



HAL
open science

Unsupervised clustering of hyperspectral images of brain tissues by hierarchical non-negative matrix factorization

Bangalore Ravi Kiran, Bogdan Stanciulescu, Jesus Angulo

► To cite this version:

Bangalore Ravi Kiran, Bogdan Stanciulescu, Jesus Angulo. Unsupervised clustering of hyperspectral images of brain tissues by hierarchical non-negative matrix factorization. *BIOIMAGING 2016*, Feb 2016, Rome, Italy. pp.8, 10.5220/0005697600770084 . hal-01280453

HAL Id: hal-01280453

<https://hal.science/hal-01280453>

Submitted on 29 Feb 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.

Unsupervised clustering of hyperspectral images of brain tissues by hierarchical non-negative matrix factorization

Bangalore Ravi Kiran¹, Bogdan Stanciulescu¹ and Jesús Angulo²

¹ *Centre de Robotique(CAOR), MINES ParisTech, PSL-Research University, Paris, France*

² *Centre de Morphologie Mathématique(CMM), MINES ParisTech, PSL-Research University, Fontainebleau, France*
{ravi.kiran, bogdan.stanciulescu, jesus.angulo}@mines-paristech.fr

Keywords: Hyperspectral Image, Hierarchical clustering, Brain tissue, Non-negative matrix factorization.

Abstract: Hyperspectral images of high spatial and spectral resolutions are employed to perform the challenging task of brain tissue characterization and subsequent segmentation for visualization of in-vivo images. Each pixel is a high-dimensional spectrum. Working on the hypothesis of pure-pixels on account of high spectral resolution, we perform unsupervised clustering by hierarchical non-negative matrix factorization to identify the pure-pixel spectral signatures of blood, brain tissues, tumor and other materials. This subspace clustering was further used to train a random forest for subsequent classification of test set images constituent of in-vivo and ex-vivo images. Unsupervised hierarchical clustering helps visualize tissue structure in in-vivo test images and provides a inter-operative tool for surgeons. Furthermore the study also provides a preliminary study of the classification and sources of errors in the classification process.

1 Introduction

Medical Hyperspectral imaging (MHSI) is a non-invasive and non-ionizing modality used in robust tissue identification, characterization achieved reasonable classification and localization accuracy in of various kinds of cancerous tissues on account of vast improvements in spectral and spatial imaging resolutions (LF14). Discriminating cancerous and normal tissues with Near Infra-Red (NIR) spectroscopy has been an active topic in the past two decades. (PYF⁺07) study the use of MHSI for identifying residual tumor in a resection bed and to indicate regions requiring more resection. (LHW⁺14) use HSI in-vivo to achieve tumor identification in tumor bearing mice. (LWL11) study localization of tumor tissue on the tongue of human patients using a tunable HSI camera. In this work we study the characterization and identification of brain tissues by performing joint clustering and un-mixing. We also demonstrate the utility of the hierarchical clustering as a visual aid for inter-operative(IoP) procedures (SC03) in the resection of cancerous tissue. Though one should note that the standard in IoP brain tumor surgery modality has been MRI, it is precise but a costly technology. Hyperspectral images are cheaper and furthermore faster to acquire.

The primary goal of this paper is to address the

common problem of lack of an exact ground truth labels in the hyperspectral images in-vivo. This produces a challenging problem of hyperspectral image segmentation and classification. Our solution consists in generating unsupervised clustering that can serve as a stable identification and labeling, for a subsequent semi-supervised classification. Further we note that, it is unrealistic to obtain a well localized marker of tumors, in-vivo, during a surgical procedure, in view of different constraints of space and time. Finally we also do not know the exact spectrum of cancerous and normal tissue at this stage of our study. We do note that from previous studies on different tumor types, normal and cancerous tissue spectra are close (see subsection 4), and are difficult to distinguish.

We thus characterize the tissue by clustering their spectrum in similar subspaces. This helps visualize different structures in the tissue as well as locate tumors and related structures. Unsupervised clustering is rarely followed, for such problems, but given that we operate at high spatial and spectral resolutions, we stand to gain some useful information. We use the clustering of ex-vivo tumors with surrounding tissue, blood and other structures as a target model we predict the location of tumors in in-vivo images. Finally, the clustering will also be used as tool by surgeons and pathology to label the tissues to generate ground truth data.

