Detection of methane plumes in hyperspectral images from SENTINEL-5P by coupling anomaly detection and pattern recognition

E. Ouerghi, T. Ehret, C. de Franchis, G. Facciolo, T. Lauvaux, E. Meinhardt, J. -M. Morel

DETECTION OF METHANE PLUMES IN HYPERSPECTRAL IMAGES FROM SENTINEL-5P BY COUPLING ANOMALY DETECTION AND PATTERN RECOGNITION

E. Ouerghi⋆,∗   T. Ehret⋆   C. de Franchis⋆,†   G. Facciolo⋆   T. Lauvaux‡   E. Meinhardt⋆   J.-M. Morel⋆

⋆Université Paris-Saclay, CNRS, ENS Paris-Saclay, Centre Borelli, 91190, Gif-sur-Yvette, France
† Kayrros SAS
‡ Laboratoire des Sciences du Climat et de l'Environnement, CEA, CNRS, UVSQ/IPSL

KEY WORDS: Methane, Hyperspectral, Pattern recognition, Time series, A contrario modeling, Anomaly detection.

ABSTRACT:

Reducing methane emissions is essential to tackle climate change. Here, we address the problem of detecting large methane leaks using hyperspectral data from the Sentinel-5P satellite. For that we exploit the fine spectral sampling of Sentinel-5P data to detect methane absorption features visible in the shortwave infrared wavelength range (SWIR). Our method involves three separate steps: i) background subtraction, ii) detection of local maxima in the negative logarithmic spectrum of each pixel and iii) anomaly detection in the background-free image. In the first step, we remove the impact of the albedo using albedo maps and the impact of the atmosphere by using a principal component analysis (PCA) over a time series of past observations. In the second step, we count for each pixel the number of local maxima that correspond to a subset of local maxima in the methane absorption spectrum. This counting method allows us to set up a statistical a contrario test that controls the false alarm rate of our detections. In the last step we use an anomaly detector to isolate potential methane plumes and we intersect those potential plumes with what have been detected in the second step. This process strongly reduces the number of false alarms. We validate our method by comparing the detected plumes against a dataset of plumes manually annotated on the Sentinel-5P L2 methane product.

1. INTRODUCTION

The detection of large methane (CH_4) leaks linked to oil and gas production is currently a major stake in order to reduce greenhouse gas (GHG) emissions. In a time lapse of 20 years, a CH_4 molecule has a global warming potential 80 times larger than carbon dioxide (CO_2) (Anthropogenic and Natural Radiative Forcing, 2013). A large part of the CH_4 emissions could be controlled or avoided, as they come from oil rigs and other oil and gas infrastructures. In order to detect GHG fossil fuel emissions produced by human activities, several satellites have been placed in orbit around the Earth over the past ten years. Here, we focus on the data provided by Sentinel-5P, launched in 2017 by ESA. Sentinel-5P provides hyper-spectral images in wavebands for which CH_4 has a significant absorption coefficient. Data from Sentinel-5P is publicly available and is already being used by ESA to quantify CH_4 emissions and other greenhouse gases (Pandey et al., 2019).

CH_4 detection belongs to the traditional field of anomaly detection in hyper-spectral imagery. Usually, detection is performed on infrared wavelengths (Scafutto and Filho, 2018, Crevoisier et al., 2009). Although the absorption features of CH_4 and water vapor (H_2O) sometimes overlap in the infrared spectrum, their separation can be addressed by a wise selection of wavelengths (Crevoisier et al., 2009).

A classic anomaly detection strategy consists in performing a background subtraction - for example with a principal component analysis (PCA) (Reitberger and Sauer, 2019) - then locally modeling the remaining residual image as following a Gaussian model. This probabilistic model then enables Neyman-Pearson tests to be carried out on the pixels to detect anomalies (Manolakis and Shaw, 2002, Matteoli et al., 2010).

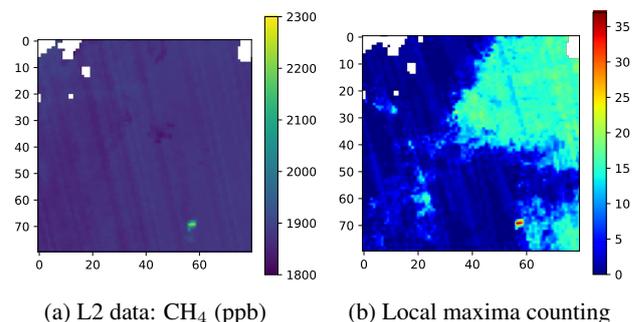(a) L2 data: CH_4 (ppb)   (b) Local maxima counting

Figure 1. Figure (a) shows the concentration of methane in parts per billion (ppb) given by the L2 methane product from Sentinel-5P. A methane plume is visible in green. Pixels in white are discarded pixels (cloudy for example). Figure (b) gives the number of CH_4 spectrum maxima counted for each pixel in the original image with the proposed method.

