

ACTIVE DIFFUSION AND VCA-ASSISTED IMAGE SEGMENTATION OF HYPERSPECTRAL IMAGES

Sam L. Polk ¹, Kangning Cui ², Robert J. Plemmons ³, and James M. Murphy ^{1*}

¹ Department of Mathematics, Tufts University

² Department of Mathematics, City University of Hong Kong

³ Departments of Mathematics and Computer Science, Wake Forest University

ABSTRACT

Hyperspectral images encode rich structure that can be exploited for material discrimination by machine learning algorithms. This article introduces the Active Diffusion and VCA-Assisted Image Segmentation (ADVIS) for active material discrimination. ADVIS selects high-purity, high-density pixels that are far in diffusion distance (a data-dependent metric) from other high-purity, high-density pixels in the hyperspectral image. The ground truth labels of these pixels are queried and propagated to the rest of the image. The ADVIS active learning algorithm is shown to strongly outperform its fully unsupervised clustering algorithm counterpart, suggesting that the incorporation of a very small number of carefully selected ground truth labels can result in substantially superior material discrimination in hyperspectral images.

Index Terms— Active Learning, Diffusion Geometry, Hyperspectral Imagery, Image Segmentation, Semi-supervised Machine Learning.

1. INTRODUCTION

Hyperspectral images (HSIs) are high-dimensional remotely-sensed images that encode rich information about a scene [1]. Despite storing reflectance in a hundred or more spectral bands, HSIs typically encode intrinsically low-dimensional structure that can be exploited by machine learning algorithms for image segmentation [2, 3, 4]. While HSIs are an important data source for material discrimination, their use for this task is complicated by at least two key factors. First, the number of pixels in typical HSIs can be very large, rendering manual labeling and analysis infeasible [1]. Second, the spatial resolution of HSIs is often coarse, so any one pixel may correspond to a spatial region that contains multiple materials [2]. Thus, efficient machine learning algorithms that rely on few expert annotations or labels are needed to capture the latent material structure in HSIs.

*Corresponding Author; email: JM.Murphy@Tufts.edu
This work was partially funded by the US National Science Foundation grants NSF-DMS 1912737, NSF-DMS 1924513, and NSF-CCF 1934553.

This article introduces the Active Diffusion and VCA-Assisted Image Segmentation (ADVIS) algorithm for material discrimination in HSIs using Vertex Component Analysis (VCA) [5]. ADVIS is an *active learning* algorithm based on the unsupervised Diffusion and VCA-Assisted Image Segmentation (D-VIS) clustering algorithm. D-VIS is closely related to Diffusion and Volume maximization-based Image Clustering (D-VIC), which has been shown to perform well at material discrimination on benchmark HSIs [2]. We show that incorporating just a few carefully-chosen expert labels in ADVIS substantially improves algorithm performance.

The rest of this article is structured as follows. In Section 2, background is provided on HSI segmentation, diffusion geometry, spectral unmixing, and D-VIS. In Section 3, the ADVIS algorithm for active material discrimination is introduced. Section 4 contains numerical experiments where ADVIS is compared against D-VIS on real HSI data. In Section 5, we conclude and discuss future work.

2. BACKGROUND

2.1. Hyperspectral Image Segmentation

An HSI segmentation algorithm partitions pixels of an HSI $X = \{x_i\}_{i=1}^n \subset \mathbb{R}^D$ (interpreted as a point cloud, where n is the number of pixels and D is the number of spectral bands) into groups $\{X_k\}_{k=1}^K$ sharing key commonalities (e.g., common materials) [6]. Unsupervised HSI segmentation (also called *clustering*) algorithms do not rely on ground truth labels to obtain the partition $\{X_k\}_{k=1}^K$. In contrast, *semi-supervised* and *active learning* HSI segmentation algorithms rely on the ground truth labels of a few pixels to partition X .

