

---

# Hyperspectral shadow removal with Iterative Logistic Regression and latent Parametric Linear Combination of Gaussians

---

**Core Francisco Park**  
Department of Physics  
Harvard University  
Cambridge, MA 02138, U.S.A.  
corefranciscopark@g.harvard.edu

**Maya Nasr**  
John A. Paulson School of  
Engineering and Applied Sciences  
Harvard University  
Cambridge, MA 02138, U.S.A.  
mayanasr@g.harvard.edu

**Manuel Pérez-Carrasco**  
Data Science Unit  
University of Concepcion  
Biobio, Chile  
maperezc@inf.udec.cl

**Eleanor Walker**  
John A. Paulson School of  
Engineering and Applied Sciences  
Harvard University  
Cambridge, MA 02138, U.S.A.  
ewalker@g.harvard.edu

**Douglas Finkbeiner**  
Department of Physics  
Harvard University  
Cambridge, MA 02138, U.S.A.  
dfinkbeiner@cfa.harvard.edu

**Cecilia Garraffo**  
AstroAI  
Center for Astrophysics | Harvard & Smithsonian  
Cambridge, MA 02138, U.S.A.  
cgarraffo@cfa.harvard.edu

## Abstract

Shadow detection and removal is a challenging problem in the analysis of hyperspectral images. Yet, this step is crucial for analyzing data for remote sensing applications like methane detection. In this work, we develop a shadow detection and removal method only based on the spectrum of each pixel and the overall distribution of spectral values. We first introduce Iterative Logistic Regression (ILR) to learn a spectral basis in which shadows can be linearly classified. We then model the joint distribution of the mean radiance and the projection coefficients of the spectra onto the above basis as a parametric linear combination of Gaussians. We can then extract the maximum likelihood mixing parameter of the Gaussians to estimate the shadow coverage and to correct the shadowed spectra. Our correction scheme reduces correction artefacts at shadow borders. The shadow detection and removal method is applied to hyperspectral images from MethaneAIR, a precursor to the satellite MethaneSAT.

## 1 Introduction

Methane is the second most important greenhouse gas following carbon dioxide. MethaneSAT is a satellite mission aiming to map the methane density in the atmosphere[1]. For this and other similar missions, careful analysis of the hyperspectral image resulting from the detector is essential to retrieve the target atmospheric components. Two significant challenges in image analysis are clouds and shadows[2]. In the case of MethaneSAT, a cloud essentially blocks the retrieval of methane density




Figure 1: Hyperspectral data from MethaneAIR. (Left) The mean radiance for all pixels. (Middle) The spectrum at each of the 3 points on the left image. (Right) The map of labels. “Sure Ground” and “Sure Cloud” are simple made by binary erosion of the shadow and ground labels.

under it. However, depending on the indirect ground illumination, shadowed pixels might carry a detectable signal. These pixels could either be processed by a different downstream pipeline or be corrected for the effect of the shadow before entering the subsequent pipelines. In this work, we present a probabilistic latent spectral space inference method to remove the average effect of a shadow on a spectrum, while preserving the original spectral features as much as possible. This method is applied to hyperspectral images from MethaneAIR, a precursor to the MethaneSAT satellite.

## 2 Method

**Data** We use hyperspectral imaging data from a MethaneAIR flight [2]. Our data is a set of hyperspectral images with 1024 spectral bins and varying width and height. An image with clouds and shadows within the field of view is in Fig. 1. The most evident feature of a shadow is the reduced mean radiance. To enhance the method to learn spectral shape changes by shadows, we explicitly remove the mean radiance ( $m_i$ ) by dividing each spectrum by its mean and taking the logarithm. Thus each spectrum  $f_i$  is normalized as  $s_i = \log(f_i/m_i)$ .  $\log$  is the natural logarithm everywhere.