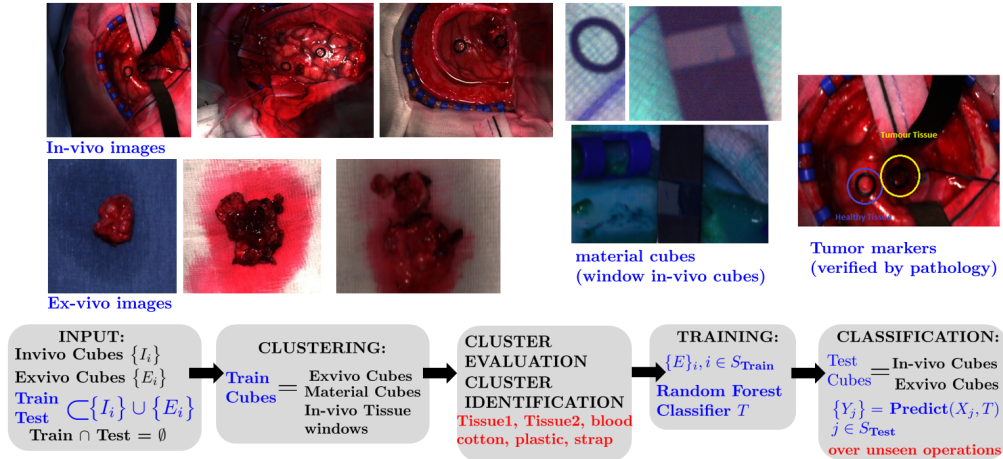


Figure 1: We classify the input images into three categories: In-vivo images (during surgery), Ex-vivo images (extracted tumor after surgery) and material cubes (carefully windowed regions containing only surgical material). Black plastic rings markers highlight the locations of healthy tissue and a tumor tissue in a scene. The ex-vivo images serve as localized targets that ensure the presence of tumor tissue in a localized window. We partition the operations into a training and test set, after which we perform hierarchical clustering on the test images. The H2NMF clustering algorithm also calculates the endmembers for each cluster. We use clusters at a particular level to train a random forest classifier. We show the classifier prediction on other ex-vivo and in-vivo images in the test set. The flow-diagram at the bottom illustrates all the steps in our study.

The primary contributions of the paper include the identification of pure-spectrum or pure-pixels of materials using VNIR camera images of brain tumors by applying the hierarchical clustering by 2-rank NMF approximation (GKP15). Pure-pixels are found at the leaves of this hierarchical structure and correspond to rank-1 matrices. This is discussed in section 2. We use the unsupervised clustering as supervised labels to identify these pure-pixels in the in-vivo images by training a decision forest. See figure 1 for a summary of the work-flow. The organization of the paper is as follows. We describe briefly the imaging setup in this section. A brief introduction to endmember extraction is provided in section 2, the separable non-negative matrix factorization(NMF) and the hierarchical NMF decomposition algorithm (GKP15) employed in this study.

Imaging Setup: The medical imaging demonstrator consists of two camera pairs covering different parts of the spectrum at different spatial resolutions, a VNIR (visible-to-near-infrared) and NIR (near-infrared cameras) cameras. The VNIR camera captures 826 spectral bands, in the spectral range of 400-1000 nm with a spectral resolution of 2-3 nm, and each pixel has a dimension of 0.1287 mm x 0.1287 mm. While the NIR camera captures 172 spectral bands, in the spectral range of 900-1700 nm, each pixel has a dimension of 0.48 mm x 0.48 mm. We shall work with the VNIR camera images mainly since they consist of precise spectral information. The increment in wavelength is 1.9285 nm for the current

setup. Images are calibrated using a standard white and dark reference.

In this study we work on hyperspectral images from a VNIR camera with a very high spatial and spectral resolution. Images were obtained of the brain tissue and brain tumors, in-vivo (brain tissue open during surgery) and ex-vivo (extracted brain tumor on surgical table). Images were obtained over 20 operations. The images from both cameras were obtained during the surgical procedure. The information of the type of tumor and other associated patient information was missing during the current paper’s study. This will be discussed more in detail. The images were taken during the surgical procedure (in-vivo) and right after the procedure outside the body (ex-vivo) with the extracted tumor placed under the camera. In this paper we restrict our study to the VNIR camera images due to their high spectral and spatial resolution.

