Aircraft used, Aircraft Aspect at time of deployment and Time-into-burn combination.

In a first instance, this paper addresses the algorithms that have been developed to characterize all these descriptor sets. In order to improve the reliability of the signature characterization, only the descriptor sets having at least three measurements were used. This has reduced the usable data set to 1864 signatures corresponding to 292 set descriptors. The goal of the work was to reduce this vast 292 dimensional parameter space to a more manageable set of flare "categories" based on the relative spectral characteristics of their IR emissions. The distinct flare categories found (distinct by their unique infrared signatures) have been implemented in a military target infrared signature library. In this study the absolute intensities of the flare emissions were not included.

Having at hand a library of potential military signatures, it was logical to aim at developing an algorithm allowing to detect and identify target present within the field of view of the sensor. Hence, any individual field measurement could be identified by its correlation with the signature categories included in the library.

The details of the classification algorithm will be the subject of section 2, while section 3 will briefly discuss the identification algorithm and a result example. Section 4 contains the summary.

2. FLARE CLASSIFICATION

2.1 Data processing

The flare signatures have been collected with the hyperspectral imager PIRATES. Its spatial resolution is provided by an 8x8 InSb detector array, covering the 1900-4800 cm^{-1} spectral band with a 5.5 milliradians (mrad) field-of-view. The spectral resolution was set at 16 cm^{-1} for this trial, corresponding to a measurement rate of 50 scans per second. Overall, the spatial extent was 2.5 degrees which corresponds approximately to the field-of-view of conventional IR seeker-

Imaging Spectrometry XIV, edited by Sylvia S. Shen, Paul E. Lewis, Proc. of SPIE Vol. 7457,
74570R · © 2009 SPIE · CCC code: 0277-786X/09/$18 · doi: 10.1117/12.828401

heads. The ranges of deployment, relatively to PIRATES, were approximately found between 750 m and 2000 m. Targets were tracked manually during this trial, permitting in some instances the measurement of emissions from multiple sources. This latter capability was exploited to significantly extend the single-flare deployment data set from under 100 to over 2000 entries. Each flare signature in this data set was labelled as follows:

- Air platform used (one of 13),
- Flare type (one of 24),
- Air platform aspect to observer at time of flare deployment (one of five possible: nose-on, 45 degrees off-nose, side-on, 45 degrees off-tail or tail-on),
- One of three times-into-burn (0.1, 0.2, or 0.5 second).

In order to obtain the final spectra for the study presented here, the following procedure was applied to each raw measurement (interferogram):

- apodization with a triangular function;
- Fast-Fourier Transform;
- phase correction;
- radiometric calibration.

A representative background spectrum was selected for each target and subtracted in order to generate the flare emission-only signatures. Each flare spectrum has been truncated (the strong carbon dioxide (CO_2) [2250-2380] cm^{-1} and water vapour (H_2O) [3500-3900] cm^{-1} atmospheric absorption bands have been removed): each spectrum is then described by N=517 frequency points. All signatures L are thereafter normalized according to :

$$\sum_{k=1}^{N} L_k = 1 \quad . \tag{1}$$

The intensity information is therefore not used in our classification.

2.2 Procedure outline

The flare categorization procedure can be simplified into two main steps: the extraction of representative coefficients and the classification of the flares based on those extracted coefficients. This is illustrated in Figure 1. The flare signature data set consists of 1864 measurements that can be initially grouped into 292 flare descriptors (i.e. a combination of flare type, platform, aspect angle and time-into-burn). Each flare descriptor has been restricted to contain at least three different measurements in order to take into account the possible fluctuations between the different measurements. In order to reduce this dimensionality further, we can apply algorithms that will find the lowest number of parameters to describe each spectrum, without any loss of information. This extraction of representative coefficients constitutes the dimensional reduction.

After the dimensional reduction step, as shown in Figure 1, each signature is represented by a set of W coefficients. This means that all flares from the data set are now associated to a point in a W-dimensional plot. A second algorithm is then applied to span this W-dimensional space and determines which points (signals) are close enough to be associated to the same category. Finally, the classification is done by associating each flare descriptor set to one of the X-categories found by the second algorithm. In other words, it means that many descriptor sets, among the initial 292, are undistinguishable. Therefore the classification regroups this descriptor set into X distinguishable classes (X << 292).

In the following subsections, the coefficient extraction method and the classification method will be shown, as well as the results obtained with those methods.

