An Adaptive Similarity Measure for Classification of Hyperspectral Signatures

Brian D. Bue, and Erzsébet Merényi, Senior Member, IEEE

Abstract—Capturing both the shape of the spectral continuum and the positions/widths of absorption bands is essential to accurately measure similarity between hyperspectral signature. Furthermore, the relative importances of these features are data dependent. In this work, we present an adaptive version of our recently proposed CICR (Continuum Intact/Continuum Removed) similarity measure which automatically determines a convex weighting between similarity measurements of Continuum Intact (CI) and Continuum Removed (CR) signatures according to input data. We describe an efficient technique to calculate an optimal weight for a linear combination of CI and CR similarity measurements. We evaluate the technique on AVIRIS spectra sampled from a well-studied urban scene, and show that our technique yields improved classification accuracy in comparison to CI or CR similarity measurements alone, and performs comparably to calculating the weight via brute-force search, at much reduced computational cost. A source code implementation of our algorithm is provided online.

Index Terms—hyperspectral, continuum removal, adaptive, similarity measure, metric learning, linear discriminant analysis

I. SPECTRAL SIMILARITY MEASURES

T HE widespread deployment of hyperspectral imaging systems on both terrestrial and planetary orbiters allows for highly detailed analysis of large scale regional surveys. Improvements in both spectral and spatial resolution of these systems drives the innovation of advanced image classification techniques. Such image classification techniques have proven an invaluable tool for analysts, as the sheer volume of the collected data renders exhaustive manual image interpretation impossible. Increasing data volumes demand algorithms which not only are sophisticated enough to yield high classification accuracies in various scenarios, but which are also efficient, capable of processing large hyperspectral datasets quickly, for example in ground-based archives, or in resource-constrained / real-time onboard applications.

A problem central to image classification is assessing similarity between spectral signatures (pixels), which is challenging due to numerous factors such as data dimensionality, noise, and environmental effects. Employing adaptive similarity measures that exploit knowledge of known classes can help mitigate these effects and improve classification accuracy, potentially with fewer training samples than a task-agnostic measure may require to achieve the same accuracy. Robust spectral similarity measures must capture both the shape of the spectra and the positions/widths of absorption bands. Most widely-used measures assume all spectral bands are of equal importance, without specific emphasis on absorption features. These measures take *Continuum-Intact* (CI) reflectance signatures as input, which often poorly capture differences in absorption features. To compensate for this, one can employ measures which characterize absorption features using *Continuum-Removed* (CR) signatures. The CR representation measures spectral absorption features by "dividing out" a spectral continuum curve. The CR representation often better captures the composition and concentration of the material a spectral signature represents than its CI counterpart. But, since signatures with considerably different continuua can have equivalent CR representations (as observed in [1], [2], and [3]), the CR representation alone can be an unreliable descriptor in classification settings.

In this work, we present an adaptive similarity measure designed specifically for hyperspectral signatures. Our measure, d_{CICR} , finds a compromise between the use of the continuum shape and absorption features by calculating a scalar weight α in a convex combination of CI and CR distance measurements that best separates a set of training classes. When spectral signatures contain significant absorption features, our measure better discriminates material classes in comparison to measures considering CI or CR signatures alone. We evaluate classification accuracy on real AVIRIS imagery from a wellstudied urban area, in scenarios consisting of prominent or subtle absorption features.

II. THE CICR MEASURE

Given a pair of hyperspectral signatures \mathbf{x}_i and $\mathbf{x}_j^{\ 1}$, each in \mathbb{R}^D , we define the d_{CICR} similarity measure as follows:

$$d_{\text{CICR}}(\mathbf{x}_i, \mathbf{x}_j, \alpha) = (1 - \alpha) d_{\text{CI}}(\mathbf{x}_i, \mathbf{x}_j) + \alpha d_{\text{CR}}(\mathbf{x}_i, \mathbf{x}_j)$$
(1)

with

$$d_{\rm CI}(\mathbf{x}_i, \mathbf{x}_j) = \left\| \frac{\mathbf{x}_i}{\|\mathbf{x}_i\|} - \frac{\mathbf{x}_j}{\|\mathbf{x}_j\|} \right\|$$
(2)

$$d_{\rm CR}(\mathbf{x}_i, \mathbf{x}_j) = \left\| \frac{\mathrm{CR}(\mathbf{x}_i)}{\|\mathrm{CR}(\mathbf{x}_i)\|} - \frac{\mathrm{CR}(\mathbf{x}_j)}{\|\mathrm{CR}(\mathbf{x}_j)\|} \right\|$$
(3)