2 Endmember Extraction

Hyperspectral unmixing is the problem of inferring which are the materials in the scene (visible directly or visible as a mixture) by the general process of dimensionality reduction. This basically consists in decomposing a hyperspectral image $X = [x_1, x_2, \dots, x_n]$ with n pixels, where each pixel x_i is p -dimensional spectrum. The two elements of the decomposition are a set of r reference spectra or endmembers and $W = [w_1, w_2, \dots, w_r]$ and their relative proportions at

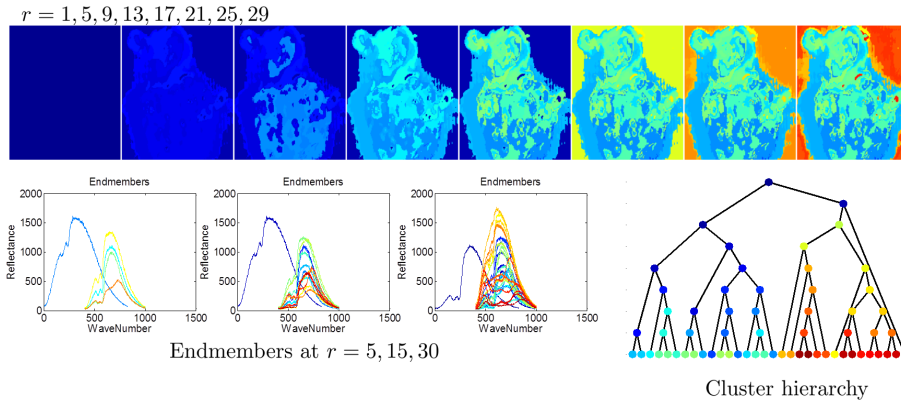


Figure 2: Segmentation Hierarchy showing a tumor ex-vivo decomposed at different levels $r = 1$ to 32 of the hierarchical rank-2 decomposition (H2NMF) algorithm. The 2-rank approximation divisional step is used to arrive with a unit rank matrices and the leaves of the hierarchical decomposition. The labels in the segmentations are consistently colored to match with the tree nodes. They help show the persistence of a subspace/clustering visually. This tree structure gives a visual idea of the structure of the tissue, labels that are close in terms subspaces branch. By viewing the endmembers (spectra) at different levels, we can infer visually that blood is a dominant material in the scene (blue peak around 400nm). The shallow clusters correspond to submatrices of low approximation error which close to rank-1 and thus correspond to a single tissue/material.

each pixel(abundances) $H = [h_1, h_2, \dots, h_n]$, where, w_i is a p -dimensional spectrum and h_i is r -dimensional coefficient weights. The spectrum at a pixel is a convex composition of endmembers. This imposes two conditions: i) the abundances need to be positive, ii) the sum of abundances by each endmember sum to unity at each pixel (MBDC⁺14). The goal of the hyperspectral unmixing(HU) or endmember extraction process is to obtain the spectrum that corresponds to “pure” materials in the scene. This can be a complex problem based on the physical interactions between the materials, rays and the reflections. We assume the Linear mixing model(LMM) for its inherent simplicity and popularity in the remote sensing community. The model assumes that at a given pixel location the incident ray interacts only with one material in the true site of the pixel. Given the spectral and spatial resolution of the VNIR camera, this assumption seems to be a reasonable.

Separable NMF: Endmember extraction is well studied as a problem of non-negative matrix factorization problem since it respects the endmember constraints. It is written as a matrix factorization problem $X \approx WH$, where this approximation is often evaluated by minimizing the Frobenius norm of the residual $\|X - WH\|_F^2$. Independent of how the decomposition is calculated, geometrically the NMF generates the vectors h_i (endmembers) whose simplicial cone contains all the data points $\{x_i\} \in X$. There can be many cones possible and the decomposition is non-unique. In case of high spatial and spectral resolution of HSI cubes, one can make the hypothesis that there exists pixels in the image which themselves form the

bounding vertices of the a convex cone, and thus directly the endmembers. Physically this means that given a pixel spectrum we can infer the material it contains and be sure that its contribution is from a single material. This is called the separability condition (GL14) under which we obtain a unique and tractable decomposition of the NMF problem. Geometrically, the separability condition reformulates the NMF problem as that of finding the extreme rays of the conical hull of a finite set of vectors. Given a HSI cube $X \in \mathbb{R}^{n \times p}$, we can write its separable NMF decomposition as

$$X = WH = W[I_r, H']\Pi \quad (1)$$

where $W \in \mathbb{R}^{m \times r}$, $H' > 0$, Π is a permutation matrix, which rotates the r different endmembers. Columns of W are the endmembers, while the entry $H(i, j)$ are the abundances/scalar coefficient of the i th endmember at the j th pixel in X . Geometrically this corresponds to the convex hull of the columns of X .

Hierarchical Two-rank NMF: The hierarchical two rank NMF (H2NMF) first introduced in the paper (GKP15) aims at obtaining a rank-1 sub-matrix by hierarchical application of a NMF decomposition. The decomposition step consists of: i) projecting the data points on to the first two eigen vectors, and thus a 2-d cone, by calculating the singular value decomposition(SVD). ii) the direction matrix is approximated by a separable NMF decompositions, namely the SPA (GL14) of the data points under decomposition.

