

USING SPATIAL CORRESPONDENCES FOR HYPERSPECTRAL CLASS KNOWLEDGE TRANSFER: EVALUATION ON SYNTHETIC DATA

Brian D. Bue, Erzsébet Merényi

Department of Electrical and Computer Engineering, Rice University, Houston, TX, 77005, U.S.A.
Email: {bbue, erzsebet}@rice.edu

ABSTRACT

We describe a proof of concept for class knowledge transfer from a labeled hyperspectral image to an unlabeled image, captured with a different (hyper-/multi-spectral) sensor, when the spatial extents of the images partially overlap. By defining a set of spatio-spectral correspondences between the labeled source image and the unlabeled target image, we create a mapping between the images we can use to propagate labels from the source to the target image. This mapping allows us to classify the target image using the source labels without manually defining training labels in the target image. We evaluate the technique using state of the art synthetic hyperspectral imagery.

Index Terms— hyperspectral, knowledge transfer, synthetic, classification, DIRSIG, HYDICE, MASTER

1. KNOWLEDGE TRANSFER FOR HYPERSPECTRAL IMAGE CLASSIFICATION

Machine learning techniques have been widely used to classify remotely sensed images into known land cover or surface material types. Once an image has been classified, however, transferring that knowledge to new imagery is often quite difficult. This is particularly challenging when the images are captured using distinct sensor types, where it is necessary to reconcile differences in spectral resolution, sensor noise, and capture geometry. Furthermore, dynamic changes in atmospheric and illumination conditions can dramatically alter the composition of the image data between capture times. Due to the numerous unknowns to be considered in the knowledge transfer problem, a sensible approach for algorithm development and validation is to analyze synthetic data. In recent years, techniques for generating synthetic hyperspectral imagery have made significant advances in modeling complex physical and chemical phenomena, and are now capable of generating imagery of exceptional detail and realism. Additionally, as the environmental conditions and surface materials within synthetic imagery are user-defined, one can generate imagery with characteristics similar to real-world scenes.

The knowledge transfer approach is not a new one. In the remote sensing domain, one of the first works employing this approach was Shashahani et al. [1]. Here, a set of unlabeled samples are used to improve the classification accuracy of a maximum a posteriori classifier (with assumptions on the distribution of the unlabeled samples) by iteratively labeling unlabeled samples, and then adding those samples with their classified labels to the training set. An analogous approach taken in other works ([2], [3], [4] for example) is to characterize spectral, spatial or temporal relationships captured by a trained classifier for a particular dataset, and then propagate class information where similar relationships in new data are found.

In this work, we present a technique for knowledge transfer between a labeled source image and an unlabeled target image, captured by different sensors. We assume the spatial extents of the images partially overlap, which provides set of natural correspondences between the images. These correspondences allow us to apply weights to the class similarity function to account for spectral differences between sensors. By defining a threshold on a class similarity function, we prevent new classes present in the target image but not in the source image from being incorrectly classified, allowing for discovery of new classes. We evaluate our technique using state of the art DIRSIG synthetic hyperspectral image data [5].

2. THE RELTRANS ALGORITHM

We seek to estimate the probability, $p(y^S | \mathbf{x}^T)$ ¹, that a target image pixel \mathbf{x}^T belongs to source image class $y^S \in [1, k_S]$. Our data consists of two disjoint sets: spectra from the source image $\mathbf{S} = \{\mathbf{x}_1^S, \dots, \mathbf{x}_{n_S}^S\}$ with class labels $\mathbf{y}^S = \{y_1^S, \dots, y_{n_S}^S\}$, and spectra from the target image $\mathbf{T} = \{\mathbf{x}_1^T, \dots, \mathbf{x}_{n_T}^T\}$, where n_S, n_T are the number of samples in \mathbf{S} and \mathbf{T} , respectively. Each $\mathbf{x}_i^S, i \in [1, n_S]$ and $\mathbf{x}_j^T, j \in [1, n_T]$ is a d -dimensional vector, where d is the number of bands in \mathbf{S} and \mathbf{T} . If the source and target images do not have the same number of image bands, we upsample the lower spectral resolution image to the wavelengths of the higher resolution image.