Here, $\|\cdot\|$ is the L² norm, $\alpha \in [0, 1]$ is a weighting parameter, and CR(\cdot) performs continuum removal. We estimate the continuum of a given spectrum by fitting a piecewise linear function to local maxima using the procedure described in [4]. Observations on the continuum are assigned values of zero, and absorption features (observations between local maxima)

B. Bue is with the Department of Electrical and Computer Engineering and E. Merényi is with the Department of Statistics, Rice University, Houston TX, 77005, USA. E-mail: {bbue,erzsebet}@rice.edu. Color figures available in the online version of this paper.

¹In this and subsequent sections, we denote scalar variables in italics, vectors in bold, and matrices in bold caps.

are assigned values in the [0, 1] range, proportional to their relative distance from the estimated continuum. Because the continuum removal procedure is sensitive to spurious local maxima, we smooth each signature using a moving average filter before performing continuum removal. While smoothing may mask small absorption features, such features are often close to the noise floor of the sensor, and we accept this loss in specificity in favor of noise reduction. In our experiments using AVIRIS data, smoothing windows ranging from three to five bands (0.03-0.05 μ m) have performed well.

Dividing the CI and CR signatures by the L^2 norm is important for several reasons. First, scaling each signature by its norm maps both CI and CR distances to a common range. This allows us to better tune the weight α to combine the CI and CR similarity measurements according to input data. Second, for CI signatures, this scaling mitigates linear effects caused by differing illumination conditions (while preserving spectral angles). Lastly, scaling CR signatures by their respective norms can accentuate discriminative absorption features and can improve performance when classifying CR signatures.

The measure described in this work differs from our original formulation in [4]. First, the convex combination of CI and CR terms yields more consistent performance than applying α to the CR term only. Second, due to the nature of continuum estimation, CR signatures contain many values near zero, which provides little discriminating information between signatures when combined with the CI distance measure. Scaling the CR signatures uniformly by the variance of all CR distances (as we originally described) does not resolve this issue, whereas scaling each signature by its norm allows the most prominent absorption features to play greater roles in discrimination.

III. COMBINING CI AND CR DISTANCE MEASUREMENTS

Figure 1 gives an overview of the methodology we use to calculate the weight parameter α in Equation 1, which we describe in detail below.



Fig. 1. Processing steps for calculating d_{CICR} weight parameter α .

Given a set of N vectors $\{\mathbf{x}_i\}_{i=1}^N, \mathbf{x}_i \in \mathbb{R}^D$ belonging to K classes, with labels $y_i \in [1, K]$, we calculate α using a method inspired by multiclass linear discriminant analysis (LDA) ([5], [6]). Multiclass LDA computes the vector w that maximizes the Rayleigh quotient (using the formulation given in [7])

$$S = (\mathbf{w}^T \mathbf{M}_B \mathbf{w}) (\mathbf{w}^T \mathbf{M}_W \mathbf{w})^{-1}$$
(4)

where \mathbf{M}_B and \mathbf{M}_W are (symmetric, positive-definite) between-class separation and within-class scatter matrices. We form the \mathbf{M}_B and \mathbf{M}_W matrices according to the capability of each of the $\{d_{CI}, d_{CR}\}$ measures to separate the given classes. Specifically, we define the between-class separation, s_b , and within-class scatter, s_w , between measures d_1, d_2 as:

$$s_b(d_1, d_2) = \sum_{j=1}^K N_j d_1(\boldsymbol{\mu}_j, \overline{\boldsymbol{\mu}}) d_2(\boldsymbol{\mu}_j, \overline{\boldsymbol{\mu}})$$
(5)

$$s_w(d_1, d_2) = \sum_{j=1}^K \sum_{i:y_i=j} d_1(\mathbf{x}_i, \boldsymbol{\mu}_j) d_2(\mathbf{x}_i, \boldsymbol{\mu}_j) \quad (6)$$

where $\{\mu_j\}_{j=1}^K$ are the mean vectors of each of the K classes, $\overline{\mu}$ is the mean of the μ_j , and N_j is the number of samples in class j. We then form the \mathbf{M}_B and \mathbf{M}_W matrices as follows:

$$\mathbf{M}_B = \frac{1}{N} \begin{bmatrix} s_b(d_{\mathrm{CI}}, d_{\mathrm{CI}}) & s_b(d_{\mathrm{CI}}, d_{\mathrm{CR}}) \\ s_b(d_{\mathrm{CI}}, d_{\mathrm{CR}}) & s_b(d_{\mathrm{CR}}, d_{\mathrm{CR}}) \end{bmatrix}$$
(7)

$$\mathbf{M}_{W} = \frac{1}{N} \begin{bmatrix} s_{w}(d_{\mathrm{CI}}, d_{\mathrm{CI}}) & s_{w}(d_{\mathrm{CI}}, d_{\mathrm{CR}}) \\ s_{w}(d_{\mathrm{CI}}, d_{\mathrm{CR}}) & s_{w}(d_{\mathrm{CR}}, d_{\mathrm{CR}}) \end{bmatrix}$$
(8)

The first (largest) eigenvector of $\mathbf{M}_{W}^{-1}\mathbf{M}_{B}$, w, maximizes Equation 4, with separation S equal to the corresponding eigenvalue [7]. The components of $\mathbf{w} = [w_{CI}, w_{CR}]$ provide a weighting of the CI and CR distances with good class separation on training data, but is not necessarily convex (as we require in Equation 1), and may not generalize well to test data. Because Rayleigh quotients are invariant with respect to scaling of w (i.e., for any c > 0, $c\mathbf{w}$ also maximizes Equation 4) [8], we can scale the components of w to a convex range by dividing each component by $\|\mathbf{w}\|_{1}$. This yields the convex pair $\{w_{CI}/\|\mathbf{w}\|_{1}, w_{CR}/\|\mathbf{w}\|_{1}\} = \{(1 - \alpha), \alpha\}$, as desired.

As Equation 4 may become ill-posed, we regularize the within-class scatter matrix via a shrinkage operator:

$$\mathbf{M}'_W = (1 - \lambda)\mathbf{M}_W + \lambda \mathcal{I} \tag{9}$$

where $\lambda \in (0, 1)$ is a regularization parameter, and \mathcal{I} is the (2×2) identity matrix. In practice, we select λ via cross-validation, using the methodology described in the next section. We provide implementations of our LDA-based algorithm and the continuum removal algorithm we used at the following url http://www.ece.rice.edu/~bdb1/#code.

IV. CASE STUDY: AVIRIS SPECTRAL SIGNATURES

The starting point of the work described here is a set of reflectance spectra sampled across distinct material species from a Low-Altitude Airborne Visible / Infrared Imaging Spectrometer (AVIRISLA) image of Ocean City, MD. This image (acquired Nov 5, 1998, with 4 m / pixel spatial resolution, in 220 spectral bands from 0.4 $2.5 \mu m$) was

analyzed in previous work to capture spectral clusters, verify them against field knowledge and identify materials they represent, as reported in detail in [9]. The 35 clusters resulting from [9] guided the extraction of a trustworthy representative subset of spectra for this study, by stratified random sampling across 14 of those 35 clusters for which material identification was unambiguous, and which served the methodology design for this work. The design is explained below. The work in [9] performed all the necessary pre-processing of this image including atmospheric correction, conversion of radiances to reflectances, and normalization to cancel illumination effects. These preprocessing steps are duly described in [9]. For this work the reflectance spectra were extracted from the already pre-processed Ocean City image. The number of samples per class is constrained by the size of the smallest cluster (≈ 145 samples), and we fix the number of samples per class to 100 for most of our analysis, and characterize classification accuracy vs. training set size later in this section (Figure 4).

This work is a deeper analysis of previous work ([3], [4]) that examines three different spectral scenarios specifically constructed to contrast the performance of the proposed method. In the first, all samples contain only minor absorptions, where we define a "minor" absorption as one with no CR band depths greater than threshold τ ; we use $\tau = 0.1$ (10%) absorption with respect to the continuum) in this work. In this case, we expect similar classification accuracies from the d_{CICR} and d_{CI} measures, since the CR signatures lack prominent absorption features (and therefore are flat and uninformative). The classes in this scenario consist of asphalt rooftop materials (class A), roads/parking lots (classes I, J, T, W, and h), and dry beach sand (class e). Figure 2 (top) shows the CI and CR mean signatures for these classes. In the second scenario, all signatures contain one or more *major* absorptions, where we define major absorptions as those with CR values greater than τ . Here, we expect a more significant boost in accuracy in comparison to the minor absorption scenario, as the CR signatures are more informative in this case. The subset of data in this scenario consists of vegetation (classes L and M), a tennis court (class C), wet sand (classes O and Q), and composite rooftop materials (classes D and U). Figure 2 (bottom) shows the CI and CR mean signatures for these classes. The 7 spectral species in each of the major and minor absorption categories are relatively "pure" representatives of their respective species. The last scenario, combined, consists of all classes from both major and minor absorption scenarios. We expect to see notable performance gains with the d_{CICR} measure in this case, as both the CI and CR signatures provide information to distinguish between classes.