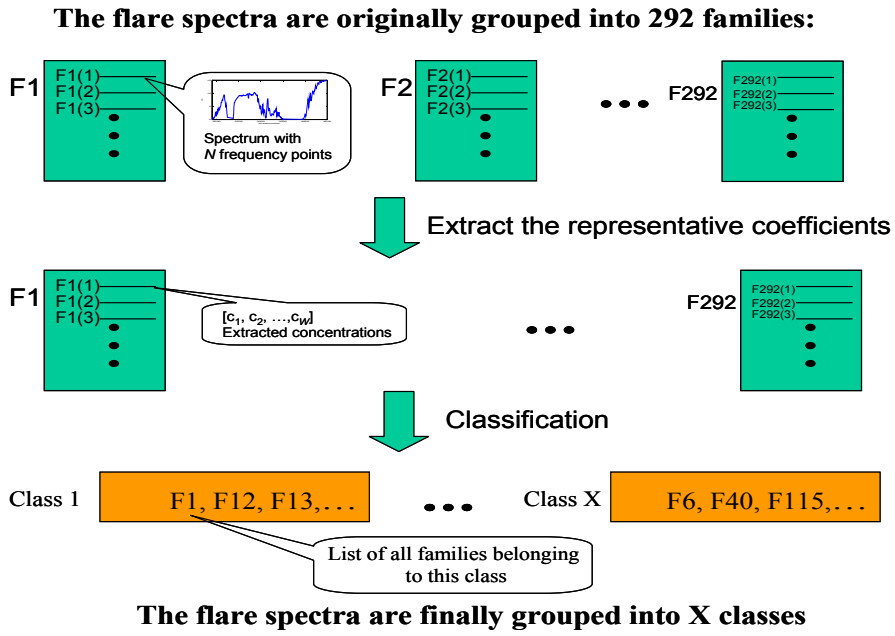

Figure 1 : Schematic plot of the flare classification

2.3 Coefficient extraction method

The technique that has been selected to extract the representative coefficients is Least Squares (LS). In this LS approach, the decomposition is done into a set of chemical elements that are included in the matrix *S*. The matrix *S* contains *W* spectra, each corresponding to a pure emission species. In order to describe as much as possible the emission pattern of the flare signature, the following constituent spectra are included:

- 3 blackbody (BB) spectra (T=700°C, 1400°C, 2500°C)
- 4 spectra of CO_2 (from hot to very hot)
- 1 spectrum of water (H_2O)
- 2 spectra of hydrochloric acid (HCl) (medium-hot and hot)
- 1 spectrum of carbon monoxide (CO)
- 1 spectrum of hydrocarbons (CH)

These spectra are illustrated in Figure 2. In order to describe the contamination of the spectra by sun glint, a typical sun glint spectrum has also been included in the constituent spectra (Figure 3). In a similar way typical background spectra could also have been included in the constituent spectra in order to isolate the flare signatures. However, for this data set the background contributions had already been removed.

We thus have W=13 spectra. The spectra displayed in Figure 2 have a 4 cm^{-1} spectral resolution and correspond to a measurement distance of 300 m. For the processing that follows, each of these spectra are degraded (by truncating the corresponding interferogram) to the flare signature resolution (16 cm^{-1}). The same procedure is also applied to the sun glint spectrum. Then, all constituent spectra are treated with the same apodization, spectral resolution and spectral grid than the flares signatures. Furthermore, an atmospheric transmission *Tr* is evaluated using the program PcModWin (MODTRAN version 3.7) and calculated for the range of the flare (minus the 300 m already included in the constituent spectra), and applied to each constituent spectra, excepted for sun glint. All constituent spectra are finally truncated (in order to remove atmospheric absorption regions) and normalized to one (see Eq. (1)).

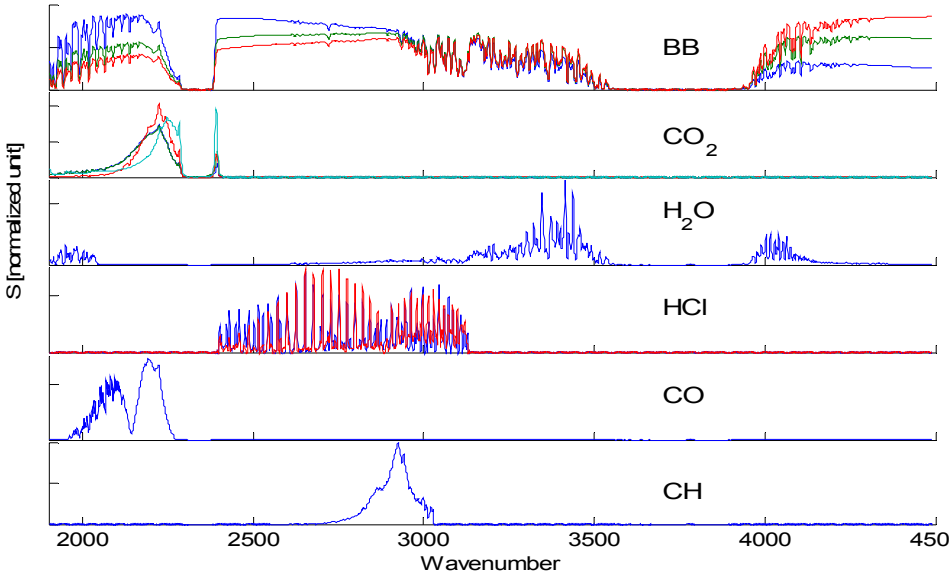

Figure 2 : Normalized constituent spectra that describe the flare signatures. The spectral resolution is 4 cm^{-1} and the range is 300 m