2.2. Diffusion Geometry

To exploit the intrinsic low-dimensionality of HSIs, graph-based HSI segmentation algorithms identify HSI pixels as nodes in an undirected graph [7]. Edges between pixels are encoded in an adjacency matrix $\mathbf{W} \in \mathbb{R}^{n \times n}$, where $\mathbf{W}_{ij} = 1$ if x_j is one of the N nearest neighbors of x_i and $\mathbf{W}_{ij} = 0$

otherwise. Define $\mathbf{P} = \mathbf{D}^{-1}\mathbf{W}$, where \mathbf{D} is the diagonal matrix with $\mathbf{D}_{ii} = \sum_{j=1}^n \mathbf{W}_{ij}$. The matrix \mathbf{P} can be identified as the transition matrix for a Markov diffusion process on HSI pixels. Assuming \mathbf{P} is reversible, aperiodic, and irreducible, there is a unique $\pi \in \mathbb{R}^{1 \times n}$ satisfying $\pi\mathbf{P} = \pi$.

Diffusion distances enable direct comparisons between pixels in the context of the diffusion process encoded in \mathbf{P} [7]. Define the diffusion distance at time $t \geq 0$ between pixels $x_i, x_j \in X$ by $D_t(x_i, x_j) = \sqrt{\sum_{k=1}^n [(\mathbf{P}^t)_{ik} - (\mathbf{P}^t)_{jk}]^2 / \pi_k}$. For datasets with well-separated and highly coherent classes, the within-class diffusion distance is bounded away from the between-class diffusion distance across a broad range of t [8]. Thus, diffusion distances are a useful tool for HSI segmentation. Diffusion distances can be related to the eigendecomposition of \mathbf{P} via $D_t(x_i, x_j) = \sqrt{\sum_{k=1}^n |\lambda_k|^{2t} [(\psi_k)_i - (\psi_k)_j]^2}$, where $\{(\lambda_k, \psi_k)\}_{k=1}^n$ are the right eigenvalue-eigenvector pairs of \mathbf{P} [7]. For t sufficiently large, eigenvectors with $|\lambda_k|^t \approx 0$ can be discarded, yielding a low-cost, accurate approximation of diffusion distances.

2.3. Spectral Unmixing

HSIs are often recorded at a coarse spatial resolution, so a single pixel may correspond to a spatial region that contains multiple materials [5, 9, 10]. *Spectral unmixing* algorithms may be used to estimate the proportions of materials within each pixel [5, 10]. Mathematically, if m is the number of materials in the scene, linear spectral unmixing algorithms learn two matrices, $\mathbf{A} \in \mathbb{R}^{n \times m}$ (called abundances) and $\mathbf{U} = (u_1 \ u_2 \ \dots \ u_m)^\top \in \mathbb{R}^{m \times D}$ (called endmembers) such that $x_i \approx \sum_{j=1}^m \mathbf{A}_{ij} u_j$ for each $x_i \in X$. Each u_i is the intrinsic spectral signature of a material, and the rows of \mathbf{A} encode the relative abundances of materials in HSI pixels [5, 10]. The *purity* of the pixel x_i , defined by $\eta(x_i) = \max_{1 \leq j \leq m} \mathbf{A}_{ij}$, will thus be large for pixels that predominantly contain just one material and small elsewhere [2].

2.4. Diffusion and VCA-Assisted Image Segmentation

D-VIS (Algorithm 1) is an unsupervised material discrimination algorithm, meaning no expert labels are used to obtain an HSI segmentation. D-VIS operates in two main stages. In the first, D-VIS learns an estimate for pixels that are exemplary of all underlying material classes (called *class modes*) and assigns these pixels unique labels. D-VIS propagates the labels of class modes to non-modal pixels in its second stage.

