



HAL
open science

Unmixing-based gas plume tracking in LWIR hyperspectral video sequences

Guillaume Tochon, Delphine Pauwels, Mauro Dalla Mura, Jocelyn Chanussot

► **To cite this version:**

Guillaume Tochon, Delphine Pauwels, Mauro Dalla Mura, Jocelyn Chanussot. Unmixing-based gas plume tracking in LWIR hyperspectral video sequences. WHISPERS 2016 - 8th Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing, Aug 2016, Los Angeles, CA, United States. hal-01356161

HAL Id: hal-01356161

<https://hal.science/hal-01356161>

Submitted on 25 Aug 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

UNMIXING-BASED GAS PLUME TRACKING IN LWIR HYPERSPECTRAL VIDEO SEQUENCES

G. Tochon¹, D. Pauwels¹, M. Dalla Mura¹, J. Chanussot^{1,2}

¹GIPSA-lab, Grenoble Institute of Technology, Grenoble, France

²Department of Electrical and Computer Engineering, University of Iceland, Reykjavik, Iceland

ABSTRACT

It is now possible to collect hyperspectral video sequences (HVS) at a near real-time frame rate. The wealth of spectral, spatial and temporal information of those sequences is particularly appealing for chemical gas plume tracking. Existing state-of-the-art methods for such applications however produce only a binary information regarding the position and shape of the gas plume in the HVS. Here, we introduce a novel method relying on spectral unmixing considerations to perform chemical gas plume tracking, which provides information related to the gas plume concentration in addition to its spatial localization. The proposed approach is validated and compared with three state-of-the-art methods on a real HVS.

Index Terms— Spectral unmixing, object tracking, hyperspectral video sequence, chemical gas plume tracking.

1. INTRODUCTION

Thanks to the recent advances in sensor design, it is now possible to acquire hyperspectral video sequences (HVS) at near real-time frame rates. The temporal information additionally provided by such sequences, with respect to traditional hyperspectral imagery, is particularly appealing to monitor both spectrally and spatially transient phenomena. There is however a need to adapt classical hyperspectral processing tools to handle equally the spectral, spatial and temporal information contained in HVS.

Chemical gas plume tracking is a particular application that would benefit from the development of such adapted tools. As a matter of fact, this application is of interest for several real life scenarios, such as the detection of gas leaks in pipelines in order to minimize their impact on the environment [1], or to prevent the use of chemical weapons in a context of defense and security [2]. Most gases are in practice invisible to the human eye and respond only in a restrained portion of the long-wave infrared (LWIR) domain, hence the necessity of a fine sampling of the electromagnetic spectrum. The gas plume to monitor may in addition be very volatile and diffuse in a short lapse of time. Real-time HVS therefore appears as the most suited type of imagery for such tracking applications. State-of-the-art methods devoted to the tracking of gas plumes

in HVS can be divided into three categories. Statistical methods consider the gas plume as anomalous with respect to the background and make use of anomaly detectors, see [3–5] and references therein. Clustering-based methods aim at separating the plume from the background thanks to their spectral dissimilarities [6–8]. Object tracking approaches combine temporal information and object detection methods to iteratively track the plume through the sequence [9, 10]. To the best of our knowledge however, all existing methods retrieve the position and shape of the plume in all frames of the HVS in a binary fashion, and none of them integrate the information related to its concentration. We propose to bridge this gap here by presenting a novel method for gas plume tracking that relies on spectral unmixing considerations, therefore providing this additional knowledge about the plume concentration.

The rest of this paper is organized as follows: section 2 recalls the basics of spectral unmixing and introduces the notations used throughout. Section 3 features the proposed unmixing-based tracking method. Results are presented in section 4 while conclusions and perspectives are drawn in section 5.