**Iterative Logistic Regression and Latent Basis Learning** Using spectra and binary labels ( $s_i \in \mathcal{R}^{1024}$ ,  $y_i \in [0, 1]$ ) from sunlit and shadowed pixels, we train our Iterative Logistic Regression (ILR), described in Algorithm 1. On each iteration, the best spectral weight separating shadowed pixels and bright pixels are learned by a logistic regression and projected out from the data, until the classification accuracy becomes negligible. The resulting weights ( $\mathbf{W} \in \mathcal{R}^{(k,1024)}$ ) form a  $k$ -dimensional spectral basis for the latent space and the projection coefficients of each spectrum onto this basis ( $\beta_i \in \mathcal{R}^k$  s.t.  $s_i = \beta_i^T \mathbf{W} + s_i^{null}$ ) becomes the features of each spectrum in this basis. In practice the dimensionality of this space is much smaller than the number of spectral bins, in our case  $k = 23$ .

---

### Algorithm 1 Iterative Logistic Regression

---

**Require:**  $x \in \mathcal{R}^{(n,p)}$ ,  $y \in [0, 1]^n$ ,  $f_{thres} \in (0, 1)$  ▷ Dataset with  $n$  rows and  $p$  features  
 $W = []$ ;  $f_1 = 1$   
**while**  $f_1 \geq f_{thres}$  **do** ▷ Repeat until the data is hard to classify  
     $x_{Tr}, y_{Tr}, x_{Te}, y_{Te} = split(x, y)$   
     $w = get\_wLR(x_{Tr}, y_{Tr})$  ▷ Get the separating hyperplane from logistic regression  
     $\hat{y}_{Te} = \sigma(x_{Te} \cdot w) > 0.5$ ;  $f_1 = get\_f_1(y_{Te}, \hat{y}_{Te})$  ▷ Get the predictions,  $\sigma$  is a sigmoid  
     $x \leftarrow x - x \cdot w / \|w\|_2$  ▷ Project out this component  
     $W.append(w)$   
**end while**  
**return** matrix( $W$ )

---

**Joint estimation of shadow fraction and mean radiance for border pixels.** Now, we construct a new feature space for each pixel by combining the log average radiance ( $\log(m_i)$ ) and  $\beta_i$ . This is our  $(k+1)$  dimensional latent spectrum,  $e_i \in \mathcal{R}^{k+1}$ , with  $e_{i,0} = \log(m_i)$ . We approximate the joint distribution for the sunlit latent spectrum (subscript  $g$  for ground) and the shadowed latent spectrum

(subscript s for shadow) as Gaussian distributions. Our model then assumes each latent spectrum on the border is a draw from a linear combination of these Gaussians. It is important to note that our model is not a Gaussian Mixture Model(GMM) but simply a Parametric Linear Combination of Gaussians(PLCG).

$$e_i^b \sim (1 - \alpha_i)\mathcal{N}_g(\boldsymbol{\mu}_g, \boldsymbol{\Sigma}_g) + \alpha_i\mathcal{N}_s(\boldsymbol{\mu}_s, \boldsymbol{\Sigma}_s) \quad (0 \leq \alpha \leq 1) \quad (1)$$

where  $e_i^b$  is a border latent spectrum,  $\alpha$  is the mixing parameter and  $\boldsymbol{\mu}_g, \boldsymbol{\Sigma}_g, \boldsymbol{\mu}_s, \boldsymbol{\Sigma}_s$  are the means and covariances for the ground and shadowed latent spectrum, respectively. The superscript  $i$  indicates that each border latent spectrum is assumed to be drawn from a distribution resulting from a different mixing parameter. When  $\alpha = 0$  the border latent spectrum is just a draw from  $\mathcal{N}_g$  and for  $\alpha = 1$  it is a draw from  $\mathcal{N}_s$ , thus  $\alpha$  is a parameter representing how much a spectrum is covered by a shadow. We can then estimate  $\alpha$  independently for each border latent spectrum using maximum likelihood estimation (MLE) under a Gaussian likelihood<sup>1</sup>.