In (Scafutto and Filho, 2018), a CH_4-specific technique is developed where the background is removed by precise atmosphere modeling. Then, CH_4 is detected by a matched filter using the CH_4 absorption spectrum. Another matched filter is used for CH_4 detection in the Cluster-Tuned Matched Filter (CTMF), originally dedicated to sulfur dioxide (Funk et al., 2001), and then applied to CH_4 detection in (Thorpe et al., 2013). The idea proposed by (Funk et al., 2001) is to use a modified version of a k-means clustering prior to the matched filter. Spatial clustering allows to obtain results over very large areas without the risk of confusing anomalies due to CH_4 with those due to other gases.

Another well-established technique in remote sensing is the band ratio technique. This method is quite usual for glacier monitoring (Citterio et al., 2010) but is also used for hydrocarbon

---

* Corresponding author, elyes.ouerghi@gmail.com

detection (Garain et al., 2019), and is highly-relevant to our goals (Varon et al., 2020, Bradley et al., 2011). Band ratios are used to enhance the spectral signature of hydrocarbons. In (Garain et al., 2019), those band ratios are used jointly with a PCA. PCA of several spectral bands is computed in order to highlight the presence of hydrocarbons in the resulting principal component image. Then, all the computed images are superimposed to perform the detection.

Many methane monitoring approaches aim at quantifying methane emissions. Besides that, we want to have a daily monitoring of methane emissions. By covering the entire earth once a day, data from Sentinel-5P satellite allow us to do that.

Our objective is to introduce a flexible CH_4 emission detection method using the level 1 (L1) data provided by Sentinel-5P. We do so by using the fine spectral sampling of Sentinel-5P data. We detect local maxima in the negative logarithmic spectra of pixels that correspond to maxima in CH_4 absorption spectrum. Observing an excess of such maxima in the same pixel after background subtraction should be the consequence of CH_4 emissions (see Figure 1). Our contribution is to use this count of spectrum extrema as a cue, to find statistically founded detection thresholds. In that way we reduce considerably the false alarms and obtain a proof that the methane fingerprint is indeed present in the pixel after background subtraction. We complete the method with a classic anomaly detector. We use the Reed-Xiaoli algorithm (RX) with two different sets of wavelength bands (one sensitive to methane, the other not) to detect independently potential methane plumes. By intersecting those potential plumes with those detected by the local maxima method, we still reduce the number of false alarms. To validate the method, we shall use a set of manually detected plumes by experts in the L2 product.

2. MATERIALS

We use hyper-spectral images from Sentinel-5P. This satellite provides a dense spectrum (nearly 4,000 wavelengths) for each pixel and covers the entire Earth once a day. All the available wavelengths are organized in eight wavelength bands. Each band is composed of a given number of channels. A channel corresponds to an image at a precise wavelength. Thus, each band can be seen as a hyper-spectral image. Here, we use the SWIR bands 7 and 8 which cover the 2,300-2,389 nm range where the main absorption features of CH_4 are located. These images are part of the level 1 (L1) product.

We also use the level 2 (L2) fully-processed data including cloud maps, albedo and X_{CH4} column mixing ratios. Cloud maps are necessary because clouds preclude any CH_4 detection. X_{CH4} images will be used to identify plumes of CH_4 only for validation purposes.

It should be noted that the L2 product provides a quantification of methane but not a plume detection. Therefore it is not a proper ground truth for plume detection. To get a methane plume mask from the L2 product, a detection must be performed on the L2 product. This detection, performed manually or automatically, can introduce its own bias in the results, with false alarms or under-detections. For example, as we will see in Section 5, the L2 product can yield high methane concentrations on some areas, which are unlikely to be real methane sources; but a detection algorithm based on the L2 product might detect those areas as methane plumes.
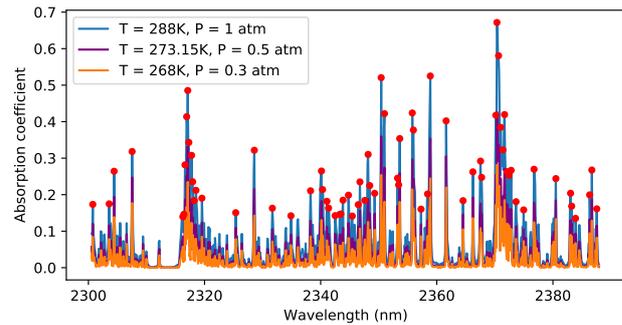

Figure 2. Comparison of methane absorption spectra under several pressure/temperature conditions. Red dots highlight the 70 maxima we use for detecting the plumes.

Lastly, we use a detailed CH_4 absorption spectrum taken from the HITRAN spectral database (Gordon et al., 2017). CH_4 spectrum varies depending on pressure and temperature. However, here we are only interested in CH_4 spectrum within the bottom layers of the atmosphere (below 1,500 meters above ground), because we want to detect CH_4 shortly after being emitted, before it rises and dilutes.

Moreover, spectrum variations are small in the conditions of the near-surface atmospheric layers. Moreover, here we use the CH_4 absorption spectrum only for detecting local maxima, and as we can see from Figure 2, its profile practically does not change for the near-surface conditions. Here, we selected the CH_4 spectrum at $15°C$ and 1 atm to represent near-surface atmospheric conditions.

3. MODELING

We shall use a simplified atmospheric absorption model to explain the value of each pixel of an image. Let us consider a pixel $\mathbf{P}$ in a Sentinel-5P image. Such a pixel is a vector in $\mathbb{R}^d$ where d is the number of channels in the image (for our case 960 channels from bands 7 and 8). Each pixel component corresponds to a wavelength λ.