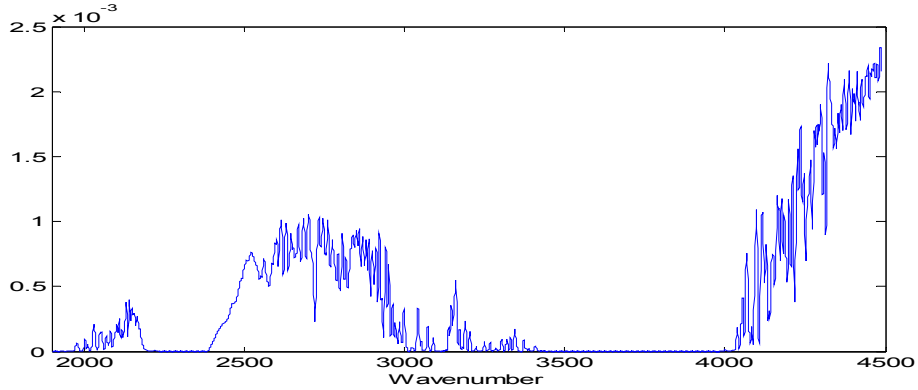

Figure 3 : Sun glint spectrum at 4 cm^{-1} resolution

A function G is then defined for each flare signature L:

$$G_k = c_1 S_{1,k} + c_2 S_{2,k} + \ldots + c_W S_{W,k} - L_k \,. \tag{2}$$

The chi-square function is thereafter defined by:

$$\chi^2 = \sum_{k=1}^{N} {G_k}^2 \,. \tag{3}$$

The G function is then minimized by looking at the minimum of the chi function, as a function of the individual relative concentrations:

$$\frac{\partial \chi^2}{\partial c_j} = 0 = \sum_{k=1}^{N} G_k S_{j,k} \,. \tag{4}$$

This corresponds to a linear set of equations that can be written into matrix form:

$$\begin{bmatrix} \sum_{k=1}^{N} {S_{1,k}}^2 & \sum_{k=1}^{N} S_{2,k}S_{1,k} & \cdots & \sum_{k=1}^{N} S_{W,k}S_{1,k} \\ \sum_{k=1}^{N} S_{1,k}S_{2,k} & \sum_{k=1}^{N} {S_{2,k}}^2 & \cdots & \sum_{k=1}^{N} S_{W,k}S_{2,k} \\ \cdots & \cdots & \cdots & \cdots \\ \sum_{k=1}^{N} S_{1,k}S_{W,k} & \cdots & \cdots & \sum_{k=1}^{N} {S_{W,k}}^2 \end{bmatrix} \begin{bmatrix} c_1 \\ c_2 \\ \cdots \\ c_W \end{bmatrix} = \begin{bmatrix} \sum_{k=1}^{N} S_{1,k}L_k \\ \sum_{k=1}^{N} S_{2,k}L_k \\ \cdots \\ \sum_{k=1}^{N} S_{W,k}L_k \end{bmatrix}. \quad (5)$$

The first matrix on the left hand side and the right hand side matrix are known. The concentration vector ${C^{LS}}_{i,1:W} = [c_1\ c_2\ \ldots\ c_W]$ is isolated from the equation using numerical algorithms: we solved this equation with the bi-conjugate gradient method [1]. The reconstructed flare signal is given by:

$${S^{LS}}_k = \sum_{h=1}^{W} {C^{LS}}_h S_{h,k}\ . \quad (6)$$

The parameter describing the efficiency of the fit is the residual between the flare spectrum and the reconstructed spectrum:

$$r_d = \frac{\sum_{k=1}^{N} \left| L_k - {S^{LS}}_k \right|}{\sum_{l=1}^{N} L_l}. \quad (7)$$

Since the range to the flare is not exactly known, an iterative procedure can be performed using Eq. (2) to Eq. (7), for different atmospheric paths (applying different Modtran transmissions *Tr*). The effective range selected is the one giving the smallest residual r_d. We highlight here the "effective" character of the evaluated ranges, since the MODTRAN transmission windows that are applied correspond to the standard configuration, which might not exactly reflect the real atmospheric conditions at the time of the flare deployments.

The sun glint contribution should not be included in the classification; its contribution is therefore removed from the coefficients and the remaining sets of coefficients are normalized:

$${C^{LS}}_{1:W-1} \rightarrow \frac{{C^{LS}}_{1:W-1}}{\sum_{j=1}^{W-1} {C^{LS}}_j}. \quad (8)$$

An example of the LS decomposition is shown in Figure 4. The residual of the fit is 0.069 in that example.