$$\alpha_i^* = \operatorname{argmax}_{\alpha_i} \mathcal{L}(\alpha_i | e_i^b) \quad (2)$$

$$\log \mathcal{L}(\alpha_i | e_i^b) = -\frac{1}{2} \log(|\boldsymbol{\Sigma}(\alpha)|) - \frac{1}{2} (e_i^b - \boldsymbol{\mu}(\alpha))^T \boldsymbol{\Sigma}^{-1}(\alpha) (e_i^b - \boldsymbol{\mu}(\alpha)) + C \quad (3)$$

$$\boldsymbol{\mu}(\alpha) = (1 - \alpha)\boldsymbol{\mu}_g + \alpha\boldsymbol{\mu}_s \quad \boldsymbol{\Sigma}(\alpha) = (1 - \alpha)\boldsymbol{\Sigma}_g + \alpha\boldsymbol{\Sigma}_s \quad (4)$$

Since the latent space’s dimension is small( $\mathcal{O}(10)$ ), and  $\alpha$  is a single parameter, a simple grid search can be performed to find the (MLE)  $\alpha$ . Then each corrected border spectrum (and its mean) can be obtained as:

$$e_i^{rb} = \mathbf{S}_i(e_i^b - \boldsymbol{\mu}(\alpha_i^*)) + \boldsymbol{\mu}_g \quad (5)$$

$$s_i^{rb} = (e_{i,1:k}^{rb})^T W + s_i^{null} \quad \text{and} \quad \log(\bar{s}_i) = e_{i,0}^{rb} \quad (6)$$

where  $\mathbf{S}$  is an identity matrix everywhere except  $S_{00} = \sqrt{\Sigma_{g00}/\Sigma(\alpha_i^*)_{00}}$ . This  $S_{00}$  term scales  $e_{i,0}$  by the ratio of standard deviation of  $\log(m_i)$  of ground spectra to shadow spectra, so that the whole region under a shadow doesn’t end up with a lower contrast. This is motivated by the fact that we only want to correct for the average effect of the shadows while keeping the spectral variances as they are.

### 3 Results

**ILR and shadow basis learning** Fig. 2 shows ILR applied to MethaneAIR data with ground and shadow labels. After  $k = 23$  iterations, the F1-score is negligible and we have extracted the basis vectors relevant for linear shadow classification. From the joint distribution of the mean radiance and the shadow coefficients of every pixel, we find that the “sure ground” and “sure shadow” can both be described as a  $(k+1)$ -dimensional Gaussian distribution, while the border pixels are scattered between them.

**Shadow Correction** Fig. 3 illustrates our correction method and its results. We find that the posterior cloud fraction is, as expected, a tight function of the mean radiance. We find that our reconstruction scheme is better than adjusting the mean radiance by a factor determined from the raw probabilities of logistic regression. This is illustrated in the bottom center panel of Fig. 3, where the cloud probability is discontinuous and does contain information about the shift in the mean radiance. A mean correction thus results in an over-correction right inside the shadow and a under-correction right outside the shadow. Our method mitigates this effect.

### 4 Discussion

**Objective of shadow removal** Our method casts shadow removal into a slightly different problem than in the literature. Previous works [3, 4, 5], focus on retrieving the radiance and spectral shape which would have been measured if the shadow was not present. Here, we focus on removing the average effect of having a shadow cast on a pixel, while conserving the original spectrum as much as possible.

<sup>1</sup>This assumes the two distributions are jointly Gaussian distributed




Figure 2: Iterative Logistic Regression and the shadow basis (Left) F1-score as components relevant to shadow classification are iteratively removed. (Middle) The learned spectral weights( $\mathbf{W}$ ) forming the shadow basis. (Right) The joint distribution of the log-mean radiance and the first shadow coefficient( $(m_i, e_{i,1}^b)$ ). We observe that the sunlit spectra is roughly  $e^{1.5} \sim 4.5$  times brighter than the shadowed spectra.




Figure 3: Our shadow correction method. (Top Left) The inferred posterior cloud fraction as a function of mean radiance. (Top Center) Our correction method illustrated. The olive colored points are samples from the inferred mixture of Gaussians. (Top Right) The distribution shift in latent space from our correction method. (Bottom Left) the mean based correction and our correction (Bottom Center) Our correction method across a shadow border(red line). (Bottom Right) Spectral correction from our method. The spectra are corrected from the magenta point and is compared with the reference spectrum from the blue point.