2. DEFINITIONS AND NOTATIONS

2.1. Spectral unmixing

Let $\mathbf{X} \in \mathbb{R}^{N_\lambda \times N_p}$ be a hyperspectral data cube composed of N_p pixels and N_λ wavelengths. The goal of spectral unmixing is to factorize \mathbf{X} as the product of an endmember matrix $\mathbf{E} \in \mathbb{R}^{N_\lambda \times d}$ with a fractional abundances matrix $\Phi \in \mathbb{R}^{d \times N_p}$, plus a noise matrix $\mathbf{N} \in \mathbb{R}^{N_\lambda \times N_p}$, with d being the number of endmembers:

$$\mathbf{X} = \mathbf{E}\Phi + \mathbf{N} . \quad (1)$$

Falling in the scope of the linear mixing model [11], each pixel spectrum \mathbf{x} is therefore written as a linear combination of the various endmember spectra \mathbf{e}_i (corresponding to pure macroscopic elements) weighted by their corresponding fractional abundances ϕ_i plus some noise:

$$\mathbf{x} = \sum_{i=1}^d \phi_i \mathbf{e}_i + \boldsymbol{\eta} , \quad (2)$$

with the fractional abundances subjects to positivity ($\phi_i \geq 0 \forall i$) and sum-to-one ($\sum_{i=1}^d \phi_i = 1$) constraints (see [12] for

an overview of spectral unmixing methods). The quality of the unmixing process for a given pixel is frequently evaluated by the root mean square error (RMSE) between the true pixel spectrum \mathbf{x} and $\hat{\mathbf{x}} = \sum_{i=1}^d \phi_i \mathbf{e}_i$, reconstructed using the induced endmembers and abundances:

$$\epsilon(\mathbf{x}, \hat{\mathbf{x}}) = \frac{1}{\sqrt{N_\lambda}} \|\mathbf{x} - \hat{\mathbf{x}}\|_2 \quad (3)$$

2.2. Hyperspectral video sequences

Manipulating HSV implies to handle the temporal dimension in addition to the spectral and spatial ones. For this purpose, \mathbf{X}_t will refer to the hyperspectral image acquired at time t , where the temporal axis has been discretized following the HSV frame-rate: \mathbf{X}_{t+1} will therefore correspond to the frame following \mathbf{X}_t . $\mathbf{E}_t = [\mathbf{e}_1^t, \dots, \mathbf{e}_{d_t}^t]$ will denote the endmember matrix at time t (note that the number of endmembers d_t may vary along the sequence). The same extension will be considered for the fractional abundances matrix Φ_t .

3. PROPOSED UNMIXING-BASED TRACKING METHOD

3.1. Pre-processing

In a context of passive surveillance, it is realistic to assume that the recorded HVS contains plenty of frames featuring only the still background \mathbf{X}_{ref} (up to some minor variations due to sensor noise). The first step of the proposed unmixing-based tracking method is to extract a set of d_{ref} reference endmembers \mathbf{E}_{ref} and fractional abundance Φ_{ref} from \mathbf{X}_{ref} . To overcome the noise influence, we propose in practice to perform this endmember induction and abundances retrieval over a time-averaged version of the still background.

3.2. Detection step

Considering again a passive surveillance scenario, the time where the plume appears in the sequence is likely to be *a priori* unknown. A detection step must therefore be performed prior to the tracking. An input frame \mathbf{X}_t either corresponds to the still background or to the plume appearing over the background in some parts of the frame. In the former case, \mathbf{X}_t should have a low reconstruction error with respect to \mathbf{E}_{ref} and Φ_{ref} , while this reconstruction error should be higher in the area where the plume has appeared in the latter case. Therefore, a change detection procedure is conducted over all pixels composing the reconstruction maps $\epsilon(\mathbf{X}_{t-1}, \mathbf{X}_{\text{ref}})$ and $\epsilon(\mathbf{X}_t, \mathbf{X}_{\text{ref}})$. More specifically, let $\epsilon(\mathbf{x}_{t-1}, \mathbf{x}_{\text{ref}})$ and $\epsilon(\mathbf{x}_t, \mathbf{x}_{\text{ref}})$ be the reconstruction errors of a given pixel spectrum at $t-1$ and t , and let $\Delta\epsilon_{t-1,t}(\mathbf{x}, \mathbf{x}_{\text{ref}}) = \epsilon(\mathbf{x}_t, \mathbf{x}_{\text{ref}}) - \epsilon(\mathbf{x}_{t-1}, \mathbf{x}_{\text{ref}})$ be the difference of such errors. If no change has occurred between $t-1$ and t within pixel \mathbf{x} , this spectrum should have relatively close reconstruction error values, while their difference should be