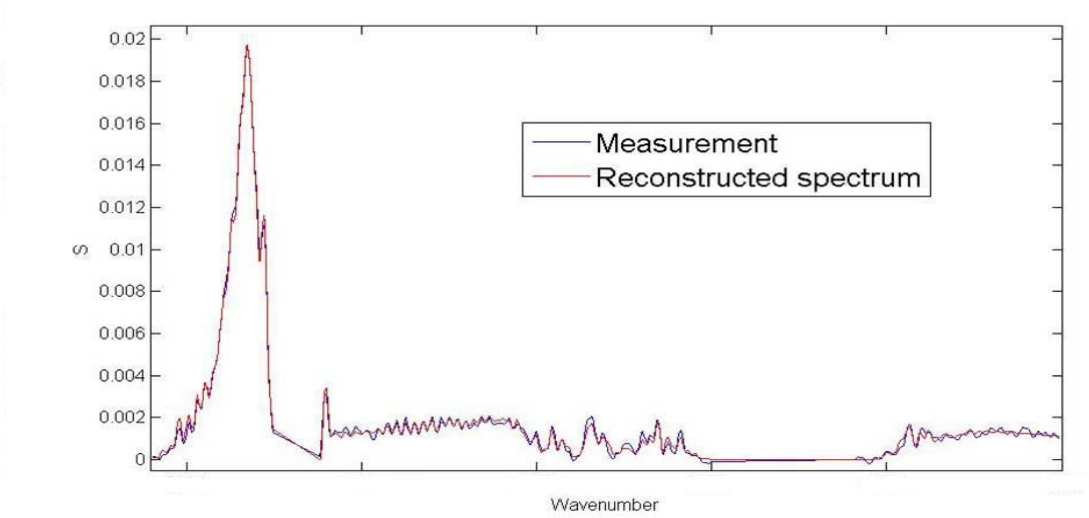

Figure 4 : Comparison of a flare signature (blue line) with the reconstructed spectrum (red line)

2.4 Classification method

The starting point is the coefficient vector $C^{LS}_{i,1:W-1}$, for each of the 1864 signatures from the data base ($1 \leq i \leq 1864$). The aim of each algorithm is to regroup all signatures into distinguishable clusters (classes). The efficiency of each algorithm is evaluated by the reproducibility factor r_p, which simply evaluate the cohesion of the classification. A perfect classification would correspond to r_p =100%. Inspired by the Gaussian Mixture model [2], we have developed the Empirically Constrained Gaussian Mixture (ECGM) model, in order to take advantage of the known signature information. This known signature information is the constituent relative concentrations and which signatures belong to each flare descriptor. The constituent concentration information is exploited by creating, from the start, different "worlds" that are defined by the presence or absence of HCl and CH, and whether the graybody[1] emissions dominate the signature (graybody concentration larger than 95%) or not. It was deemed unnecessary to define a "world" for CO_2, since the CO_2 and the graybody components are always present in the flare emission spectra and are, in some sense, complementary. Thus, by combining these different criteria, four possible worlds are defined. This is illustrated in Figure 5, step 2, although only the HCl, non-HCl and non-HCl graybody dominated worlds have been displayed (no flares have been associated to the other world).

As shown in the step 3 from Figure 5, the initial clusters in each world correspond to the flare descriptors associated to that world. The average concentration vector $\mu^{\{k\}}$, for each initial cluster k, is given by the average over the V_k measurements belonging to that flare descriptor:

$$\mu^{\{k\}}{}_{1:W-1} = \sum_j \frac{C^{LS}{}_{j,1:W-1}}{V_k}, \text{ with } j \in k . \tag{9}$$

[1] From this point, the graybody terminology will be used, rather than blackbody (BB). A BB has an emissivity of one, while a graybody has an emissivity smaller than one, but constant over all frequencies. The BB vectors included in the LS model are normalized; they represent basis vectors for the graybody emissions. The BB concentration extracted from the LS model represents the apparent graybody concentration, i.e the product of the real graybody concentration with the graybody emissivity, which are both unknown.