Our model takes into account: the effect of sun irradiance $F_I(\lambda)$, the albedo A, the absorption coefficients of the dry atmosphere $K_{atm}(\lambda)$, water vapor $K_{H_2O}(\lambda)$ and methane $K_{CH_4}(\lambda)$. We denote by e_{gas} the thickness of gas crossed by the radiation before reaching the sensor. Making implicit the dependence on λ we can write the absorption model for the whole vector $\mathbf{P}$ as

$$\mathbf{P} = A \exp(F_I - K_{atm}e_{atm} - K_{H_2O}e_{H_2O} - K_{CH_4}e_{CH_4}). \quad (1)$$

Similarly to other works (Coakley, 2002), we assume here that the albedo is roughly constant over the part of the infrared spectrum we use, as it is extremely regular near 2000nm (Montpetit et al., 2012). We take into account the absorption by the dry atmosphere, as a single gas whose absorption spectrum is well known. This spectrum includes absorption from methane that is always present in the atmosphere. The term $K_{CH_4}e_{CH_4}$ represents the excess of emitted methane over the one already present in the dry atmosphere.

From now on, we shall work with $-\log(\mathbf{P})$ instead of $\mathbf{P}$. This allows us to have a linear model where excesses of methane comes out positively. We denote by P_0 the new pixel value.
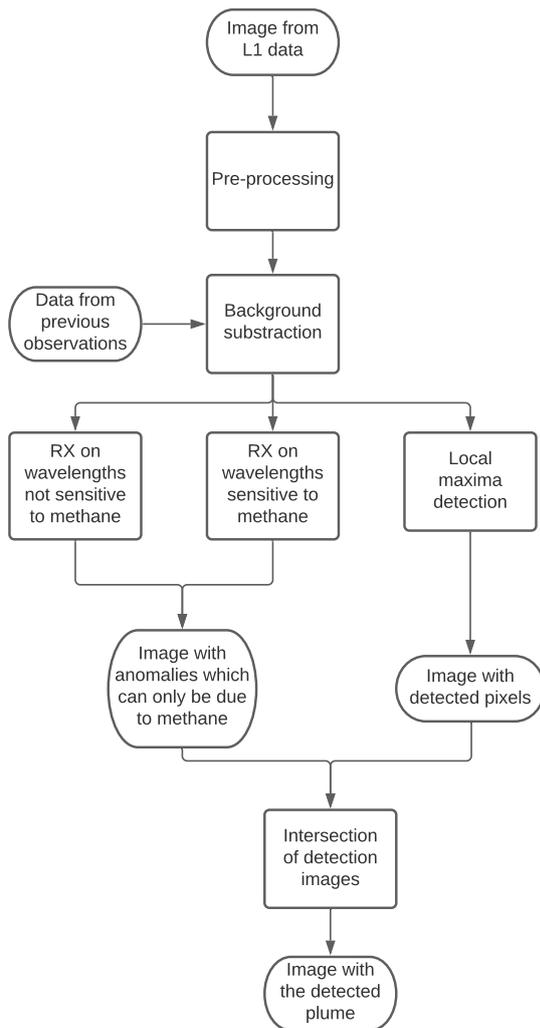
Figure 3. Diagram summarizing the method.

4. METHODS

Our methane plume detection method consists of three steps: first a background subtraction to remove the contribution of albedo and atmosphere, then a counting step where we count for each pixel the number of local maxima that correspond to a shortlist of local maxima in methane absorption spectrum, and finally Reed-Xiaoli algorithm (RX) is used on the background-free image to isolate potential methane plumes. We use RX both to highlight pixels which might present an excess of methane and to remove anomalies which are not due to an excess of methane. A diagram summarizing the proposed method is shown in Figure 3.

4.1 Background subtraction

Our observations are dominated by the absorption of the atmosphere and the albedo. In order to detect potential local excess of CH_4, we start by performing a background subtraction. This background subtraction has two advantages. First, it removes the contribution of albedo and atmosphere from the spectrum of the current pixel. Second, it sets the mean CH_4 concentration to zero. Indeed, there is a nearly-constant concentration of CH_4 in the atmosphere and CH_4 plumes rarely exceed 3% of this concentration. So, we must make sure that this mean CH_4

concentration is completely removed during background subtraction. To do so background subtraction applies three steps.

I. Albedo removal. First, albedo values from the L2 data are used to remove the albedo component from each pixel. Given a pixel P_0 from a pre-processed hyperspectral image I and the albedo A corresponding to this pixel, we compute the albedo-corrected pixel

$$P_1 = P_0 + \log(A). \tag{2}$$

The albedo is assumed to be identical for each channel, as variations of albedo are minor in the infrared spectrum (Coakley, 2002, Montpetit et al., 2012). After this first subtraction, P_1 contains only contributions from irradiance, atmosphere, water vapor and CH_4.