We present results using a minimum distance to class means (MinDist) classifier with 5-fold random stratified sampling, using 50% for training and the remaining 50% for testing. We calculate α by maximizing $S(\alpha)$ as described in Section III. We compare this α value to the α_{LS} value obtained by line search (LS) on a uniformly spaced range of 100 points, $\alpha_{LS} \in (0, 1)$, which yields the highest classification accuracy. We calculate classification accuracy according to accuracy = (# of True Positives)/(# of Samples) and we report accuracy on test data only. Accuracies produced via line search are an



Fig. 2. Top: Mean CI and CR signatures for minor absorption classes. Top inset: detail view of minor absorption signatures, wavelengths $1.5-2.5\mu$ m. Bottom: Mean CI and CR signatures for major absorption classes. The large disconnected regions near 1.3-1.5 and 1.7-2.0 μ m consist of bands removed due to water saturation. CI signatures are scaled by their L² norms to compensate for varying illumination conditions.

approximate upper bound on achievable accuracy. We choose the regularization parameter λ for each scenario on a hold-out set comprising 50 samples per class, which is not included in the training/testing samples. For our data, we select the λ with the best classification accuracy on the hold-out set from 10 uniformly spaced values in [0.001,0.1]. We chose this range because smaller λ values tended to yield ill-posed solutions, and larger values did not improve classification accuracy in any of the three scenarios – regardless of α . We calculate λ once for each of the minor, major and combined scenarios, and use the same value for each cross-validation fold. We also reject any λ values which produce solutions to $\mathbf{M}_W^{-1}\mathbf{M}_B$ with no positive eigenvalues, as such λ yield rank-deficient \mathbf{M}'_w (Equation 9).

Figure 3 gives the overall and per-class classification accuracies for $\alpha \in [0, 1]$. The vertical magenta dashed line marks the α value determined by maximizing Equation 4, and the black vertical line gives α_{LS} . Table I provides average accuracies for each measure. In all three scenarios, small alpha values (< 0.3) yield the highest classification accuracies (though we do not constrain the search to this range). This indicates that, for this dataset, CI signatures are more robust descriptors than CR signatures for classification. This is particularly obvious in the minor absorption scenario (Figure 3, top), where the CR signatures lack discriminative features. Here, classification accuracy using d_{CI} is close to d_{CICR} , and both our method and the line search produce α values near zero.

For major absorption classes (Figure 3, middle), note that the rate of decrease in classification accuracy is less dramatic as α approaches one, by comparison to the minor (top) and combined (bottom) absorption scenarios. This indicates that the CR signatures provide additional discriminating information, which increases the α values yielding higher classification accuracy. Correspondingly, the maximum separation also shifts towards larger α values. While α and α_{LS} differ the most in this scenario, their corresponding classification accuracies are not far apart (97.4% vs. 98.4%). Both are improvements over the baseline d_{CI} accuracy (1.5% and 2.5% relative improvements for our LDA-based α and α_{LS} , respectively).

In the combined scenario (Figure 3, bottom), due to potentially increased class confusion between signatures (compared to the other two scenarios), locating a compromise between the CI and CR terms is challenging. As we see in Figure 3, the mean classification accuracy for this scenario generally falls between the mean accuracies of the minor and major absorption scenarios. However, we see the most significant improvement, over the baseline d_{CI} method, in classification accuracy in this scenario (4.5%, by comparison of the thick black line to the horizontal red dashed line in Figure 3, bottom), vs. the other two scenarios, since both the CI and CR representations provide complementary information to discriminate the classes. This is noteworthy given that the CI and CR classification accuracies in the combined scenario are close to those of the minor absorption scenario (88.5% vs. 88.2% and 66.1% vs. 66.4%, in the combined vs. minor scenarios, respectively), yet the relative improvement in the minor scenario is, not surprisingly, lower 1.7%.