This yields pure-pixels in the HSI cube approximating the first principle component. The goal is to finally arrive with a set of leaf sub-matrices of rank

one. This corresponds to pure-pixel or a single material. Though this hierarchical method is stated, the stopping condition for the growth of the hierarchy is not provided. Instead this is a parameter of the algorithm that determines the number of endmembers being extracted. At each step of this divisional hierarchical clustering step, the approximation error minimized is (GKP15):

$$E_k = \sigma_1^2(X(:, K_i^1)) + \sigma_1^2(X(:, K_i^2)) - \sigma_1^2(X(:, K_i)) \quad (2)$$

where σ_1 refers to the first singular value of the matrices, and K_i corresponds to the indices of the parent and K_i^1, K_i^2 are indices of the child clusters in the binary split. This approximation error corresponds to the rank approximation of input cluster or matrix.

The split is decided by choosing line after the projection on to a 2-d cone. This is chosen by parameter δ^* (see algorithm 1) so as to have a balanced set of points across the split (GKP15) as well a clusters with elements that are homogeneous.

Cluster Endmember: The H2NMF algorithm associates an endmember with each cluster. The endmember of a cluster is the spectrum best approximating the first eigen value of submatrix, we follow (GKP15)'s measure. Mean removed spectral angle (MRSA) is defined for any two pairs of spectra \mathbf{x}, \mathbf{y} as:

$$\phi(\mathbf{x}, \mathbf{y}) = \frac{1}{\pi} \arccos \left(\frac{(x - \bar{x})^T (y - \bar{y})}{\|x - \bar{x}\|^2 \|y - \bar{y}\|^2} \right) \in [0, 1] \quad (3)$$

where for a vector $z \in \mathbb{R}^m, \bar{z} = \sum_{i=1}^m z_i$. The MRSA provides the a measure to compare the pure-pixel in the input sub-matrix and its endmember spectrum produced by a rank-2 NMF decomposition. We use MRSA to evaluate the homogeneity of the hierarchical clustering at different levels.

We use the H2NMF algorithm (see Algo. 1) to provide a hierarchical subspace decomposition of the ex-vivo cubes to recover the pure-pixels by varying the level of decomposition. The decomposition is robust to noise and calculates quickly, given the huge sizes of the datasets (X 's size is several 100000 pixels \times 800 wavelengths). Secondly we visualize the hierarchical clustering structure which provides us a way to localize different tissue based on their paths in the hierarchy. In figure 2 we can observe from the clustering and the endmembers that, blood (endmember spectrum in blue, peaking around 400 nm for $r = 4, 8, 12$) at various levels of the hierarchy. Continuing further $r > 8$ we see that the tissues themselves are seen to have different layers and structuring of these layers.

3 Experiments and analysis

We partition the surgeries (numbered 1-20) into two disjoint sets: first the training set over which we perform unsupervised clustering and subsequent classifier training, second the test set over which we predict the cluster labels. We predict both on in-vivo and ex-vivo images from different operations in the test set to evaluate the validity of the different cluster labels. The training set consists of ex-vivo images, material cubes and certain windows from In-vivo images that contain synthetic surgical materials such as cotton, plastic cables and straps, as demonstrated in figure 1. The training set also contains supplementary tissue information such skull and surrounding tissues. Ideally this process captures the dictionary of materials to be clustered in the same subspace.

Training set : The training set is a union of pixels from images of the tumor ex-vivo and dictionary of materials such as plastic and cotton, from the training subset of operations. We denote the training images by $X_{\text{train}} \in \mathbb{R}^{N \times p}$, where N are the total number of pixels in the training set images together, and $p = 826$ is the dimensionality of the spectra at each pixel. The ex-vivo and material cubes in the test set will be clustered together so as to separate the different pure-pixels corresponding to different materials. The training set does not contain any in-vivo images since images containing the target tumor samples shall remain unlabelled and not part of a class.

Test set : The test set consists of both types of images: Images ex-vivo of the tumor, and images in-vivo from operations not in the training set. We refer to the different test cubes/images as X_i and their corresponding classification map by the random forest as Y_i .

We use algorithm (1) described in section 2 to perform hierarchical clustering of the training-images. The principal goal here is to visualize the tissue structure as decomposed by the hierarchical clustering and aid a surgeon or expert in annotating these images. Another important goal is to determine the depth of the hierarchy at which we are able to separate the tumor, blood, tissues and materials and extract their endmembers or spectral signatures.