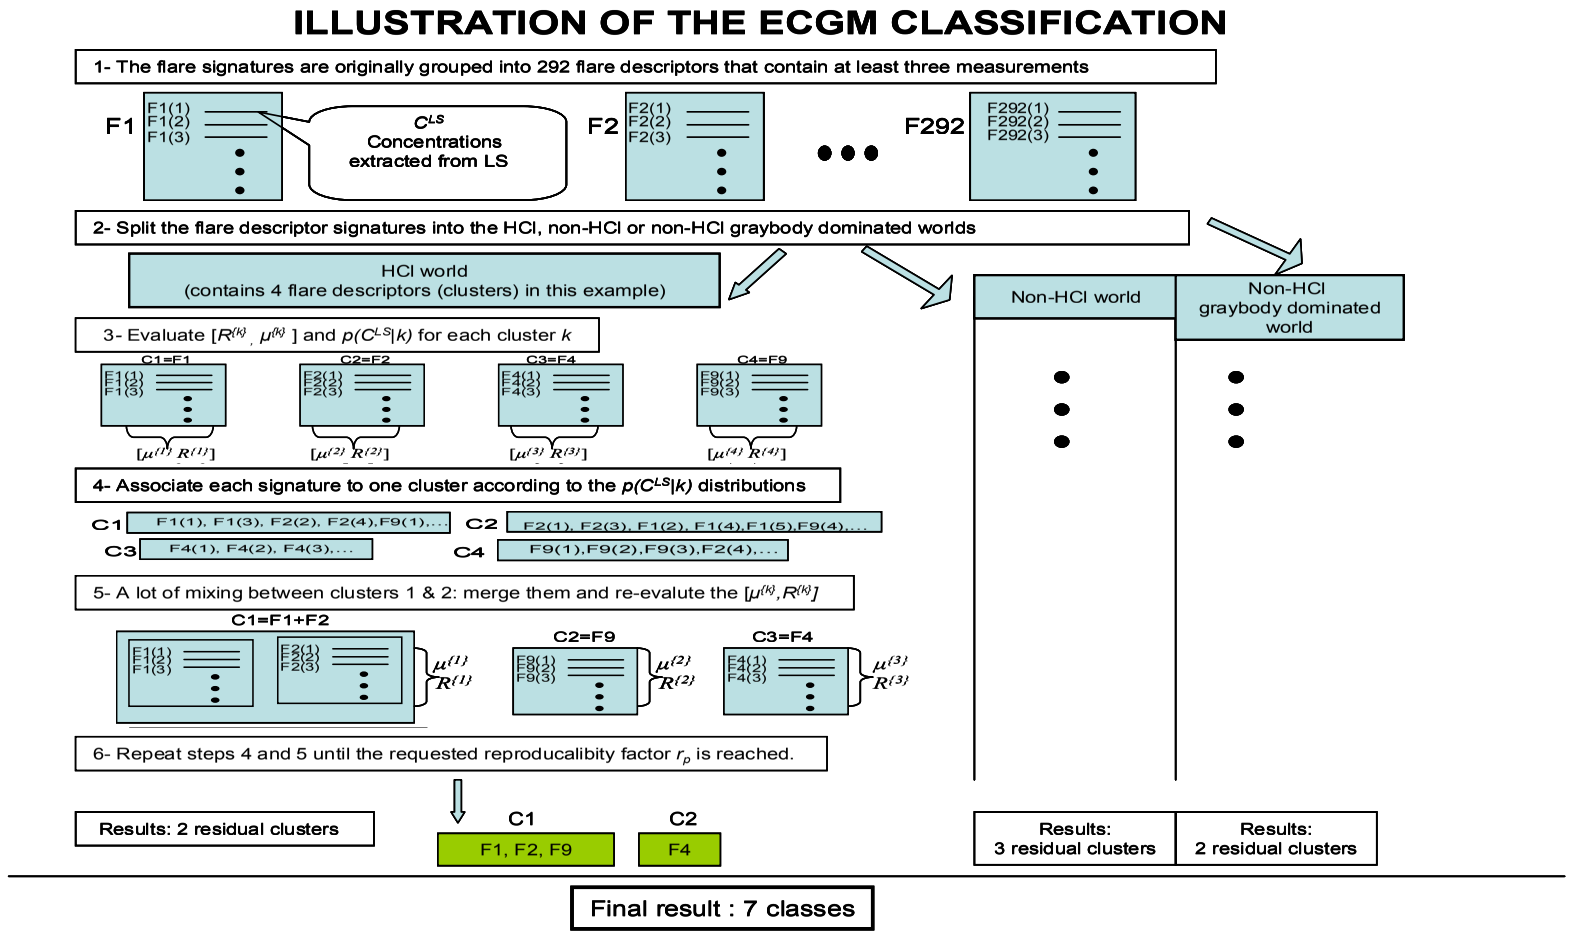

Figure 5: The Empirically Constrained Gaussian Mixture method

The covariance matrix, however, is evaluated differently since it might not be well defined in cases where the cluster contains only a few measurements. The covariance matrix for a cluster k is thus evaluated according to:

$$R^{\{k\}} = \left(R^{\{world\}} + I \frac{\sum_{l=1}^{W-1} R^{\{world\}}{}_{l,l}}{(W-1)10^5} \right) w + (1-w)\overline{R}^{\{k\}} + I \frac{\sum_{l=1}^{W-1} R^{\{world\}}{}_{l,l}}{(W-1)10^5}, \tag{10}$$

where I is the identity matrix, and $R^{\{world\}}$ and $\overline{R}^{\{k\}}$ are respectively the covariance matrices evaluated from all measurements belonging to the corresponding "world" and from all measurements belonging to the cluster k:

$$R^{\{world\}}{}_{i,j} = \sum_{l} \frac{\left(C^{LS}{}_{l,i} - \sum_{h} \frac{C^{LS}{}_{h,i}}{V_{world}} \right) \left(C^{LS}{}_{l,j} - \sum_{h} \frac{C^{LS}{}_{h,j}}{V_{world}} \right)}{V_{world}}, \tag{11}$$

and

$$\overline{R}^{\{k\}} = \sum_{j \in k} \frac{\left(C^{LS}{}_{j,1:W-1} - \mu^{\{k\}} \right)^T \left(C^{LS}{}_{j,1:W-1} - \mu^{\{k\}} \right)}{V_k}. \tag{12}$$

The sum in Eq. (11) runs over all V_{world} measurements belonging to the corresponding "world". The weight factor w is defined by:

$$\begin{aligned} w &= \exp\left(-\left(V_k - 10\right)/4\right) \text{ for } V_k > 10, \\ w &= 1 \text{ otherwise.} \end{aligned} \tag{13}$$

All signatures are then associated (step 4 in Figure 5) to the D cluster giving the highest association probability

$$C^{LS}{}_{i,1:W-1} \in D_j \quad \text{if} \quad p(C^{LS}{}_{i:W-1} \mid j) > p(C^{LS}{}_{i:W-1} \mid k) \quad \text{for } k=1,2,\ldots,K_0,\ k \neq j\ , \tag{14}$$

where the number of clusters is defined by K_0. The association probability is defined by:

$$p(C^{LS}{}_{i,1:W-1} \mid k) = \frac{1}{(2\pi)^{W-1}} \left| R^{\{k\}} \right|^{-1/2} \exp\left(-\frac{1}{2} (C^{LS}{}_{i,1:W-1} - \mu^{\{k\}}{}_{1:W-1})^T R^{\{k\}^{-1}} (C^{LS}{}_{i,1:W-1} - \mu^{\{k\}}{}_{1:W-1}) \right). \tag{15}$$