higher if the plume has contaminated \mathbf{x}_t , since \mathbf{E}_{ref} and Φ_{ref} alone cannot properly reconstruct the plume signature. We can therefore design the following hypothesis test:

$$\begin{aligned} \mathbf{H}_0 : \quad & \Delta\epsilon_{t-1,t}(\mathbf{x}, \mathbf{x}_{\text{ref}}) \sim \mathcal{N}(0, \sigma_{\text{ref}}^2) \\ \mathbf{H}_1 : \quad & \Delta\epsilon_{t-1,t}(\mathbf{x}, \mathbf{x}_{\text{ref}}) \sim \mathcal{N}(\mu \neq 0, \sigma_{\text{ref}}^2) \end{aligned} \quad (4)$$

where the null hypothesis \mathbf{H}_0 corresponds to the absence of change in \mathbf{x} between $t-1$ and t and the alternative hypothesis \mathbf{H}_1 is declared when significant change has occurred. The reference variance σ_{ref}^2 can be reliably estimated using the differences between the reconstruction maps of several background frames. The hypothesis test (4) is solved using the following Z-test statistic [13]:

$$z(\mathbf{x}) = \frac{\overline{\Delta\epsilon_{t-1,t}(\mathcal{V}(\mathbf{x}), \mathcal{V}(\mathbf{x}_{\text{ref}}))}}{\sigma_{\text{ref}} / \sqrt{|\mathcal{V}(\mathbf{x})|}} \underset{\mathbf{H}_0}{\overset{\mathbf{H}_1}{\geq}} \gamma_z \quad (5)$$

where $\overline{\Delta\epsilon_{t-1,t}(\mathcal{V}(\mathbf{x}), \mathcal{V}(\mathbf{x}_{\text{ref}}))}$ denotes the average difference reconstruction error of pixels within the neighborhood $\mathcal{V}(\mathbf{x})$ centered at \mathbf{x} (set in practice to be a 3×3 neighborhood). It is known that $z(\mathbf{x}) \sim \mathcal{N}(0, 1)$ under \mathbf{H}_0 , which allows to set the threshold γ_z to achieve a predefined probability of false alarm p_{FA} . Performing test (5) over all pixels of $\Delta\epsilon_{t-1,t}(\mathbf{x}, \mathbf{x}_{\text{ref}})$ finally leads to the binary change map $C_{t-1,t}$. The plume is declared to have appeared in the sequence if $C_{t-1,t}$ is not empty after removal of single isolated true pixels.

3.3. Tracking step

The tracking step is launched once the gas plume has been detected in the sequence, that is when the binary change mask $C_{t-1,t}$ output by the detection step contains pixels where some change has been flagged. Likewise classical object tracking methods (see [14] for a review), the proposed approach is decomposed in two inner steps addressed sequentially.

3.3.1. Motion prediction step

The goal of the motion prediction step is to estimate an area in the current frame \mathbf{X}_t where the gas plume can be found with a high probability. Here, the change detection procedure previously described in section 3.2 is applied to \mathbf{X}_{t-1} and \mathbf{X}_t , to produce a binary change map $C_{t-1,t}$ featuring areas that have changed between the two consecutive frames. More particularly, $C_{t-1,t}$ describes the regions that have either been left by the plume, or invaded by it between $t-1$ and t . Denoting by O_{t-1} the binary mask of the plume position in \mathbf{X}_{t-1} , it is possible to obtain an estimate position of the plume in \mathbf{X}_t by

$$\hat{O}_t = O_{t-1} \oplus C_{t-1,t} \quad (6)$$

where \oplus stands for the binary XOR operation, following [9]. Note that $\hat{O}_t = C_{t-1,t}$ for the frame where the plume appears, as O_{t-1} is empty in this case. The proposed motion prediction step is illustrated by the workflow presented by figure 1.




Fig. 1: Workflow of the proposed motion prediction step.