D-VIS first performs spectral unmixing of X to calculate $\eta(x)$, using HySime to learn m [9] and VCA to learn end-members [5, 11]. This differs slightly from D-VIC, which relies on Alternating Volume Maximization to learn end-members [2, 10]. Next, D-VIS calculates empirical density: $p(x) = \sum_{y \in NN_N(x)} \exp(-\|x-y\|_2^2 / \sigma_0^2)$, where $NN_N(x)$ is the set of N nearest neighbors of x in X and $\sigma_0 > 0$ is a *density scale* that controls the interaction radius between pixels.

Algorithm 1: Diffusion and VCA-Assisted Image Segmentation (D-VIS)

Input: X (HSI), N (# neighbors), K (# classes)
 σ_0 (density scale), t (diffusion time),
Output: $\hat{\mathcal{C}}$ (HSI segmentation)

- 1 Compute $\eta(x)$, using HySime [9] to estimate m and VCA [5, 11] for spectral unmixing;
- 2 For each $x \in X$, compute $\zeta(x) = \frac{2\bar{p}(x)\bar{\eta}(x)}{\bar{p}(x)+\bar{\eta}(x)}$;
- 3 Sort X according to $\mathcal{D}_t(x) = \zeta(x)d_t(x)$ in non-increasing order. Denote this sorting $\{x_{m_k}\}_{k=1}^n$. Label $\hat{\mathcal{C}}(x_{m_k}) = k$ for $1 \leq k \leq K$;
- 4 In order of non-increasing $\zeta(x)$, for each unlabeled $x \in X$, assign the label $\hat{\mathcal{C}}(x) = \hat{\mathcal{C}}(x^*)$, where $x^* = \operatorname{argmin}_{y \in X} \{D_t(x, y) | \zeta(y) \geq \zeta(x) \wedge \hat{\mathcal{C}}(y) > 0\}$;

D-VIS incorporates pixel purity and empirical data density into a single measure of pixel quality $\zeta(x) = \frac{2\bar{p}(x)\bar{\eta}(x)}{\bar{p}(x)+\bar{\eta}(x)}$, where $\bar{p}(x) = \frac{p(x)}{\max_{y \in X} p(y)}$ and $\bar{\eta}(x) = \frac{\eta(x)}{\max_{y \in X} \eta(y)}$. Thus, $\zeta(x)$ is the harmonic mean of pixel purity and density, normalized so that each is on the same scale. Note $\zeta(x)$ will be large for pixels that are modal (with high p -value) and representative of a single material class (with high η -value).

The second main function used for mode detection is

$$d_t(x) = \begin{cases} \max_{y \in X} D_t(x, y) & x = \operatorname{argmin}_{y \in X} \zeta(y), \\ \min_{y \in X} \{D_t(x, y) | \zeta(y) \geq \zeta(x)\} & \text{otherwise,} \end{cases}$$

which returns the diffusion distance at time t between x and its D_t -nearest neighbor of higher density and purity for pixels that are not ζ -maximizers. Maximizers of $\mathcal{D}_t(x) = \zeta(x)d_t(x)$ are high-density, high-purity pixels that are far in diffusion distance at time t from other high-density, high-purity pixels, making them reasonable choices as exemplars for underlying material class structure. The K maximizers of $\mathcal{D}_t(x)$ are assigned unique labels and are treated as class modes. Non-modal labels are assigned in order of non-increasing $\zeta(x)$ according to the label of their D_t -nearest neighbor that is already labeled and has a higher ζ -value.

3. ACTIVE DIFFUSION AND VCA-ASSISTED IMAGE SEGMENTATION

Though HSI segmentation can be performed without the aid of ground truth labels, incorporating the labels of just a few carefully-chosen pixels may significantly improve the predictive capacity of an HSI segmentation algorithm. *Active learning* algorithms query the ground truth labels (denoted \mathcal{C}_{GT}) of $B \in \mathbb{N}$ (called the *budget*) pixels. It is typically desired that these pixels exemplify underlying class structure, as queried points' labels are propagated to unlabeled pixels. It has been

Algorithm 2: Active Diffusion and VCA-Assisted Image Segmentation (ADVIS)