In most traditional classification settings, it is assumed that both the training (source) and testing (target) data come from the same distribution. If both the source and target image were captured by the same sensor (under similar environmental conditions), one may make this assumption and use these methods to estimate $p(y^S | \mathbf{x}^T)$. But, since we compare spectra from images captured using different sensor types, we cannot assume that the distributions of the source and target images are the same. Thus, we need additional information to map the domain of the target image to the domain of the source image. Here, that information comes in two forms:

(1) Correspondence Spectra: We assume that some commonalities exist between the two images - otherwise there would be nothing to gain in knowledge transfer. To exploit these commonalities, we define a set of n_C correspondence spectra $\mathbf{C} = \{(\mathbf{c}_1^S, \mathbf{c}_1^T, l_1^S), \dots, (\mathbf{c}_{n_C}^S, \mathbf{c}_{n_C}^T, l_{n_C}^S)\}$, where \mathbf{c}_i^S and \mathbf{c}_i^T are pixels with matching spectral footprints (i.e., matching materials) in the source and target images, respectively, and $l_i^S \in [1, k_S]$ is the label assigned to \mathbf{c}_i^S , for $i \in [1, n_C]$. In this work, we assume that these spectra are pixels at the same spatial locations in an overlapping sub-region shared by the source and target images, but, in principle, these correspondences need not be determined by spatial relationships.

(2) Relation Vectors: Defining a robust comparison function between samples generated from different distributions can be a difficult task. In such cases, it can be advantageous to capture structured,

B. Bue is supported by NASA GSRP Fellowship NNX0AR79H. E. Merényi is partially supported by grant NNG05GA94G from the Applied Information Systems Research Program, NASA, Science Mission Directorate.

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

relative relationships between classes within each dataset (i.e., *intra*-dataset relationships), and use these relationships to evaluate similarity between datasets (i.e., *inter*-dataset relationships). In this work, we characterize intra-dataset relationships in the following manner: given a d -dimensional data vector \mathbf{v} (pixel) and an $k_S \times d$ matrix of prototype (e.g., class mean) vectors \mathbf{M} , we define the *relation vector* between \mathbf{v} and \mathbf{M} as

$$\text{rel}(\mathbf{v}, \mathbf{M}) = \left[\frac{d(\mathbf{v}, \mathbf{m}_1)}{\sum_j^{k_S} d(\mathbf{v}, \mathbf{m}_j)}, \dots, \frac{d(\mathbf{v}, \mathbf{m}_{k_S})}{\sum_j^{k_S} d(\mathbf{v}, \mathbf{m}_j)} \right] \quad (1)$$

where $d(\cdot, \cdot)$ is a distance measure (e.g., Euclidean distance). $\text{rel}(\mathbf{v}, \mathbf{M})$ is k_S -dimensional unit vector whose i -th element is the spectral discriminatory probability [6] which estimates the likelihood of distinguishing \mathbf{v} from each representative $\mathbf{m}_i \in \mathbf{M}$, $i \in [1, k_S]$. We can then evaluate the *relation similarity* $\text{relsim}(\mathbf{r}_1, \mathbf{r}_2)$ between a pair of relation vectors \mathbf{r}_1 and \mathbf{r}_2 as follows:

$$\text{relsim}(\mathbf{r}_1, \mathbf{r}_2) = 1 - \frac{1}{2} \|\mathbf{r}_1 - \mathbf{r}_2\|_2 \quad (2)$$

where $\|\cdot\|_2$ is the L^2 norm. $\text{relsim}(\mathbf{r}_1, \mathbf{r}_2)$ yields values in the $[0, 1]$ range with similar relation vectors taking values near one.

Algorithm 1 describes the **RelTrans** procedure, which calculates the class similarity matrix Σ and predictions \mathbf{p} for target data \mathbf{T} . The algorithm first calculates the class means of the source (training) data and correspondence spectra. The class means of the source pixels \mathbf{S}_μ are used to calculate relation vectors for the means of the correspondence spectra, and for the \mathbf{S}_μ vectors themselves. A target pixel $\mathbf{T}(i)$, $i \in [1, n_T]$ is classified by selecting the source label with the highest relsim score between the pixel relation vector \mathbf{r}_i and the source relation vectors \mathbf{S}_R , weighted by the relsim scores between \mathbf{r}_i and the source/target relation vectors \mathbf{C}_R^S and \mathbf{C}_R^T .