3.3.2. Matching step

The goal of the matching step is to retrieve the position of the plume O_t in the current frame \mathbf{X}_t given the estimation \hat{O}_t provided by the motion prediction step. Here, a set of $d_{ref} + 1$ endmembers \mathbf{E}_t is generated using some endmember induction algorithm (EIA), where d_{ref} is the number of reference endmembers used to describe the static background. Ideally, the EIA should generate one endmember associated to the gas plume signature, and the others corresponding to the background. Being in an unsupervised context where the theoretical gas plume signature is unknown, the corresponding endmember \mathbf{e}_p^t is identified as the one that falls in the estimated plume area \hat{O}_t . If several or no endmembers are induced within \hat{O}_t , then the EIA can be run again until exactly one endmembers falls in \hat{O}_t , or \mathbf{e}_p^t can be enforced as the spectrum of the pixel having the highest reconstruction error within \hat{O}_t in $\epsilon(\mathbf{X}_t, \mathbf{X}_{ref})$.

Once the current endmembers \mathbf{E}_t have been generated, their corresponding fractional abundances Φ_t are retrieved. Following, a classification map Θ_t of \mathbf{X}_t is issued using a maximum voting rule on the abundances Φ_t . More specifically, let $\phi_t = [\phi_1^t, \dots, \phi_{d_{ref}+1}^t]^T$ be the vector of $d_{ref} + 1$ abundances associated to a given pixel \mathbf{x}_t , and let $l(\mathbf{x}_t) \in \{1, \dots, d_{ref} + 1\}$ be its associated label. The maximum voting rule

$$l(\mathbf{x}_t) = \underset{i \in \{1, \dots, d_{ref}+1\}}{\operatorname{argmax}} \phi_i^t \quad (7)$$

assigns each pixel \mathbf{x}_t to the class corresponding to the endmember whose fractional abundance is the highest in \mathbf{x}_t .

Following, the classification map Θ_t is used to produce the binary map O_t featuring only the pixels classified in the gas plume class. This binary map is applied over the abundance map associated to the gas plume endmember \mathbf{e}_p^t in order to finally obtain the data of the position, shape, and concentration of the plume O_t in \mathbf{X}_t . This whole matching process is summarized by the workflow in figure 2.

4. EXPERIMENTS

4.1. Data set

The data set used to validated the proposed unmixing-based tracking methodology described in the previous section 3 was provided by John Hopkins Applied Physics Laboratory. The




Fig. 2: Workflow of the proposed matching step.




Fig. 3: (left) False color composition of the 20th frame of the aa12 HVS and (right) corresponding ground truth data.

video sequence was acquired in 2006 at the Dugway Proving Ground in Utah (USA) with a FIRST sensor [15] at a frame-rate of 0.2 Hz. Each frame is a data cube composed of 128×320 pixels in the spatial domain, and 129 spectral channels spanning $7.81 \mu\text{m}$ to $11.97 \mu\text{m}$ in wavelength. The HVS features the explosive release of an acetic acid canister and the resulting sequence is composed of 30 frames, with the gas release occurring at the 11th frame.

Ground truth data was created in order to quantitatively evaluate the performances of the proposed method. A false color composition of each frame was first derived, and two classes (corresponding to strongly concentrated and more diffused areas within the plume) were then carefully delineated. A false color composition of the 20th frame along with its ground truth map are displayed by figure 3.

4.2. Experimental methodology

The critical parameter to operate the proposed unmixing-based tracking methodology is the definition of the EIA, as inconsistent endmembers would not yield proper results. Among the broad literature of EIA [12], we choose to work with the VCA algorithm [16] as it already proved to efficiently unmix LWIR hyperspectral images with faint gas signatures [17]. Fractional abundance are then estimated using the SUNSAL [18] method. Following [4], the number of reference endmembers d_{ref} to described the background is set to 4. It is assumed that the first 5 frames of the sequence are known not to contain the plume. Consequently, the reference variance σ_{ref}^2 is estimated over the first 4 differences of reconstruction error maps. The Z-test statistic (5) is operated under a probability of false alarm set to $p_{FA} = 0.01$.

Due to the stochasticity of the VCA, the proposed method is run 20 times, and the average percentages of correct detections (amount of ground truth pixels correctly detected) and false detections (amount of ground truth pixels wrongly detected)




Fig. 4: Top row: frames #12, 16, 20, 24 and 28 of the sequence. Bottom row: corresponding detection \mathcal{O}_t .