**Impact towards tackling climate change** Our approach increases the raw amount of area from which methane emission can be detected[2]. MethaneAIR has two spectrometers targeting 1236-1319 nm for  $O_2$  and 1592-1697 nm for  $CO_2$  and  $CH_4$  (1249-1305 nm and 1598-1683 nm respectively for MethaneSAT). MethaneAIR detects the total column-averaged dry-air mole fractions of  $CH_4$  ( $XCH_4$ ) by the  $CO_2$  proxy method [6] using the second (longer wavelength) spectrometer. This requires accurate spectra at the P and R bands of  $CO_2$  (1595–1618 nm) and the  $2\nu_3$  band of  $CH_4$  (1629–1654 nm) [7]. These bands correspond roughly to spectral channel indices 50-250 and 350-550 in the x-axis of Fig. 3, bottom right. We find that our correction method is significantly different from

the mean based correction in these regions and thus demonstrates that it might affect the methane detection significantly. We plan to study this more carefully in a future work.

**Limitations and Future Work** A major limitation of the presented model is that it is agnostic about any spatial priors, and this limits the estimation of the shadow coverage at the borders. This is the main reason our shadow removal, although better than merely correcting the mean of the distribution, has spatial artifacts at the shadow borders. We believe the introduction of spatial priors added to the likelihood function in Eq. 3 will improve our method and further reduce over/under-corrections.

Our method doesn't make use of any assumptions specific to shadows, so the same model can be used to correct for other continuous artifacts.

## 5 Conclusion

We have developed a method to detect and remove the average effect of shadows from a hyperspectral image using only binary shadow and ground labels. We plan to enhance the model using spatial priors and apply it to MethaneAIR and MethaneSAT images for additional detection of methane in shadowed areas.

## Acknowledgments and Disclosure of Funding

CFP gratefully acknowledges Dr. Steven Wofsy and Christopher Chan Miller for providing the data used in this work and for useful discussions. This work is supported by the Environmental Defense Fund and by MethaneSAT.org, a wholly-owned subsidiary of the Environmental Defense Fund. The computations in this paper were run on the FASRC Cannon cluster supported by the FAS Division of Science Research Computing Group at Harvard University. This collaboration was facilitated by EarthAI and AstroAI.

## References

- [1] Reuben R Rohrschneider, Steven Wofsy, Jonathan E Franklin, Joshua Benmergui, Juancarlos Soto, and Spencer B Davis. The methanesat mission. 2021.
- [2] Apisada Chulakadabba, Maryann Sargent, Thomas Lauvaux, Joshua S Benmergui, Jonathan E Franklin, Christopher Chan Miller, Jonas S Wilzewski, Sébastien Roche, Eamon Conway, Amir H Sour, et al. Methane point source quantification using methaneair: A new airborne imaging spectrometer. *EGUsphere*, 2023:1–22, 2023.
- [3] Guichen Zhang, Daniele Cerra, and Rupert Müller. Shadow detection and restoration for hyperspectral images based on nonlinear spectral unmixing. *Remote Sensing*, 12(23):3985, 2020.
- [4] Min Zhao, Jie Chen, and Susanto Rahardja. Hyperspectral shadow removal via nonlinear unmixing. *IEEE Geoscience and Remote Sensing Letters*, 18(5):881–885, 2021.
- [5] Puhong Duan, Shangsong Hu, Xudong Kang, and Shutao Li. Shadow removal of hyperspectral remote sensing images with multiexposure fusion. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–11, 2022.
- [6] Christian Frankenberg, U Platt, and T Wagner. Iterative maximum a posteriori (imap)-doas for retrieval of strongly absorbing trace gases: Model studies for ch 4 and co 2 retrieval from near infrared spectra of sciamachy onboard envisat. *Atmospheric Chemistry and Physics*, 5(1):9–22, 2005.
- [7] Christopher Chan Miller, Sebastien Roche, Jonas S Wilzewski, Xiong Liu, Kelly Chance, Amir H Sour, Eamon Conway, Bingkun Luo, Jenna Samra, Jacob Hawthorne, et al. Methane retrieval from methaneair using the co 2 proxy approach: A demonstration for the upcoming methanesat mission. *EGUsphere*, 2023:1–40, 2023.