Input: X (HSI), K (# classes), N (# neighbors), σ_0 (density scale), t (diffusion time), B (budget)

Output: $\hat{\mathcal{C}}$ (HSI segmentation)

- 1 Compute $\eta(x)$, using HySime [9] to estimate m and VCA [5, 11] for spectral unmixing;
- 2 For each $x \in X$, compute $\zeta(x) = \frac{2\bar{p}(x)\bar{\eta}(x)}{\bar{p}(x)+\bar{\eta}(x)}$;
- 3 Sort X by $\mathcal{D}_t(x) = \zeta(x)d_t(x)$ in non-increasing order. Denote this sorting $\{x_{m_k}\}_{k=1}^n$. Assign $\hat{\mathcal{C}}(x_{m_k}) = \mathcal{C}_{GT}(x_{m_k})$ for $1 \leq k \leq B$;
- 4 Let $I = \{i_1, i_2, \dots, i_L\} \subset \{1, 2, \dots, K\}$ be the set of classes without a labeled point. If I is nonempty, label $\hat{\mathcal{C}}(x_{m_{B+k}}) = i_k$ for $1 \leq k \leq L$;
- 5 In order of non-increasing $\zeta(x)$, for each unlabeled $x \in X$, assign the label $\hat{\mathcal{C}}(x) = \hat{\mathcal{C}}(x^*)$, where $x^* = \operatorname{argmin}_{y \in X} \{D_t(x, y) | \zeta(y) \geq \zeta(x) \wedge \hat{\mathcal{C}}(y) > 0\}$;

shown that active learning algorithms often substantially outperform their unsupervised counterparts [3, 4, 12, 13, 14, 15].

In this section, we introduce the Active Diffusion and VCA-Assisted Image Segmentation (ADVIS) algorithm for material discrimination (see Algorithm 2). ADVIS is similar to the D-VIS clustering algorithm, with a crucial difference in the manner in which class modes are labeled. ADVIS queries the labels of the B pixels that maximize $\mathcal{D}_t(x)$. If any classes remain unlabeled after the budget expires, ADVIS reverts to unsupervised D-VIS mode estimation. By ensuring that class modes are correctly labeled in its first stage, the ADVIS algorithm improves all labeling downstream with computational complexity identical to that of D-VIS.

4. NUMERICAL EXPERIMENTS

This section illustrates the efficacy of ADVIS by comparing it against its unsupervised clustering counterpart (D-VIS) on the

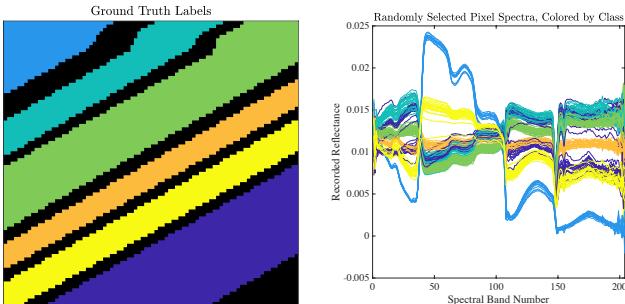


Fig. 1: Ground truth labels and randomly-selected pixel spectra for the Salinas A HSI. The class indicated in purple (8-week romaine) has high intra-class spectral variability.

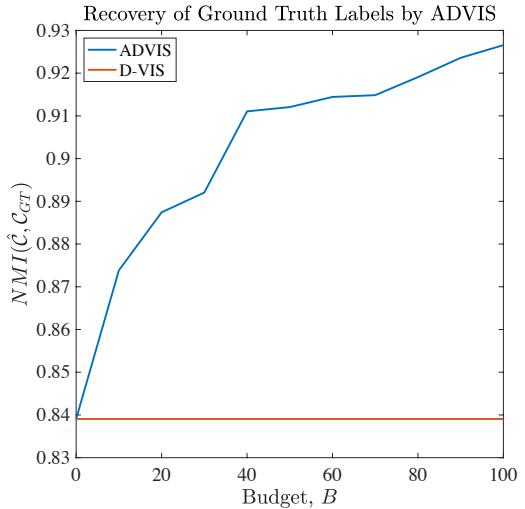


Fig. 2: Comparison of the performance of D-VIS and ADVIS as a function of budget, B . The performance of ADVIS quickly overtakes the performance of D-VIS and monotonically increases as a function of B .