Algorithm 1 RelTrans

Input: $n_S \times d$ matrix of source pixels \mathbf{S} . $n_T \times d$ matrix of target pixels \mathbf{T} . Set of n_C correspondence points $\mathcal{C} = \{(\mathbf{c}_1^S, \mathbf{c}_1^T, l_1^S), \dots, (\mathbf{c}_{n_C}^S, \mathbf{c}_{n_C}^T, l_{n_C}^S)\}$. Similarity threshold τ .

Output: $n_T \times k_S$ class similarity matrix Σ , length n_T prediction vector \mathbf{p} .

- 1: Calculate class means for source data, source and target correspondence spectra: $\{\mathbf{S}_\mu, \mathbf{C}_\mu^S, \mathbf{C}_\mu^T\}$
 - 2: Calculate the relation vectors for source class means and source and target correspondence spectra:
 $\mathbf{S}_R \leftarrow [\text{rel}(\mathbf{S}_\mu(j), \mathbf{S}_\mu)]_{j=1}^{k_S}$, $\mathbf{C}_R^S \leftarrow [\text{rel}(\mathbf{C}_\mu^S(j), \mathbf{C}_\mu^S)]_{j=1}^{k_S}$,
 $\mathbf{C}_R^T \leftarrow [\text{rel}(\mathbf{C}_\mu^T(j), \mathbf{C}_\mu^T)]_{j=1}^{k_S}$
 - 3: **for** $i = 1$ **to** n_T **do**
 - 4: Calculate relation vector for current target pixel:
 $\mathbf{r}_i \leftarrow \text{rel}(\mathbf{T}(i), \mathbf{C}_\mu^T)$
 - 5: $\Sigma(i) \leftarrow \text{relsim}(\mathbf{r}_i, \mathbf{S}_R) \cdot \text{relsim}(\mathbf{r}_i, \mathbf{C}_R^S) \cdot \text{relsim}(\mathbf{r}_i, \mathbf{C}_R^T)$
 - 6: $\mathbf{p}(i) = \begin{cases} \underset{j}{\text{argmax}} \Sigma(i, j) & \text{if } \Sigma(i, j) > \tau \\ 0 & \text{otherwise} \end{cases}$
 - 7: **end for**
-

3. DIRSIG SYNTHETIC HYPERSPECTRAL IMAGERY

We analyze synthetic hyperspectral imagery generated with the RIT Digital Imaging and Remote Sensing Image Generation (DIRSIG) [5] model. We study a subregion of the the RIT “Megascene” [7], with 400x400 pixels at 4m/pixel resolution. Spectral responses are modeled after the HYDICE [8] instrument, with 210 bands over 0.4-2.5 microns. Radiance values are converted to reflectance using empirical line correction in ENVI [9], and illumination normalization

is performed by dividing each pixel by its Euclidean norm. Additional preprocessing details are described in [10]. We extract two partially overlapping sub-images (“Source” and “Target” in Fig. 1). The source image remains at HYDICE spectral resolution, while the target image is downsampled to MASTER [11] spectral resolution. Initial experiments using spectral responses modeled after the 128-band HyMap [12] instrument proved trivially classifiable with a baseline (Minimum Distance) classifier. Thus, we opted for the lower spectral resolution of the MASTER instrument, with 23 bands in the 0.4-2.5 micron range. After upsampling the MASTER spectra back to the HYDICE wavelengths and removing saturated water absorption bands in both images, 159 wavelengths remain for evaluation. Examples of HYDICE spectra and their MASTER equivalents are shown in Fig. 2.