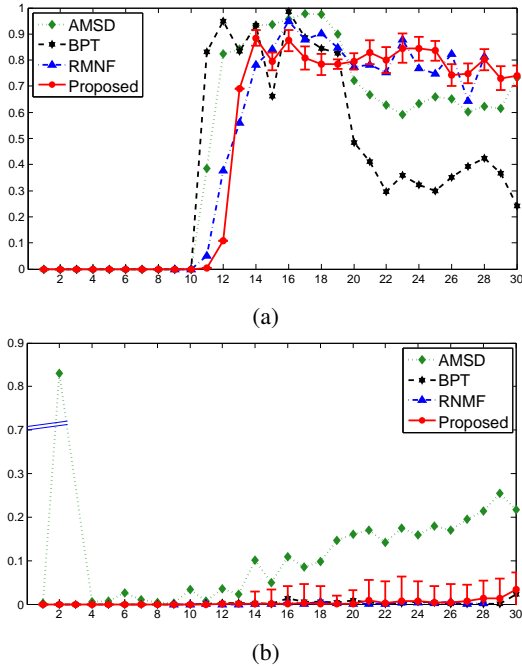


Fig. 5: Percentage of (a) correct and (b) false detections with respect to the ground truth.

with respect to the ground truth are computed for all frames of the sequence.

4.3. Results

Figure 5 shows the average percentage of correct and false detections for the proposed unmixing-based method (with standard deviation), as well as for three state-of-the-art approaches, namely the AMSD [4] (statistical-based), RMNF [6] (clustering-based) and BPT [10] (object detection based) methods. As it can be seen in figure 5a, the proposed approach performs equally well than the RMNF approach, and outperforms both AMSD and BPT-based methods for most frames of the sequence. While the release instant of the plume is correctly identified, the proposed method fails at properly detecting the gas plume in the early frames #11 and #12. The reason is that the EIA is not able to automatically induce an endmember in the estimated area $\hat{\mathcal{O}}_t$, which is then forced to

be the pixel with the highest reconstruction error. Due to the explosive release of the canister containing the gas, the enforced plume endmember is spectrally aberrant and unable to correctly describe the gas plume, hence the failing of the subsequent classification step. Nonetheless, the average percentage of correct detections for the other frames of the sequence remains around 80% while the AMSD and BPT methods suffer a drop as the diffusion is becoming more severe. It is also worth adding that the RMNF method processes all frames of the sequence in a stacked manner, therefore being incompatible with a real-time implementation scenario where incoming frames would have to be processed on the fly.

Figure 5b shows the percentage of false detections, corresponding to the pixels whose spectra are supposedly not corrupted by the plume but which are nonetheless detected by the various methods. The y-axis has been cut for a better visualization. Apart for the AMSD approach that produces a significant amount of false detections (notably before the release of the plume), all three other competing approaches correctly identify the release instant, and their number of false detections remain relatively low across the sequence.

The previous observations are confirmed by the qualitative analysis of figure 4. While the proposed approach is not able to properly detect the plume for the early frame #12, the retrieved shape and concentration appear visually very consistent for all the other frames of the sequence.

5. CONCLUSION

We presented here a novel method to perform the detection and tracking of a diffusing gas plume in HVS. Unlike classical state-of-the-art methods, the proposed approach is not only able to retrieve the position and shape of the plume in all frames of the sequence, but also provides information related to its concentration thanks to the integration of a spectral unmixing process within the tracking procedure. The proposed method was investigated over a real HVS and compared with three state-of-the-art approaches, and showed to perform equally or better while providing a richer information. Extending the proposed approach to accommodate non linear spectral mixing models is part of our research perspectives, as well as its optimized implementation to make it operable under real-time processing constraints.

ACKNOWLEDGMENTS

We thank Andrea Bertozzi, Department of Mathematics, University of California, Los Angeles (UCLA), for her collaboration on the dataset provided by the US Defense Threat Reduction Agency and the National Science Foundation through NSF grant DMS- 1118971. This work was also partially funded through DGA, under grant 2015-60-001200.470.75.01.