Fig. 3. α vs. per-class d_{CICR} classification accuracy for minor (top), major (middle) and combined (bottom) absorption classes. Colored lines indicate per-class accuracies, and the black solid line gives the overall classification accuracy. The black vertical bar gives α_{LS} , and the magenta vertical bar gives α . The horizontal lines give the CI (red, $\alpha = 0$) and CR (blue, $\alpha = 1$) classification accuracies. Because the CI representation is generally more informative than CR, α values tend towards zero, but larger values occur in cases when the CR representation provides additional discrimination information (as in the major and combined absorption class scenarios).

Figure 4 characterizes the relationship between the number of labeled samples available for training vs. classification accuracy. As before, classification accuracy is measured over 5 stratified 50/50 splits into training/test data. In every case,

	$d_{\mathbf{CI}}$	$d_{\rm CR}$	$d_{\text{cicr}} (\alpha \pm \sigma_{\alpha})$	$d_{\text{cicr}} (\alpha_{\text{LS}} \pm \sigma_{\alpha_{\text{LS}}})$
Minor	88.5	66.1	90.1 (0.0510±0.0375)	90.4 (0.0485±0.0003)
Major	95.9	76.7	97.4 (0.1493±0.0041)	98.4 (0.0770±0.0195)
Combined	88.2	66.4	92.4 (0.0903±0.0006)	92.6 (0.0670±0.0057)
TABLE I				

Average classification accuracy obtained with each of the $d_{\rm CI}$, $d_{\rm CR}$, $d_{\rm CICR}$ measures shown in Figure 3. Mean and standard deviation (σ) of α values for $d_{\rm CICR}$ measures are given in parentheses. The most accurate measure is given in bold text.

the d_{CICR} -based classifiers match or outperform the d_{CI} -based classifier. Additionally, our LDA-based technique for calculating α performs comparably to brute force search when a sufficient quantity (about 50 / class) of training samples are available. We observe the most significant performance gains of the three scenarios in the combined scenario, where the d_{CICR} measure can exploit absorption features to separate the classes belonging to the major and minor absorption scenarios, and also can capitalize on the absorption characteristics of individual classes. d_{CR} performs the worst in all three cases, and is not shown in Figure 4 to emphasize the performance of the better performing measures.



Fig. 4. Classification accuracy vs. number of samples per class for minor absorption (top), major absorption (middle) and combined (bottom) scenarios. In each scenario, the LDA-based $d_{\rm CICR}$ measure (magenta line) outperforms the baseline $d_{\rm CI}$ -based classifier (blue line), and with a sufficient number of training samples (~20-30, scenario dependant) achieves classification accuracy comparable to line search (red line). $d_{\rm CR}$ -based classification accuracy not shown above due to significantly lower accuracies (~65-77%) in comparison to the $d_{\rm CI}$ and $d_{\rm CICR}$ -based classifiers.

Determining α using our LDA-based method is significantly less computationally expensive than via brute force search (i.e., α_{LS}). Quantitatively, assuming N samples of dimensionality D belonging to K classes, we first compute the continuumremoved representation of each spectrum using our piecewise linear continuum estimation procedure – an O(D) operation per spectrum. Then, given the set of (precomputed) class means, a MinDist classifier must compare each signature to each class mean, an O(DNK) operation. Let A be the number of values α_{LS} can take in (0,1) (in this work, we choose A=100). Using brute force search, we apply the O(DNK)MinDist classifier A times. With the LDA-based method, calculating the (symmetric) \mathbf{M}_B involves three O(DK) operations and \mathbf{M}_W involves three O(DN) operations, and calculating the eigendecomposition of the $(2 \times 2) \mathbf{M}_W^{-1} \mathbf{M}_B$ matrix can be performed in constant time. This amounts to roughly an A-fold improvement in performance by the LDA-based method over line search. Because A must be large enough to adequately cover the weight parameter space, our method is an order of magnitude faster than brute-force search.

V. DISCUSSION AND FUTURE WORK

As hyperspectral sensors improve in sensitivity and spectral resolution, the choice of spectral similarity measure will play a significant role in resolving subtle differences between classes. We have shown in this work that with a small amount of preprocessing, we can improve classification performance over traditional, task-agnostic similarity measures using our adaptive, hyperspectral domain-specific similarity measure. We also provided an efficient method to calculate the weight parameter used in our similarity measure which yields classification accuracies within 1% of brute-force search over the range of possible weight values.