II. Atmosphere removal. For removing the contribution of the atmosphere we assume the irradiance and the absorption spectrum of the dry atmosphere to be roughly constant over a short period of time (in practice this analysis is performed over two weeks or less). Therefore, we can estimate those two components using a time series. For each pixel P_1, we gather observations $X_1, ... X_n$ of the same area at earlier dates and without clouds. The background is then modeled as the principal component of $X_1, ..., X_n$, which we denote F. To remove the background of P_1 we then remove its projection on the subspace directed by F, i.e.

$$P_2 = P_1 - \langle P_1, F \rangle F. \tag{3}$$

In practice, we perform this substraction on a pixel P_1 only if we can gather at least ten observations for the time series, otherwise we discard this pixel. If pixels in previous dates contain excesses of methane, it could affect the principal component F and the background substraction might remove a potential excess of methane in P_1. With a time series long enough, an excess of methane on one or two dates should not impact the principal component.

Furthermore, we take our observations $X_1, \cdots, X_n$ only from images where there is at most 50% of clouds. If in several images there is more than 50% of clouds, a part of the image under test might have a background quite different from the rest. Thus, the background will not be homogeneous and the background-free image will not be either. This leads to poorly estimated parameters for our model. A direct consequence of this is that the detection threshold will be too high or too low which can lead to under-detection or over-detection.

III. Methane equalization. The last part of the background subtraction is an equalization of the level of CH_4 in the current image. This subtraction works both spatially and spectrally. After the second part of the background subtraction, only CH_4 and water vapor should be left. However, this is not enough to detect CH_4 anomalies. Indeed, there are about 1900 ppb (particles per billion) of CH_4 in the atmosphere, but we want to detect variations in the order of 40 to 80 ppb. When the background is subtracted, a variable fraction of those 1900 ppb are removed, depending on the atmospheric CH_4 concentrations of earlier observations. This difference in background CH_4 can prevent a detection using local maxima. To address this issue we first compute a spatial average of CH_4 concentration M by projecting each pixel on the CH_4 direction

$$M = \frac{1}{|I|} \sum_{P \in I} \left\langle P_2, \frac{K_{CH_4}}{\|K_{CH_4}\|} \right\rangle. \tag{4}$$

Then, we remove this mean M from each pixel

$$P_3 = P_2 - M \times \frac{K_{CH_4}}{\|K_{CH_4}\|}. \tag{5}$$

After this last operation each pixel should only display a mix of water vapor and excess CH_4. Thus, CH_4 detection should be possible when the concentration of water vapor is not too high.

(a) L2 data: CH_4 (ppb)

(b) Original image
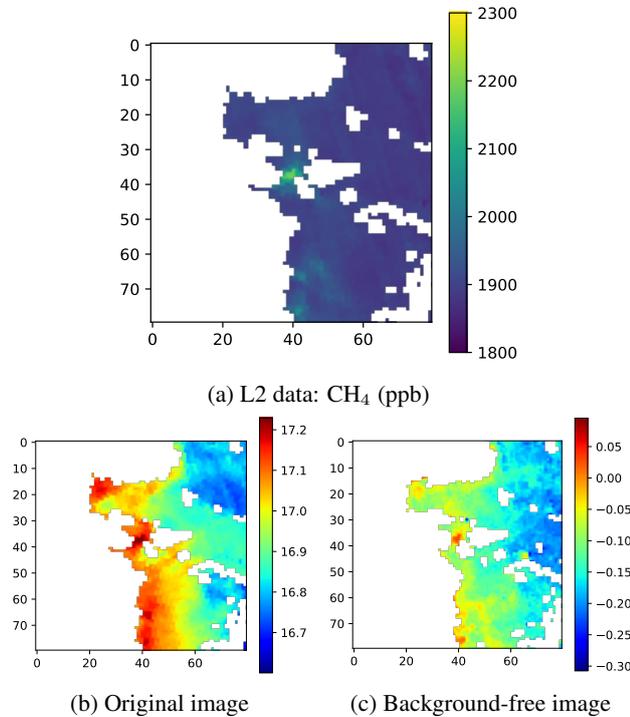
(c) Background-free image

Figure 4. In the image (a) we see a methane plume in green. Images (b) and (c) show the channel at wavelength 2370.448nm before and after background substraction respectively.

We can see on Figure 4 an example of result of the background substraction. On image (b) we have the L1 data in negative logarithmic from the channel at wavelength 2370.448nm and on image (c) the same image but after background substraction. We can see that background substraction highlights the methane plume. Even if the plume is slightly visible on image (b), the contrast between the plume and the background is far better on image (c). In particular, in image (c) the local maxima are almost only excesses of methane as seen in image (a).

4.2 Local maxima detection

Once the background has been removed, the spectrum of a pixel should mostly be composed of excess CH_4, water vapor and sensor noise. CH_4 and water vapor have different spectra in the wavelengths we are interested in. Thus, we expect that the wavelengths where CH_4 absorption is maximal, should also appear as maxima in the negative logarithmic spectrum of a pixel with excess CH_4. Otherwise, there should be only noise on those wavelengths. Therefore, to detect CH_4 plumes, we count for each pixel the number of local maxima coinciding to local maxima in the CH_4 absorption spectrum. We also refer to this step as pattern recognition as the goal is to recognize the specific absorption pattern due to methane.