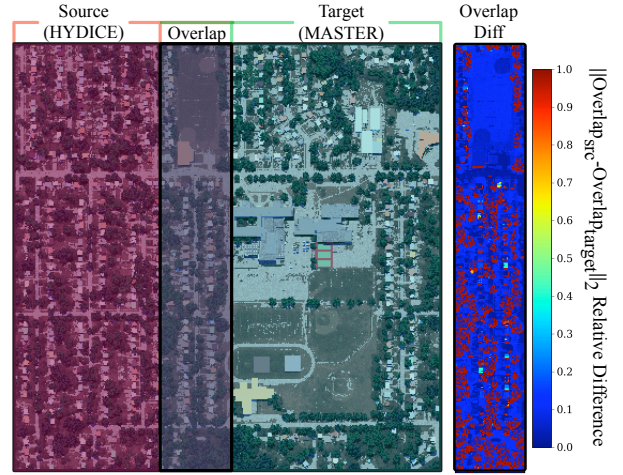


Fig. 1: Source (left, red tint), and Target (right, green tint) sub-images of the RIT DIRSIG synthetic image. The source image remains at HYDICE spectral resolution, and the target image is downsampled to MASTER spectral resolution. The target image is then upsampled back to HYDICE spectral resolution. Correspondence spectra are selected from the overlap region (blue tint). The relative difference between source and target pixels (right) is greatest for shaded pixels (see Fig. 2, class C).

4. HYDICE VS. MASTER RESULTS

We consider three knowledge transfer scenarios. In the first scenario (I ($k_S = k_T$) in Table. 1), all classes present in the source data are present in the target data. Here, the source and training datasets share a similar class structure, so we expect the highest accuracy in this case. In the second scenario (II ($k_S > k_T$) in Table. 1), the source data contains several classes not present in the target data. Here, we potentially increase inter-class confusion by including extraneous source classes. In the third scenario (Table. 1, III ($k_S < k_T$)), the target data contains classes not present in the source data. In this case, the classes present only in the target data represent “undiscovered” classes. Here, we expect the lowest overall classification accuracy, since these classes cannot be labeled correctly given the source data. This can be improved by applying a threshold to the class similarity vector to flag samples as “unknowns.” In each scenario, we perform random stratified sampling of the Self-Organizing Map-based segmentation described in [10] and collect 2000 labeled spectra from each of the source and target images (target labels are used only in validation). An additional 300 labeled spectra are selected from the overlap region as correspondence points. For simplicity, we assume that at least one correspondence point for each class is available, but we are experimenting with situations when a

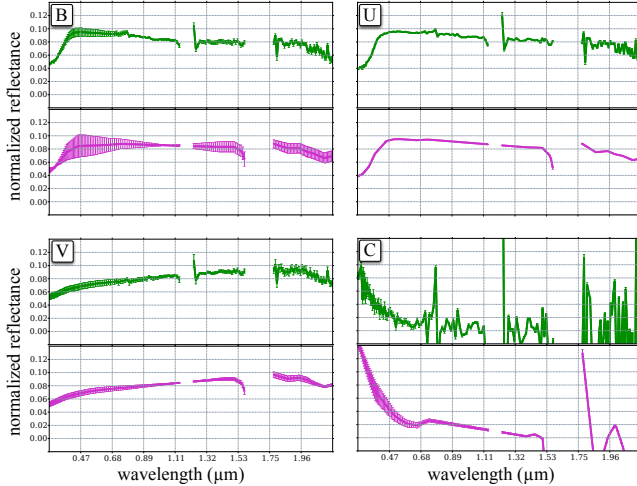


Fig. 2: Mean and std. deviation for classes B,U,V and C from source (HYDICE, green) and target (MASTER, magenta) images. Class B consists of a combination of tan asphalt shingle and gray gravel roof spectra, and is often confused with class U (brown asphalt shingles) and class V (black and gray asphalt materials). Class C is an example of a shadow class consisting of several (heterogeneous) materials. Due to the heterogeneity, such classes tend to have low relsim scores.

	I ($k_S = k_T$)	II ($k_S > k_T$)	III ($k_S < k_T$)
MinDist	0.858 (1.898e-03)	0.825 (4.922e-04)	0.579 (2.054e-04)
MinDist _{rel}	0.947 (3.355e-04)	0.877 (4.197e-04)	0.640 (1.038e-04)
RelTrans _{src}	0.947 (3.355e-04)	0.877 (4.197e-04)	0.640 (1.038e-04)
RelTrans _{corr}	0.990 (3.454e-06)	0.933 (1.736e-05)	0.664 (2.036e-06)
RelTrans _{thresh}	0.990 (8.343e-08)	0.933 (1.599e-06)	0.744 (1.210e-07)
Improvement (%)	13.2%	10.8%	8.5%

Table 1: Mean and variance of cross-validated classifier (test) accuracies for HYDICE (source) vs. MASTER (target) data, along with the improvement in accuracy between **MinDist** and **RelTrans_{corr}**.