As mentioned before, the ECGM approach is inspired by the Gaussian Mixtures approach: the cluster distributions are assumed to be Gaussian and the number of clusters is reduced by one at each iteration by merging together two clusters. The difference between the two approaches lies in the way the initial clusters are chosen (they are arbitrary in the GM approach), and also the criteria for determining which pair of clusters will be merged together. In the ECGM approach, the selected pair of clusters minimizes the reproducibility factor r_p. In the example shown in Figure 5, after the step 4, a lot of mixing between the clusters 1 and 2 is observed, i.e, many signatures belonging to the first flare descriptor (first cluster) have been associated according to Eq. (14), to the second cluster, and vice-versa. In step 5, these two clusters will thus be merged together. At this point, the average concentration vector and the covariance matrix for this merged cluster are recalculated: the new concentration vectors and covariance matrices are calculated from the signatures belonging to the corresponding flare descriptors. In Figure 5 for example, it means that $\mu^{\{1\}}$ will be calculated from the signatures belonging to the flare descriptors F1 and F2, as shown after step 5. This subtlety will allow a more reliable estimate of the average concentration and avoid situations in which the average concentration vector could be affected by an incorrect association of a signature into the merged cluster.

Finally, by repeating steps 4 and 5, the number of clusters is reduced until a satisfactory reproducibility factor has been reached. The final number of classes is obtained by collecting the residual classes in each world. Each final class is

defined by a covariance matrix R and an average concentration vector μ. All 1864 signatures have been associated (according to Eq. (14)) to one of the final clusters (classes). Finally, each flare descriptor is associated to the class that contains most of its signatures.

Any new flare measurement can also be associated to one of these seven classes in the following way: extract the constituent concentration C^{LS} with the LS algorithm, then associate the signature to a class according to Eq.(14). However, if the new measurement contains some trace of a chemical that has not been included as a constituent vector, (HF as an example), the LS algorithm will not detect that chemical and might in fact try to fit the corresponding spectral features to other constituent spectra. In this case, the quality of the fit would be poor and would be indicated by a large residual r_d factor. In the situation of a large residual, a flag in the program will be raised, indicating to the user that additional constituent vectors are probably needed in the LS algorithm.

2.5 Results

Seven distinct classes were obtained using the LS algorithm for the extraction of the constituent relative concentrations, and the ECGM algorithm for the classification of the flare signatures. The 1864 signatures from the flare data set are plotted into their respective classes in Figure 6. These classes are distributed into three groupings: signatures containing HCl (classes 1 and 2), signatures without HCl content having a prominent CO_2 contribution (classes 5 to 7), and signatures without HCl content that are graybody dominated (classes 3 and 4). The threshold graybody concentration for the graybody dominated group was set to 95 %. The reliability of the classification of the flare signatures into these 7 classes is defined by a very acceptable reproducibility factor of 91%. The average relative concentrations for each class are finally stored in a signature library.

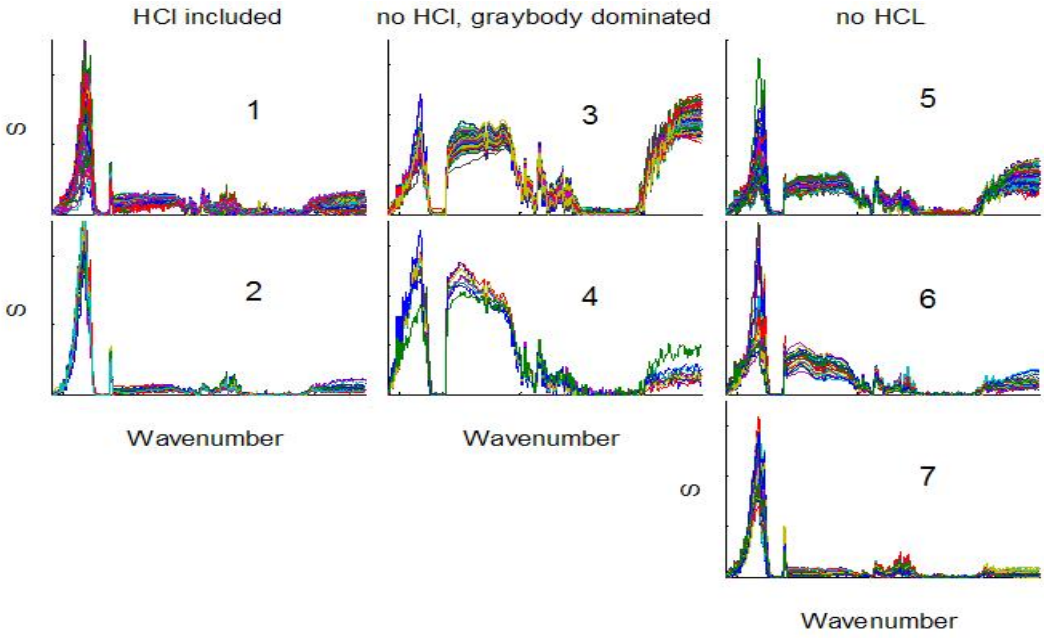

Figure 6: All classes found by the LS+ECGM classification method. The 1864 signatures are plotted into their corresponding classes. The seven classes more generally belong to three groups: HCl included (classes 1 and 2), no HCl and grabody dominated (classes 3 and 4), and no HCl CO_2 dominated (classes 5 to 7).