For this, we cannot use all of CH_4 local maxima. Some of them correspond to small absorption coefficients, which could be eas-

ily confused with noise. Thus, we can only use the highest maxima in the methane spectrum. We selected the 70 highest maxima in the CH_4 absorption spectrum between 2300nm and 2380nm. This number of maxima is the one that empirically gives the best results with the proposed method. With more than 70 maxima we end up with wavelengths for which methane has a low absorption coefficient and with less than 70 maxima the gap between methane plumes and background can be too small to perform detection without false alarms. We can see the chosen maxima in Figure 2. In addition, we define two adapted thresholds for the detection of each of these maxima.

The first threshold $\tau_1(P)$ is the median of the spectrum of the pixel P; this prevents low maxima from being detected. Since we use only the 70 highest maxima of the CH_4 spectrum, we should also have high maxima in P. So this threshold will not hinder the detection of CH_4-related maxima.

We then compute a threshold adapted to each of the 70 chosen maxima. With the first threshold some maxima can appear in almost every pixel in the image, just because a specific wavelength usually shows high values. For the wavelength λ we set the threshold $\tau_2(\lambda)$ as the 70% quantile of all the values of the image at that wavelength; i.e for an image I the 70% quantile of $\{P_\lambda \mid P \in I\}$. To summarize, for a maximum at wavelength λ in a pixel P, our detection threshold is $\max(\tau_1(P), \tau_2(\lambda))$.

We need to define a final decision threshold to tell apart CH_4 plume (excess CH_4) pixels from background pixels based on the maxima that were counted. In order to do this, we rely on an *a contrario* model. We take the *a contrario* assumption that the image has no excess CH_4, and compute the probability of false detection under this assumption (Desolneux et al., 2003). To do so, we index the 70 highest CH_4 spectrum maxima by i, going from 1 to 70, and we denote by p_i the empirical probability that a spectrum maximum occurs at i in a "normal" image. If CH_4 anomalies occur in the image under study, they are generally concentrated on very few pixels. Hence we can estimate p_i from the image itself, and this will lead to a tiny overestimation of this probability if some pixels have excess CH_4. The random variable X_i, which is equal to 1 if the i-th maximum appears and 0 otherwise, follows a Bernoulli distribution with parameter p_i. To complete the *a contrario* model, we assume that $X_1, ..., X_n$ are independent (in the absence of CH_4). We denote by $S(P) = X_1 + ... + X_n$ the number of counted maxima on a given pixel P. We then compute a detection threshold τ which guarantees a given false alarm rate p_{fa}

$$p_{fa}(\tau) = \mathbb{P}(S(P) > \tau | P \text{ without excess } CH_4). \tag{6}$$

The value of p_{fa} is set to 10^{-6} in our experiments, which amounts statistically to less than 0.01 false alarm per image.

4.3 Anomaly detection

To further reduce the false alarms, we are going to combine the counting of local maxima with an anomaly detection algorithm. The anomaly detection algorithm used here is the classic Reed-Xiaoli algorithm (RX) (Matteoli et al., 2010). The goal of this algorithm is to detect anomalies in hyper-spectral images. The algorithm learns a model around a pixel of interest then checks if this pixel follows the model. The input of this algorithm is an image after the background subtraction described above.

4.3.1 The Reed-Xiaoli algorithm. For each pixel, RX first estimates a model from its neighbors. It assumes that all pixels

in a neighborhood of the pixel under test (PUT) are independent and stem from the same random variable which follows a multivariate normal distribution. Those pixels are used to compute the parameters of the normal law. Here, the neighborhood is a rectangular block centered on the PUT, deprived of a guard window. The size of the block is a parameter which needs to be chosen adequately; it depends mostly on the image resolution and on the size of the anomaly we want to detect. Inside the block centered on the PUT, we need to select the pixels used to compute distribution parameters. The guard window is a smaller block centered on the PUT. Thus, the Gaussian parameters are computed with pixels outside the guard window. As a rule of thumb, the guard window should be larger than the expected anomaly size.

We compute the empirical mean μ and covariance matrix C of the multivariate distribution with their usual non biased empirical estimators. Then we compute the likelihood of each pixel with respect to the multivariate distribution. Finally, we compare the likelihood of each pixel to a detection threshold η which depends on μ and C. The detection criterion on the likelihood for a pixel P can be written as follows :

$$\text{Pixel } P \text{ is an anomaly iff } (P - \mu)^T C^{-1}(P - \mu) > \eta. \quad (7)$$

We want a detection threshold τ which guarantees a given false alarm rate p_{fa}. For a given threshold η, we denote by H_0 the hypothesis : the pixel P follows the model ; and H_1 the hypothesis : the pixel P is an anomaly i.e $(P-\mu)^T C^{-1}(P-\mu) > \eta)$. The probability of false alarm is given by $\mathbb{P}(H_1|H_0)$. By imposing a probability of false alarm per pixel, we can compute the corresponding detection threshold.

Unfortunately, this method has some drawbacks, quite well detailed in (Matteoli et al., 2010). These drawbacks caused by fixed shape learning detection windows can lead to over-detection and sometimes to under-detection. To address this issue, we use RX twice : a first time to detect potential methane plumes and a second time to remove over-detection caused by other parts of the spectrum.