subset of the source classes lack correspondence points. We compare results between scenarios using the following classifiers: Minimum Euclidean Distance (MinDist), MinDist between relation vectors instead of image pixels (MinDist_{rel}), RelTrans algorithm using source data only (i.e., $\Sigma(i) = \text{relsim}(\mathbf{r}_i, \mathbf{S}_R)$) in evaluating the class similarity function (RelTrans_{src}), RelTrans algorithm using source and target correspondences in addition to training data (RelTrans_{corr}), and finally, a thresholded version of RelTrans_{corr} with pixels with relsim scores < 0.95 flagged as “unknown” (RelTrans_{thresh}). Classification accuracy reported on unflagged pixels only. Predictions for each classifier are determined by majority vote over ten randomized subsamples.

Classification results for the selected scenarios and classifiers are summarized in Table. 1. We see a dramatic performance increase in all scenarios by simply classifying relation vectors instead of target pixels. This is not surprising, since the spectral shapes of pixels of the same material type in the two images can be smoothed significantly (particularly at longer wavelengths where downsampling from HYDICE to MASTER spectral resolution causes aliasing, see Fig. 2 for examples), but the change in spectral resolution does not significantly alter inter-class relationships within each image. Thus, by characterizing these class relationships within each image, we are able to form a more robust descriptor for inter-image comparisons than the pixels themselves (a similar observation was also made by Rajan et al. in [4]). The prediction rules for MinDist_{rel} and RelTrans_{src} are functionally equivalent and as a result these classifiers behave equivalently. Correcting for spectral differences be-

tween the source and target images using the correspondence points gives another significant performance boost. This is further improved in all three scenarios after thresholding by making a compromise between classification accuracy and the number of labeled samples.

To evaluate capabilities for discovery of “new” classes, we exclude several classes (specifically, classes B, E, M, P, S and k) from the source data, forcing the classifier to choose the best matching class when the “true” class is not present, or assign a label of “unknown” to such samples. These classes were selected since both the source and target images contained pixels belonging to these classes, so we can easily observe the effects on the classification accuracy caused by their exclusion. Table. 2 gives class statistics before and after thresholding for this scenario using the RelTrans_{corr} classifier. Of the 303 pixels marked as unknowns, 227 are from classes not present in the source data. 113 of these pixels are from class P (red tennis court) and another 113 belong to class E (glass). Both of these classes are significantly different in character from the other spectra, and are flagged appropriately as unknowns.

Of the remaining flagged pixels, 34 from class K (green and brown grass) are flagged due to close similarity to class K (Norway and silver maple trees). Class V has trace elements of gray gravel rooftop spectra (along with several asphalt-based materials), and is often confused with class k (containing only gray gravel rooftop spectra). The pairings of class M (also gray gravel rooftops) with class Q (red weathered stained wood), and class S (gray tarp) with class j (brown mixed brick) are rather odd, given their respective material compositions. Nonetheless, these spectra are extremely similar in terms of spectral shapes, even at full HYDICE resolution, and as a consequence will generally receive high relsim scores.

More interesting are the results for classes B and C. The material composition of class B (a class removed from the source data) is a combination of tan asphalt shingles (73.9%) and gray gravel roof (23.6%) spectra. This combination explains the large variance in reflectance values for this class at wavelengths shorter than 0.9 μm (see Fig. 2). Class U consists entirely of brown asphalt roof shingles, and class V is primarily composed of black (25.6%) and gray (73.9%) asphalt surfaces, with trace elements of gravel rooftop materials. Of the 113 target pixels in class B, the RelTrans classifier assigns 84 (74.3%) to source class U, and 28 (24.7%) to source class V, not surprisingly reproducing the true proportions of U and V, with high relsim scores. Class C, which is not excluded from the source data, is small (64 pixel), consisting of heterogeneous material in shadows (specifically, gray and black asphalt roof shingles (53.1%, 1.6%), brown plank wood siding (18.8%), concrete cinder blocks (23.4%), and dark gray BMW Paint (1.6%)). Due to this heterogeneity, these pixels are assigned low relsim scores and are flagged as unknowns accordingly.

5. DISCUSSION AND FUTURE WORK

A problem with working with synthetic data is that it is almost always much cleaner than real image data. In this study, cross-validated classification of the source data and the target data (independently) both yield overall classification accuracies near 99% (test accuracy). While it is certainly possible to achieve such high accuracy when classifying real data, validation of these results on less pristine, real image data is currently in progress.