3. TARGET IDENTIFICATION

3.1 Procedure outline

In order to identify the targets present within the field of view of the detectors, three major steps were applied on each image: determine the non-background pixels, identify the content of those pixels, and source grouping.

At first, the background content in each pixel was evaluated. Knowing the shape of the background spectrum B (that has been evaluated from previous images containing no target information), the background content f in a pixel total signature S becomes:

$$f = \frac{S_{k_1} \sum_{k=1}^{N} B_k}{B_{k_1} \sum_{k=1}^{N} S_k}, \tag{16}$$

where k_1=floor((2200-σ(1))/$d\sigma$)+1, σ and $d\sigma$ being respectively the frequency axis and the spectral spacing (in [cm^{-1}] units). In a pure background spectrum, the residual (1-f) would be made of sun glint contribution. Denoting the normalized sun glint spectrum by g, the reconstructed spectrum is:

$$S^{reconstructed}{}_k = f\left(\frac{B_k}{\sum_{h=1}^{N} B_h}\right) + (1-f)g_k . \tag{17}$$

The residual between the measured and reconstructed spectra is:

$$res = \sum_{k=1}^{N} \left| S^{reconstructed}{}_k - \frac{S_k}{\sum_{h=1}^{N} S_h} \right| . \tag{18}$$

The measured spectrum is identified as a pure background spectrum if the residual is below a determined cut-off d_1. The evaluation of the residual can also be limited to the band [2150-2450] cm^{-1} where the background signal is stronger. In that case, a second cut-off d_2 is needed. We obtained good results with the combination d_1=0.15 and d_2=0.06. The pixels that do not fulfill these two conditions are therefore classified as being potentially target like. The background is thereafter subtracted from the spectrum S, and the result S_{noBg} is normalized according to Eq. (1).

The resulting signal should be made of a combination of target signals, sun glint and noise. All target signatures that are expected to be seen in the field of view are then theoretically reconstructed from the concentrations stored in the library (see previous section). A Least-Square model is then applied, like in section 2, in order to extract the relative contributions of each constituent present in the signal S_{noBg}; at the difference from section 2, the constituent elements are target signatures (aircraft and flares) rather than chemical signatures. Then, the output coefficients c_i of the LS model, Eq. (5), now corresponds to the relative concentrations of target signals and sun glint in the spectrum S_{noBg}. The residual between the spectrum S_{noBg} and the reconstructed spectrum (reconstructed from the projection of the c_i's on the known target signatures) is evaluated, in order to eliminate noisy spectrum.

Finally, as the target might cover more than one pixel in the field of view, a source grouping algorithm is applied. This last step is done in order to connect together, into a cluster, the neighbouring pixels containing the same type of target, and then reconstruct as much as possible the full signature of the targets.

3.2 Results

The target identification algorithm has been tested on flare deployments events (not the same events used to extract the 1864 signatures). An example is displayed in Figure 7. The results of the IR signature identification algorithm, for four different images of the same event, are displayed on the top panels. On the bottom panel are displayed the images of the visible camera that was set on PIRATES. The field of view of the camera is wider that that of PIRATES; in order to

compare both, an 8x8 grid has been drawn on top of the camera's image. In that example, four flares (two from class 1 and two from the class 3) were deployed from an aircraft moving side-on. The presence of an aircraft in a pixel is indicated by a green square, while the presence of a flare is indicated by a circle within the pixel. The size of the circle is proportional to the intensity of the signature within the pixel. The first flare of type F1 to be detected is indicated by a solid (blue) line while the second to be detected is indicated by a dash (cyan) line. The flares of type F3 are indicated by red and magenta lines.

On the first panel, the aircraft is detected in pixel 22. In the following frame, two deployed flares are visible. The second flare F1 is detected by the algorithm on the third panel, pixel 23. The algorithm has been able to detect two type of flares mixing in that pixel (a part of the first F3 flare is also visible in that pixel), because these two flares belong to two different classes, and so they are spectrally different. It's however impossible for the program to indicate the exact apparition time of this second F1 flare: it might have been already present in the second panel, overlapping with the first F1 flare, on pixels 28 and 29. Since these two F1 flares have the same spectral signatures, they can only be discriminated when their signatures cease to overlap. This is a limitation of the program. Finally, the second F3 flare is detected by the algorithm in the last panel, pixels 22 and 23.

The presence of a flare in a pixel can be validated by the images of the camera, although the alignment between the camera and PIRATES is not perfect due to a parallax problem. A flare appears on the camera as a bright spot; the F3 flares appear brighter than the F1 flares in the camera for that event. The flares in Figure 7 are then easily identified with the camera. However, it's impossible to discriminate the two types of flares overlapping in the third panel (pixel 23 top panel) using the camera, since it displays only the total intensity. This is the main advantage of the IR identification, as it is based on the spectral content of the signatures, which are still detectable even when they are mixing within a pixel. A future improvement of the target identification algorithm would be to incorporate as well the total intensity of all signatures as a possible discriminatory factor.