6. REFERENCES

- [1] M. Hinnrichs, “Imaging spectrometer for fugitive gas leak detection,” in *Photonics East’99*. International Society for Optics and Photonics, 1999, pp. 152–161.
- [2] V. Farley, A. Vallières, A. Villemaire, M. Chamberland, P. Lagueux, and J. Giroux, “Chemical agent detection and identification with a hyperspectral imaging infrared sensor,” in *Optics/Photonics in Security and Defence*. International Society for Optics and Photonics, 2007, pp. 673918–673918.
- [3] D. Manolakis, S. Golowich, and R. DiPietro, “Long-Wave Infrared Hyperspectral Remote Sensing of Chemical Clouds: A focus on signal processing approaches,” *Signal Processing Magazine, IEEE*, vol. 31, no. 4, pp. 120–141, 2014.
- [4] J. B. Broadwater, D. Limsui, and A. K. Carr, “A Primer for Chemical Plume Detection Using LWIR Sensors,” Tech. Rep., John Hopkins Applied Physics Laboratory, 04 2011.
- [5] E. M. O’Donnell, D. W. Messinger, C. Salvaggio, and J. R. Schott, “Identification and detection of gaseous effluents from hyperspectral imagery using invariant algorithms,” in *Defense and Security*. International Society for Optics and Photonics, 2004, pp. 573–582.
- [6] E. Merkurjev, J. Sunu, and A. L. Bertozzi, “Graph MBO method for multiclass segmentation of hyperspectral stand-off detection video,” in *Image Processing (ICIP), 2014 IEEE International Conference on*. IEEE, 2014, pp. 689–693.
- [7] J. Qin, T. Laurent, K. Bui, R. Tan, J. Dahilig, S. Wang, J. Rohe, J. Sunu, and A. L. Bertozzi, “Detecting plumes in LWIR using robust nonnegative matrix factorization with graph-based initialization,” in *SPIE Defense+ Security*. International Society for Optics and Photonics, 2015, pp. 94720V–94720V.
- [8] H. Hu, J. Sunu, and A. L. Bertozzi, “Multi-class graph mumford-shah model for plume detection using the mbo scheme,” in *Energy Minimization Methods in Computer Vision and Pattern Recognition*. Springer, 2015, pp. 209–222.
- [9] G. Tochon, J. Chanussot, J. Gilles, M. Dalla Mura, J-M. Chang, and A. L. Bertozzi, “Gas plume detection and tracking in hyperspectral video sequences using binary partition trees,” in *IEEE Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS 2014)*, 2014.
- [10] G. Tochon, J. Chanussot, M. Dalla Mura, and A. L. Bertozzi, “Hierarchical representation of hyperspectral video sequences: application to chemical gas plume tracking,” *Pattern recognition*, (submitted).
- [11] N. Keshava and J. F. Mustard, “Spectral unmixing,” *Signal Processing Magazine, IEEE*, vol. 19, no. 1, pp. 44–57, 2002.
- [12] J. M. Bioucas-Dias, A. Plaza, N. Dobigeon, M. Parente, Qian D., P. Gader, and J. Chanussot, “Hyperspectral Unmixing Overview: Geometrical, Statistical, and Sparse Regression-Based Approaches,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 5, no. 2, pp. 354–379, 2012.
- [13] L. L. Scharf, *Statistical signal processing*, vol. 98, Addison-Wesley Reading, MA, 1991.
- [14] A. Yilmaz, O. Javed, and M. Shah, “Object tracking: A survey,” *ACM computing surveys (CSUR)*, vol. 38, no. 4, pp. 13, 2006.
- [15] V. Farley, A. Vallières, M. Chamberland, A. Villemaire, and J-F. Legault, “Performance of the first: a long-wave infrared hyperspectral imaging sensor,” in *Optics/Photonics in Security and Defence*. International Society for Optics and Photonics, 2006, pp. 63980T–63980T.
- [16] J. M-P. Nascimento and J. M. Bioucas-Dias, “Vertex component analysis: A fast algorithm to unmix hyperspectral data,” *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 43, no. 4, pp. 898–910, 2005.
- [17] S. Henrot, J. Chanussot, and C. Jutten, “Dynamical spectral unmixing of multitemporal hyperspectral images,” *arXiv preprint arXiv:1510.04238*, 2015.
- [18] M-D. Iordache, J. M. Bioucas-Dias, and A. Plaza, “Total variation spatial regularization for sparse hyperspectral unmixing,” *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 50, no. 11, pp. 4484–4502, 2012.