4.3.2 Methane specific anomaly detection. Indeed, RX detects anomalies but does not decide if they are due to local methane excess or to other causes. But we can adapt RX to detect specifically methane emissions. The method described below is used to reduce the number of false detection and should only be used jointly with an other detection method such as the pattern recognition method described in Section 4.2.

The main idea is to execute the RX algorithm several times but with disjoint groups of channels. First we apply RX with wavelengths for which the methane absorbs practically no radiation. We obtain a first binary mask A_0. Then, we apply RX with channels that are the most sensitive to methane. We obtain a second mask $A_{methane}$.

Anomalies in $A_{methane}$ are most likely due to methane and anomalies in A_0 are most likely not. However some anomalies detected in $A_{methane}$ are also detected in A_0. Therefore, we can assume that those anomalies are not due to methane. Indeed, methane has no effect on channels used to compute A_0. So in $A_{methane}$ we remove anomalies detected in A_0. Finally, we multiply the resulting mask $(A_{methane} - A_0)^+$ with the anomaly mask obtained by pattern recognition.

In practice, this process systematically reduces the number of false alarms. However, this method can only be used in con-

(a) L2 data: CH$_4$ (ppb)

(b) Local maxima counting   (c) Detection with maxima only

(d) Image A_0 (e) Image $A_{methane}$

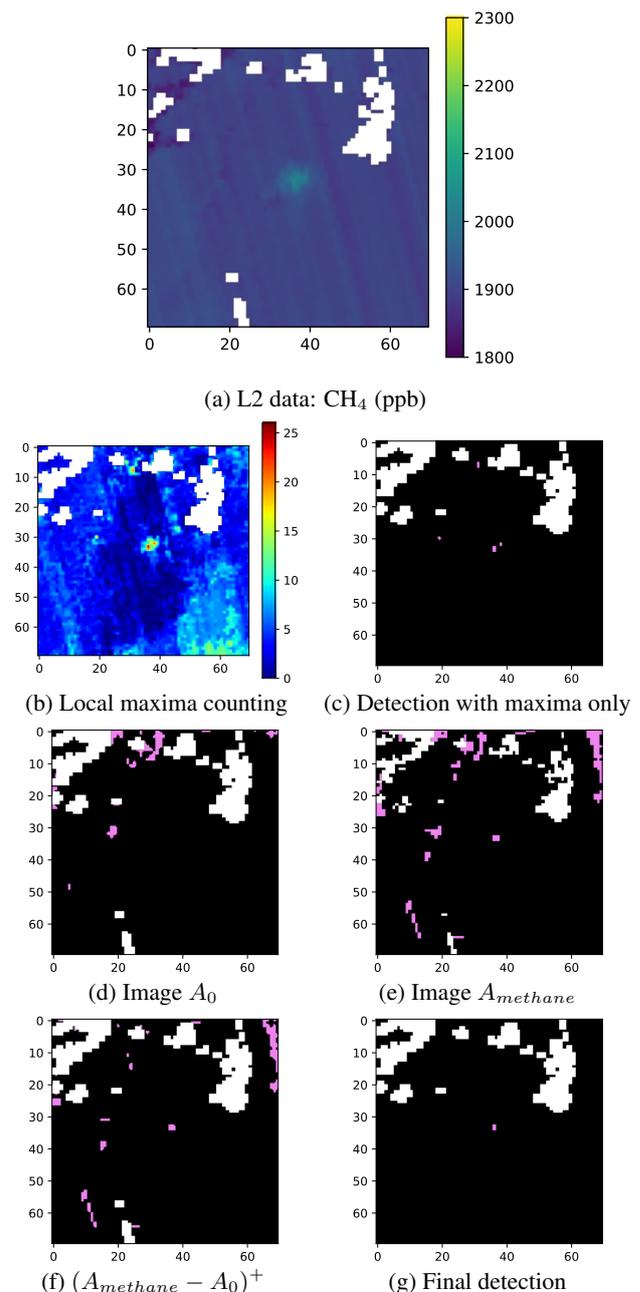(f) $(A_{methane} - A_0)^+$ (g) Final detection

Figure 5. In the image (a), a methane plume is visible in green. Each other image correspond to a part of the detection algorithm, from the counting of maxima to the final detection. Detected pixels are in purple.

junction with another detection technique. As we can see in Figure 5, in the image (f) - which represents $(A_{methane} - A_0)^+$ - many areas are potential plumes but only one of them is actually a methane plume. Even if we use RX with only methane sensitive channels, we detect many anomalies which are not due to an excess of methane. For example, a lack of methane could induce an anomaly.

5. RESULTS AND DISCUSSION

As shown in Figure 1, the results obtained with our method are consistent with the excess in X_{CH4} observed in the L2 product. We detect the plumes seen in L2 data. There is a significant gap

Location	Number of tested plumes	Number of detected plumes
North of Africa	11	8
Middle-east	9	7
Eastearn asia	7	5
Total	27	20

Table 1. Results of the method in several locations.