Another important issue is how we measure confidence in the quality of the spectral correspondences. Since in this work the only difference between the source and target images is spectral resolution, we can assume that the pixels in the overlap region share similar spectral footprints in both images. However with real data, environmental conditions or the underlying scenery itself may have changed

RelTrans _{corr}								RelTrans _{thresh}					
Class	Primary Materials	?	n	PA(%)	EO(%)	EC(%)	CA(%)	?	n	PA(%)	EO(%)	EC(%)	CA(%)
A	Roof Shingle, Asphalt, Brown and Red Blend	0	109	100.0	0.0	50.9	49.1	0	109	100.0	0.0	0.0	100.0
B*	Roof Shingle, Asphalt, Tan (73.9%), Roof, Gravel, Gray (23.6%)	0	113	0.0	100.0	0.0	100.0	1	112	0.0	100.0	0.0	100.0
C	Shadow Materials	0	30	100.0	0.0	0.0	100.0	30	0	100.0	0.0	0.0	100.0
E*	Glass	0	113	0.0	100.0	0.0	100.0	113	0	100.0	0.0	0.0	100.0
J	Tree, Maple, Silver (46.7%), Norway (53.3%)	0	113	100.0	0.0	9.6	90.4	0	113	100.0	0.0	2.6	97.4
K	Grass, Green, Healthy (91.1%), Brown (8.9%)	0	113	89.4	10.6	0.0	100.0	34	79	96.2	3.8	0.0	100.0
M*	Roof, Gravel, Gray (98.9%)	0	92	0.0	100.0	0.0	100.0	0	92	0.0	100.0	0.0	100.0
P*	Tennis court, Playing Surface, Red	0	113	0.0	100.0	0.0	100.0	113	0	100.0	0.0	0.0	100.0
Q	Wood, Stained, Red, Old, Weathered	0	113	100.0	0.0	45.4	54.6	0	113	100.0	0.0	44.9	55.1
R	Roof Shingle, Asphalt, Brown, Black, New (86.5%), Roadway Surfaces, Asphalt, Old, Gray (8.7%)	0	113	100.0	0.0	0.0	100.0	8	105	100.0	0.0	0.0	100.0
S*	Gray Tarp	0	112	0.0	100.0	0.0	100.0	0	112	0.0	100.0	0.0	100.0
U	Roof Shingle, Asphalt, Mix Brown	0	113	100.0	0.0	42.9	57.1	0	113	100.0	0.0	42.6	57.4
V	Roadway Surfaces, Asphalt, Old, Gray (73.9%), Asphalt, Black, New (25.6%)	0	113	97.3	2.7	56.0	44.0	4	109	100.0	0.0	56.2	43.8
Y	Grass, Brown and Green w/Dirt	0	113	100.0	0.0	0.0	100.0	0	113	100.0	0.0	0.0	100.0
a	Roof Shingle, Asphalt, Black, Weathered	0	110	100.0	0.0	0.9	99.1	0	110	100.0	0.0	0.0	100.0
c	Sheet Metal, White, Fair (72.8%), Saturn Hood Paint, White (18.5%)	0	79	100.0	0.0	0.0	100.0	0	79	100.0	0.0	0.0	100.0
d	Roof Shingle, Asphalt, Black	0	113	100.0	0.0	50.0	50.0	0	113	100.0	0.0	0.0	100.0
j	Brick, Siding, Mix Brown, Fair (98.8%)	0	113	100.0	0.0	49.8	50.2	0	113	100.0	0.0	49.8	50.2
k*	Roof, Gravel, Gray	0	112	0.0	100.0	0.0	100.0	0	112	0.0	100.0	0.0	100.0
Totals		0	2000	OVR=66.4%	AVG =67.6%	$\kappa=0.6445$		303	1697	OVR=74.4%	AVG =78.5%	$\kappa=0.7277$	

Table 2: RelTrans class statistics for scenario III before thresholding (left table) and after thresholding (right table). ?=Unknown class counts, n=Labeled class counts, PA=producer’s accuracy, CA=consumer’s accuracy, EO=omission errors, EC=commission errors, OVR=#correct/#samples, AVG=mean producer accuracy, κ =kappa statistic. Classes with an asterisk* were not included in the source (training) data. Green cells indicate excluded classes that were correctly flagged as unknowns, orange cells indicate flagged source classes which were included in source data, and red cells indicate excluded source classes which were not flagged. Classes B, U, V, and C shown in Fig. 2

between image capture times. Thus, an additional filtering step may be necessary to discard/weight correspondences based on how well they match between images. This also involves an assessment of the robustness of the class similarity function in the presence of noisy correspondences, and is a crucial focus of future work.