Salinas A benchmark HSI (Fig. 1). The Salinas A HSI was generated using the Airborne Visible/Infrared Imaging Spectrometer sensor over Salinas Valley, CA, USA and encodes $D = 204$ spectral bands across 83×86 pixels. D-VIS and ADVIS were evaluated on labeled pixels of the Salinas A HSI with parameters $N = 320$, $\sigma_0 = 1.14 \times 10^{-3}$, and $t = 2^5$. For ADVIS, the budget B ranged $\{10, 20, \dots, 100\}$. Performance was measured using $NMI(\hat{\mathcal{C}}, \mathcal{C}_{GT})$: the normalized mutual information between an estimated partition $\hat{\mathcal{C}}$ and the ground truth labels \mathcal{C}_{GT} . Before labeling, pixel purity was averaged across 100 runs to account for VCA's stochasticity.

In Fig. 2, the performance of ADVIS is plotted against the budget B . Fig. 3 visualizes a sample of learned partitions. These results make clear that an active learning framework with even a small budget offers a major improvement in material discrimination. Indeed, though D-VIS erroneously splits the purple class (8-week romaine) in two, ADVIS correctly groups these pixels with just $B = 20$ ground truth labels. ADVIS labelings quickly converge to \mathcal{C}_{GT} as B increases, and when $B = 100$, there is little difference between the ground truth labels and the partition estimated by ADVIS. Importantly, ADVIS does not rely on spatial information, so much of the remaining error may be corrected in a spatially regularized regime [15, 16, 17]. Nevertheless, it is clear that the inclusion of a few carefully-chosen labels in ADVIS results in image segmentations of the Salinas A HSI that are substantially close to its ground truth labels.

Software to replicate numerical experiments is available on GitHub at <https://github.com/sampolk/D-VIC>.

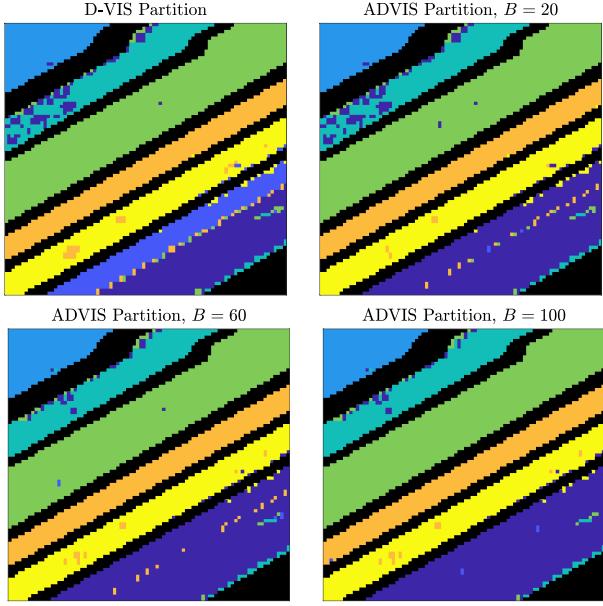


Fig. 3: Comparison of D-VIS and ADVIS on Salinas A. The major error present in the unsupervised labeling (splitting 8-week romaine) is corrected using just $B = 20$ label queries.

5. CONCLUSIONS

We conclude that an active learning framework that enables the incorporation of a few ground truth labels for material discrimination substantially improves the performance of an HSI segmentation algorithm. We expect that ADVIS can be extended for multiscale HSI segmentation, wherein a hierarchy of image segmentations is learned [8, 16]. In addition, ADVIS is likely to benefit from a modified, spatially-regularized graph, wherein edges between pixels are restricted to spatial nearest neighbors [15, 16, 17].