(a) L2 data

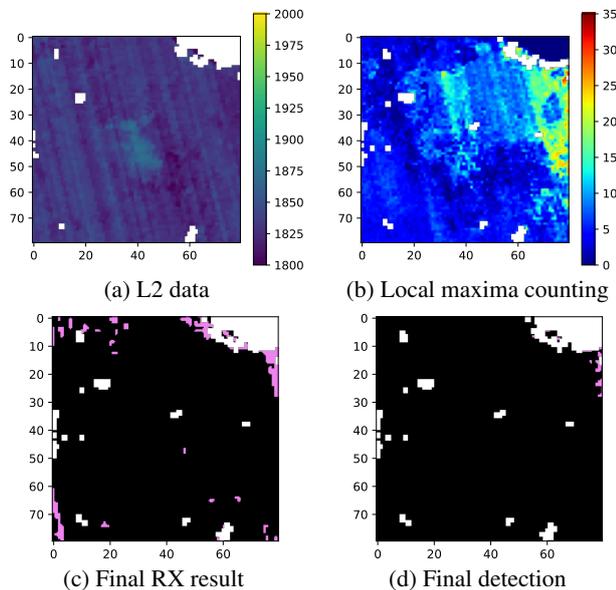(b) Local maxima counting

(c) Final RX result

(d) Final detection

Figure 6. On the top right of (d) we detect a methane plume that cannot be seen in the L2 product. Figures (b) and (c) show that this plume is found with the counting of maxima and with RX.

of the number of maxima between the pixels with excess CH_4 and the others. This allows us to have few false detections.

It should be noted that the image obtained with the local maxima counting is not proportional to the methane concentration. The number of maxima provides an information about the existence of a plume but not about its methane concentration. Thus, this method allow us to detect methane plumes but not to say if a methane plume has a higher methane concentration than another, even if they are in the same image.

Furthermore, as illustrated in Figure 5, the false alarms remaining after the detection based on the spectrum local maxima count are removed by the anomaly detection step. We see in Figure 5 that after the detections based on the spectrum local maxima count, two areas are potential false alarms; but those areas are also detected as anomalies when using the methane-insensitive wavelengths. Therefore, we remove those areas from our methane plume candidates. Finally, we detect some pixels in the center of the image which correspond to the center of the plume seen in the L2 image.

In absence of published benchmarks with ground truth, we validated the proposed method by comparing our results with plumes detected by Kayrros experts on the L2 product. This detection did not involve just L2, but also information on wind, on the potential presence of oil & gas installations, etc. The L2 product by itself does not provide a methane plume mask, it only quantifies the dry column mixing ratio of methane. Local extrema in the L2 product are detected and then the experts checked manually these potential plumes using information mentioned above. As we shall see, only a few (potentially) false alarms remain by

(a) L2 data (08/31/2020)

(b) Local maxima counting

(c) L2 data (03/09/2020)
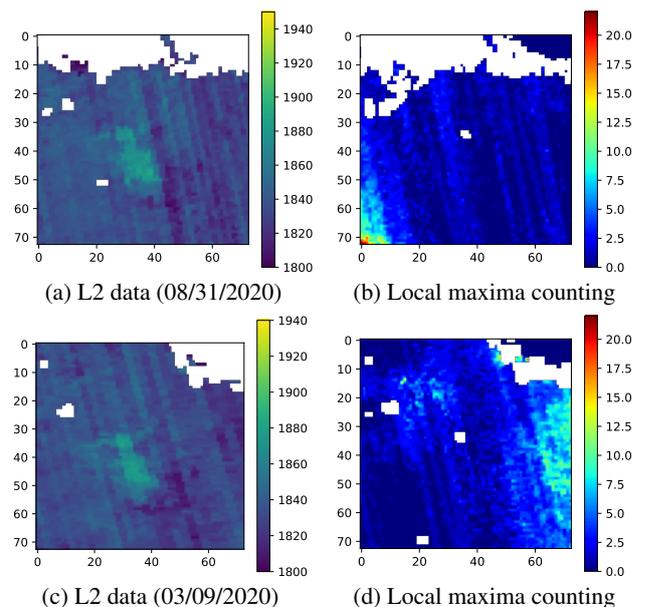
(d) Local maxima counting

Figure 7. A potential methane plume in Australia at two different dates that is discarded by our method.

our method. The plumes detected by experts on the L2 product were considered ground truth true positives. These plumes are taken from several locations, the results are summarized in the Table 1.

Each plume was tested within a 200km × 200km squared area. In total, to test all the plumes, the algorithm was applied to more than 800,000 km². Over those 800,000 km² we found only 7 potential false alarms. Without any proper measurement on the ground, we cannot tell if these detected pixels are indeed false alarms or missed methane plumes. For each tested pixel, the probability of false alarm was set to 10^{-6}.

Conversely, as illustrated in Figure 6, we detected plumes that were not found in the L2 product. Nonetheless, without any proper ground truth we cannot tell if those detections are real. A way to to tell apart false and new detections could be to look at the number of spectrum maxima. If a pixel is detected but with a small number of spectrum maxima (indeed, the detection threshold can be small if the background shows nearly no maxima in most of pixels) we can classify them as low confidence detections.

We compared our method with false alarms typically deduced from the L2 product. In Figure 7 we see such potential false alarms. The plume visible on this figure appears on every observation of this area, independently of the day or weather conditions, and has always the same shape. This plume found in the L2 product is arguably a false alarm. We tested our algorithm on this potential false alarm at several dates and we never detected it as a methane plume. In images (b) and (d) we see the number of counted maxima corresponding respectively to images (a) and (c). Images (b) and (d) show that the pattern recognition step of the algorithm prevents the detection of the shape seen in the L2 product.