Since this method allows for the discovery of new classes in the target data, one has to decide how to annotate these new classes. In previous work [13], [14], we employed external spectral libraries to annotate clusters in hyperspectral imagery. One could use a similar approach here as a post-processing step: after we transfer the source labels to the target image, we could compare pixels flagged as unknowns to a database of known material signatures to automatically annotate newly discovered classes.

Acknowledgements: We would like to thank David Pogorzala and Prof. John Kerekes at the RIT Digital Imaging and Remote Sensing (DIRS) Laboratory for their gracious assistance in generating the DIRSIG image data used in this work. We would also like to thank Dr. Maj. Mike Mendenhall and Dr. Bill Farrand for their helpful input on topics related to atmospheric correction.

6. REFERENCES

- [1] BM Shahshahani and DA Landgrebe, “The effect of unlabeled samples in reducing the small sample size problem and mitigating the Hughes phenomenon,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 32, no. 5, pp. 1087–1095, 1994.
- [2] L Bruzzone and DF Prieto, “Unsupervised retraining of a maximum likelihood classifier for the analysis of multitemporal remote sensing images,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 39, no. 2, pp. 456–460, 2001.
- [3] Yang-Chi Chen, “Knowledge-based learning for classification of hyperspectral data,” *Doctoral Dissertation - University of Texas at Austin*, pp. 1–134, 2007.
- [4] S Rajan, J Ghosh, and M Crawford, “Exploiting class hierarchies for knowledge transfer in hyperspectral data,” *IEEE Transactions on Geoscience and Remote Sensing*, Jan 2006.
- [5] JR Schott, SD Brown, RV Raqueno, HN Gross, and G Robinson, “Advanced synthetic image generation models and their application to multi/hyperspectral algorithm development,” *Proceedings of SPIE*, vol. 3584, pp. 211, 1999.
- [6] Chein-I Chang, “An information-theoretic approach to spectral variability, similarity, and discrimination for hyperspectral image analysis,” *IEEE Trans. Info. Theory*, vol. 46, no. 5, pp. 1927–1932, 2000.
- [7] C Salvaggio, L Smith, and E Antoine, “Megacollect 2004: hyperspectral collection experiment of terrestrial targets and backgrounds of the RIT megascene and surrounding area (Rochester, New York),” *Proceedings of SPIE*, 2005.
- [8] R Basedow, D Carmer, and M Anderson, “HYDICE system: implementation and performance,” *spie.org*, vol. 2480, pp. 258–267, 1995.
- [9] Research Systems Inc, *ENVI 4.6 Users Guide*, 2008, 1196 pp.
- [10] E Merényi, K Tasdemir, and W Farrand, “Intelligent information extraction to aid science decision making in autonomous space exploration,” *spie.org*, vol. 6960, 2009.
- [11] Simon J Hook, Jeffrey Myers, Kurtis Thome, Michael Fitzgerald, and Anne Kahle, “The MODIS/ASTER airborne simulator (MASTER) - a new instrument for earth science studies,” *Remote Sensing of Environment*, vol. 76, no. 93, pp. 102, 2001.
- [12] T Cocks, R Jenssen, A Stewart, I Wilson, and T Shields, “The HyMap airborne hyperspectral sensor: the system, calibration and performance,” *Proceedings of the 1st EARSeL workshop on Imaging Spectroscopy*, pp. 37–42, 1998.
- [13] B Bue, B Csathó, and E Merényi, “Automated labeling of segmented hyperspectral imagery via spectral matching,” *Proceedings IEEE WHISPERS 2009*, Aug 2009.
- [14] B. Bue, E. Merényi, and B. Csathó, “Automatic identification of materials in hyperspectral imagery,” *IEEE Transactions on Geoscience and Remote Sensing*, 2009, (submitted).