6. REFERENCES

- [1] M. T. Eismann, *Hyperspectral remote sensing*, SPIE, 2012.
- [2] S. L. Polk, K. Cui, R. J. Plemmons, and J. M. Murphy, “Diffusion and volume maximization-based clustering of highly mixed hyperspectral images,” *arXiv preprint arXiv:2203.09992*, 2022.
- [3] Z. Wang, B. Du, L. Zhang, L. Zhang, and X. Jia, “A novel semisupervised active-learning algorithm for hyperspectral image classification,” *IEEE Trans Geosci Remote Sens*, vol. 55, no. 6, pp. 3071–3083, 2017.
- [4] J. M. Murphy and M. Maggioni, “Unsupervised clustering and active learning of hyperspectral images with nonlinear diffusion,” *IEEE Trans Geosci Remote Sens*, vol. 57, no. 3, pp. 1829–1845, 2018.
- [5] J. M. P. Nascimento and J. M. Bioucas-Dias, “Vertex component analysis: A fast algorithm to unmix hyperspectral data,” *IEEE Trans Geosci Remote Sens*, vol. 43, no. 4, pp. 898–910, 2005.
- [6] J. Friedman, T. Hastie, and R. Tibshirani, *The elements of statistical learning*, Springer Series in Statistics, 2001.
- [7] R. R. Coifman and S. Lafon, “Diffusion maps,” *Appl Comput Harm Anal*, vol. 21, no. 1, pp. 5–30, 2006.
- [8] J. M. Murphy and S. L. Polk, “A multiscale environment for learning by diffusion,” *Appl Comput Harmon Anal*, vol. 57, pp. 58–100, 2022.
- [9] J. M. Bioucas-Dias and J. M. P. Nascimento, “Hyperspectral subspace identification,” *IEEE Trans Geosci Remote Sens*, vol. 46, no. 8, pp. 2435–2445, 2008.
- [10] Tsung-Han Chan, Wing-Kin Ma, ArulMurugan Ambikapathi, and Chong-Yung Chi, “A simplex volume maximization framework for hyperspectral endmember extraction,” *IEEE Trans Geosci Remote Sens*, vol. 49, no. 11, pp. 4177–4193, 2011.
- [11] R. Bro and S. De Jong, “A fast non-negativity-constrained least squares algorithm,” *J Chemom*, vol. 11, no. 5, pp. 393–401, 1997.
- [12] D. Tuia, M. Volpi, L. Copa, M. Kanevski, and J. Munoz-Mari, “A survey of active learning algorithms for supervised remote sensing image classification,” *IEEE J Sel Top Appl Earth Obs Remote Sens*, vol. 5, no. 3, pp. 606–617, 2011.
- [13] J. M. Haut, M. E. Paoletti, J. Plaza, J. Li, and A. Plaza, “Active learning with convolutional neural networks for hyperspectral image classification using a new Bayesian approach,” *IEEE Trans Geosci Remote Sens*, vol. 56, no. 11, pp. 6440–6461, 2018.
- [14] M. Maggioni and J. M. Murphy, “Learning by active nonlinear diffusion,” *Found Data Sci*, vol. 1, no. 3, pp. 271, 2019.
- [15] J. M. Murphy, “Spatially regularized active diffusion learning for high-dimensional images,” *Pattern Recognit Lett*, vol. 135, pp. 213–220, 2020.
- [16] S. L. Polk and J. M. Murphy, “Multiscale clustering of hyperspectral images through spectral-spatial diffusion geometry,” in *Proc IEEE Geosci Remote Sens Symp*, 2021, pp. 4688–4691.
- [17] J. M. Murphy and M. Maggioni, “Spectral–spatial diffusion geometry for hyperspectral image clustering,” *IEEE Geosci Remote Sens Lett*, vol. 17, no. 7, pp. 1243–1247, 2019.