In short, our CH_4 detection method detects anomalous pixels backed by a statistical model. It demonstrates the presence of methane in a pixel after background subtraction by a statistical proof of an excess of methane spectrum extrema. The probability of detection is computed as the probability that such a num-

ber of extrema would be observed just by chance in the residual spectrum after background subtraction. The lower this probability the more secure the detection. We saw that fixing an a priori threshold on this probability guarantees a very low number of false alarms, which may help experts focus on meaningful detection. However, the method does not work in presence of clouds or over water where albedo is very small. Moreover, the model used here is a simplified model ignoring the Sentinel-5P instrument noise and the influence of other greenhouse gases like NO_2. Further atmospheric modeling could improve the background subtraction and the final detection. Beside these modeling improvements, two main ideas could be developed. First, more information about weather conditions could be used. In particular, the use of the wind direction could help validate plumes by looking at their shape. Second, instead of detecting the plume pixel by pixel, we could detect groups of pixels. Detection of groups of pixels would lower the probability of false alarms and reinforce the confidence in the detection.

ACKNOWLEDGEMENTS

REFERENCES

Anthropogenic and Natural Radiative Forcing, 2013. Intergovernmental Panel on Climate Change (ed.), *Climate Change 2013 - The Physical Science Basis*, Cambridge University Press, Cambridge, 659–740.

Bradley, E. S., Leifer, I., Roberts, D. A., Dennison, P. E., Washburn, L., 2011. Detection of marine methane emissions with AVIRIS band ratios. *Geophysical Research Letters*, 38(10).

Citterio, M., Paul, F. et al., 2010. Remote sensing of glacier change in West Greenland: Accounting for the occurrence of surge-type glaciers. *Ann Glaciol*, 50, 70-80.

Coakley, J., 2002. Reflectance and albedo. *Encyclopedia of the Atmosphere*, Academic Press, 1914–1923.

Crevoisier, C., Nobileau, D., Fiore, A. M., Armante, R., Chédin, A., Scott, N., 2009. Tropospheric methane in the tropics–first year from IASI hyperspectral infrared observations. *Atmospheric Chemistry and Physics*, 9(17), 6337–6350.

Desolneux, A., Moisan, L., Morel, J., 2003. Maximal Meaningful Events and Applications to Image Analysis. *The Annals of Statistics*, 31(6), 1822–1851. http://www.jstor.org/stable/3448438.

Funk, C., James, T., Roberts, D., Borel, C., 2001. Clustering to Improve Matched Filter Detection of Weak Gas Plumes in Hyperspectral Thermal Imagery. *Geoscience and Remote Sensing, IEEE Transactions on*, 39, 1410 - 1420.

Garain, S., Mitra, D., Das, P., 2019. Detection of hydrocarbon microseepage-induced anomalies by spectral enhancements of Landsat 7 ETM+ images in part of Assam–Arakan Fold Belt, India. *J Pet Explor Prod Technol*, 9.

Gordon, I. E., Rothman, L. S., Hill, C., Kochanov, R. V., Tan, Y. et al., 2017. The HITRAN2016 molecular spectroscopic database. *J Quant Spectrosc Ra*, 203, 3 - 69. https://hal.archives-ouvertes.fr/hal-01765945.

Manolakis, D., Shaw, G., 2002. Detection algorithms for hyperspectral Imaging applications. *Signal Processing Magazine, IEEE*, 19, 29 - 43.

Matteoli, S., Diani, M., Corsini, G., 2010. A tutorial overview of anomaly detection in hyperspectral images. *IEEE Aero El Sys Mag*, 25, 5 - 28.

Montpetit, B., Royer, A. et al., 2012. New shortwave infrared albedo measurements for snow specific surface area retrieval. *Journal of Glaciology*, 58, 941-952.

Pandey, S. et al., 2019. Satellite observations reveal extreme methane leakage from a natural gas well blowout. *PNAS*, 116(52), 26376–26381.

Reitberger, G., Sauer, T., 2019. Background Subtraction using Adaptive Singular Value Decomposition. *CoRR*, abs/1906.12064. http://arxiv.org/abs/1906.12064.

Scafutto, R., Filho, C. D. S., 2018. Detection of Methane Plumes Using Airborne Midwave Infrared Hyperspectral Data. *Remote Sensing*, 10(8). https://www.mdpi.com/2072-4292/10/8/1237.

Thorpe, A., Roberts, D., Bradley, E., Funk, C., Dennison, P., Leifer, I., 2013. High resolution mapping of methane emissions from marine and terrestrial sources using a Cluster-Tuned Matched Filter technique and imaging spectrometry. *Remote Sens Environ*, 134, 305 - 318.

Varon, D. J., Jervis, D., McKeever, J., Spence, I., Gains, D., Jacob, D. J., 2020. High-frequency monitoring of anomalous methane point sources with multispectral Sentinel-2 satellite observations. *Atmospheric Measurement Techniques Discussions*, 2020, 1–21. https://amt.copernicus.org/preprints/amt-2020-477/.