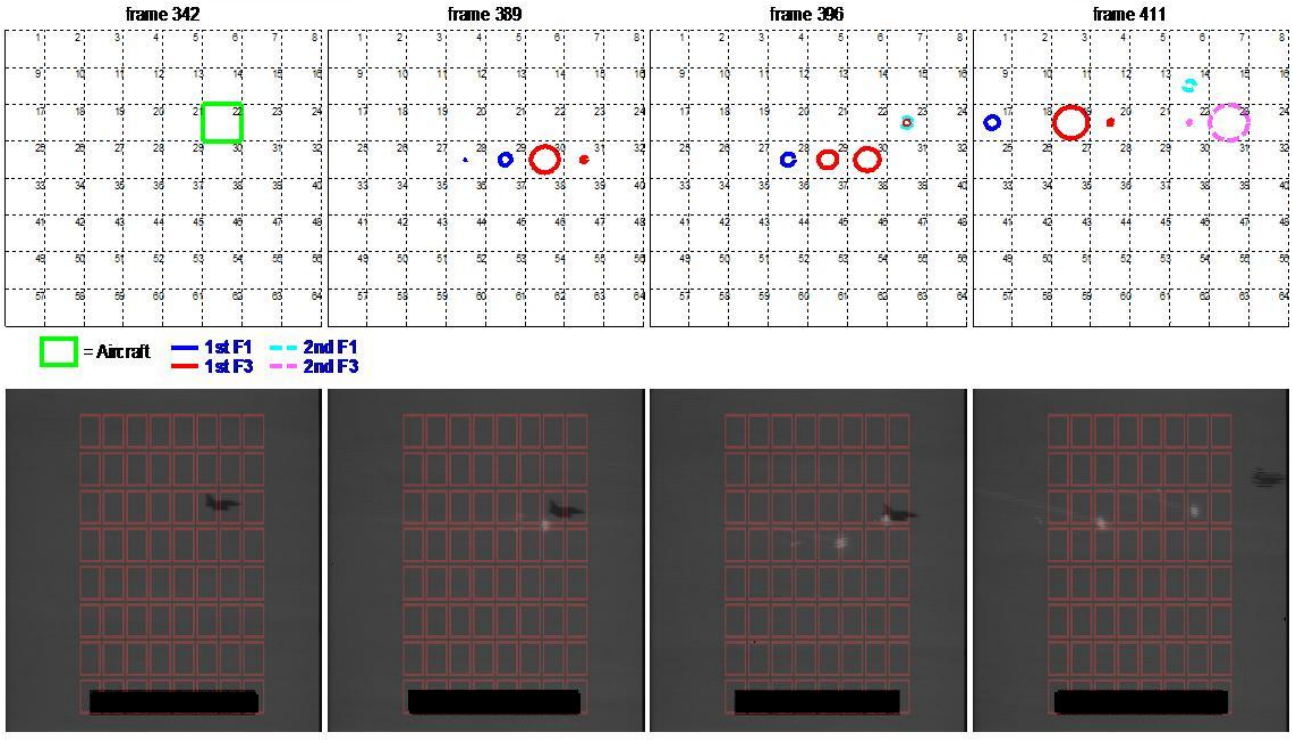

Figure 7 : Synchronization of the target identification result from PIRATES (top panels) with the visible camera (bottom panels). The field of view of the camera is wider than that of PIRATES. An 8x8 array (in red) has been displayed on the camera image in order to compare the two fields of view.

4. CONCLUSIONS

The LS method was applied on the flare IR signatures, collected during a previous international trial, in order to reduce the dimensionality of the data base and extract the representative coefficients. Those coefficients turned out to be relative chemical concentrations, such as CO_2, H_2O and graybody emissions. The ECGM method has been developed to regroup all flare signatures into the largest number of distinguishable classes, based on the extracted relative concentrations. By combining the LS and ECGM methods, it was possible to reduce the 292 flare groupings into seven flare categories for a reproducibility factor of 91%. A library of signature has then been created, from the average relative concentration of each category.

A second algorithm, also based on the LS model, was developed to exploit this signature description and interrogate individual field measurements for target detection and categorization. The algorithm is robust enough to track the individual space-time trajectory of multiple targets, with identical signatures or not, all moving within the field of view of the imager.

Although the categorization scheme has been developed for aircraft IR decoy flares, it could be extended to include other target types such as aircraft signatures, allowing for the partitioning of aircraft types. In addition to extending to more target types there are other spectroscopic parameters that also need to be investigated. In particular:

- the effect of spectral resolution (level of detail),
- the exploitation of temporal history (full flare burn profile, for example),
- spatially resolved targets and target-to-target separations (Spectral Imagery exploitation), and
- absolute emission amplitudes.

The signature characterization undertaken in this report was only possible because of the large number of IR signatures collected at the trial. It is necessary to have large coherent data sets in order to assess probabilities of categorization effectively and determine the spreads in observations. This can only be achieved with trials involving well-defined goals with controls on the experimental parameters and willingness to perform repetitions.

REFERENCES

[1] W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery, "Numerical Recipes in Fortran 77-Second Edition", Cambridge University Press, (1992).

[2] C. A. Bouman, "CLUSTER : An Unsupervised Algorithm for Modeling Gaussian Mixtures", https://engineering.purdue.edu/~bouman/software/cluster/manual.pdf, (2005).