Leveraging techniques employed by spectroscopists (i.e., continuum removal) is a natural approach to characterize spectral similarity. However, determining the best manner in which to exploit this information is challenging. First, numerous methods for continuum estimation and removal exist, ranging from fitting piecewise linear functions [1] to more involved techniques that exploit high-order derivatives of spectral signatures [10]. Determining which method performs best in terms of classification accuracy has not been widely explored in the literature (to the best of our knowledge) and is a subject of future work.

We experimented with other methods for learning the α weighting in the CICR measure, but our initial experiments in those directions produced mixed results. We believe an issue with such methods (e.g., [11]) is that the learning problem is transformed from a multiclass classification problem to a binary classification problem where spectra from "similar" vs. "dissimilar" classes are grouped into two meta-classes, thereby discarding discrimination information with respect to individual classes. We observed a similar issue with other multiclass metric learning techniques in previous work [12].

When continuum removal detects spurious absorption features L^2 -normalization may exacerbate noise. In such cases, it may be advantageous to adopt alternative normalization schemes to map the CI and CR signatures to a common range, such as scaling signatures by their respective (CI or CR) standard deviations. Using such a "global" scaling factor, we would not accentuate noise on individual signatures. However, this would not provide the same degree of contrast between classes as L^2 normalization provides. The tradeoffs involved with these normalization schemes is a subject of future investigation.

As mentioned in Section III, normalization of spectral signatures plays a significant role in the performance of this technique. In particular, scaling CR signatures by their L^2 norms often improved the classification accuracy by giving additional weight to discriminative absorption features. However, in some cases this normalization may not be desirable, for

example, if the CR signatures fail to capture the most relevant absorptions, perhaps due to noise and/or poorly estimated continua. Other normalization techniques such as scaling by the within-class variances may be beneficial in such cases.

ACKNOWLEDGEMENTS

We would like to thank Dr. David Thompson, Dr. Kiri Wagstaff, and Prof. Devika Subramanian for their invaluable input and helpful discussions, and Prof. Bea Csathó for providing the Ocean City imagery and help with the interpretation of spectral clusters.

REFERENCES

- R. N. Clark and T. Rousch, "Reflectance spectroscopy: quantitative analysis techniques for remote sensing applications," *Journal of Geophysical Research-Solid Earth*, vol. 89, no. B7, pp. 6329–6340, 1984.
- [2] E. Howell, E. Merényi, and L. Lebofsky, "Classification of asteroid spectra using a neural network," *Journal of geophysical research*, vol. 99, no. E5, pp. 10847–10,865, 1994.
- [3] B. Bue, E. Merényi, and B. Csathó, "Automated labeling of materials in hyperspectral imagery," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 48, no. 11, pp. 4059–4070, 2010.
- [4] B. D. Bue, E. Merényi, and B. Csathó, "Automated Labeling of Segmented Hyperspectral Imagery Via Spectral Matching," *Proc. IEEE* WHISPERS, Aug. 2009.
- [5] R. Fisher, "The use of multiple measurements in taxonomic problems," Annals of eugenics, vol. 7, no. 2, pp. 179–188, 1936.
- [6] C. Rao, "The utilization of multiple measurements in problems of biological classification," *Journal of the Royal Statistical Society*, vol. 10, no. 2, pp. 159–203, 1948.
- [7] T. Hastie, R. Tibshirani, and J. H. Friedman, *The Elements of Statistical Learning*, 2nd ed. Springer, Feb. 2011.
- [8] R. A. Horn and C. A. Johnson, *Matrix Analysis*, 1st ed. Cambridge University Press, 1985.
- [9] E. Merényi, B. Csathó, and K. Tasdemir, "Knowledge discovery in urban environments from fused multi-dimensional imagery," *Proc. 4th IEEE GRSS/ISPRS Joint Workshop on Remote Sensing Data Fusion over Urban Areas*, pp. 1–13, 2007.
- [10] A. Brown, "Spectral curve fitting for automatic hyperspectral data analysis," *IEEE Trans. on Geoscience and Remote Sensing*, vol. 44, no. 6, pp. 1601–1608, Jun. 2006.
- [11] A. Woznica, A. Kalousis, and M. Hilario, "Learning to combine distances for complex representations," *Proceedings of the 24th International Conference on Machine Learning*, 2007.
- [12] B. Bue, D. Thompson, M. Gilmore, and R. Castaño, "Metric learning for hyperspectral image segmentation," *Proceedings of the 3rd IEEE* WHISPERS, Jun 2011.