3.1 Cluster Evaluation

In this section we compare the performance of H2NMF, the clustering method used in our study, with respect to other clustering algorithms: hierarchical K-means(HKM), hierarchical spherical K-means(HSKM). For the cases of NMF when used as a clustering method, one does not explicitly use a distance function or dissimilarity (Kua14). The pre-

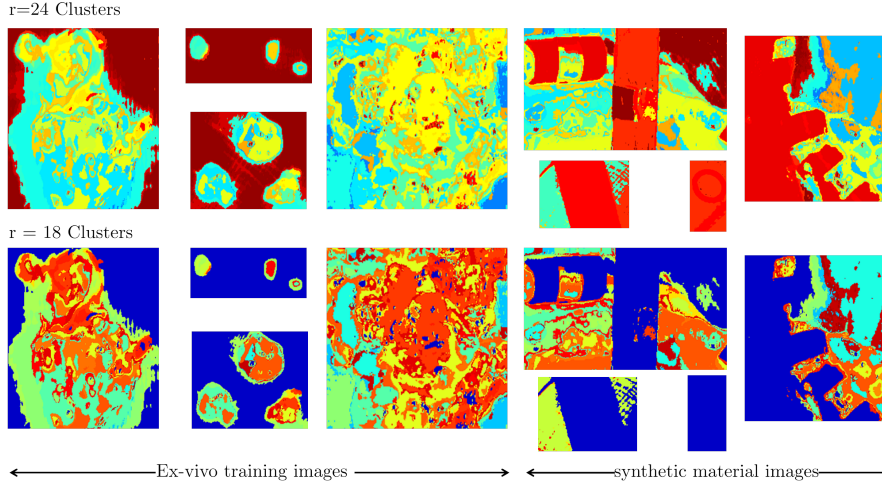


Figure 3: H2NMF clustering on the training set X_{train} which consists of input ex-vivo tumor tissue HSI cube and on materials. We have in this image the ring on a cotton, and a cable in scene with open cranium and healthy tissues. This labeling with the corresponding HSI cubes serves as the training set for constructing the random forest classifier.

Algorithm 1 H2NMF (GKP15)

- 1: **procedure** H2NMF($X \in \mathbb{R}_+^{p \times n}$)
 - 2: $\mathcal{K}_1 \leftarrow \{1, 2, \dots, n\}$ and $\mathcal{K}_i \leftarrow \emptyset$ for $2 \leq i \leq r$
 - 3: $(\mathcal{K}_1^1, \mathcal{K}_1^2) = \text{2-rank-Split}(X, \mathcal{K}_1)$
 - 4: **while** $k < r$ **do**
 - ▷ Iterate until r-clusters are reached
 - 5: $j = \arg \max_{i=\{1,2,\dots,r\}} E_k$ from eq(2)
 - ▷ pick cluster j with largest error
 - 6: $\mathcal{K}_j = \mathcal{K}_j^1, \mathcal{K}_k = \mathcal{K}_j^2$
 - ▷ update cluster j into j and k
 - 7: $(\mathcal{K}_j^1, \mathcal{K}_j^2) = \text{splitting}(X, \mathcal{K}_j)$
 - 8: $(\mathcal{K}_k^1, \mathcal{K}_k^2) = \text{splitting}(X, \mathcal{K}_k)$
 - 9: **return** \mathcal{K}_i ▷ cluster indices forming hierarchy
 - 10: **procedure** SPLITTING($X \in \mathbb{R}_+^{p \times n}, \mathcal{K} \subset \{1, 2, \dots, n\}$)
 - 11: $[W, H] = \text{rank2NMF}(X(:, \mathcal{K}))$
 - ▷ Projection onto 2-d cone
 - 12: $x(i) = \frac{H(1,i)}{H(1,i)+H(2,i)}$
 - 13: Compute split parameter δ^* ▷ parameter trades-off balanced split and cluster homogeneity
 - 14: $\mathcal{K}^1 = \{\mathcal{K}(i) | x(i) \geq \delta^*\}$
 - 15: $\mathcal{K}^2 = \{\mathcal{K}(i) | x(i) < \delta^*\}$
 - 16: **return** $\mathcal{K}^1, \mathcal{K}^2$ ▷ return 2-sets of indices corresponding to two sub-matrices of the input-matrix X .
-

diction strength (TW05) is a well suited measure for cluster validation for NMF-based clustering because it does not rely on a distance function. Here we introduce a method similar to the one described in (Kua14) to evaluate the NMF. Given matrix X we split them into $X_{\text{train}}, X_{\text{test}}$ over which we obtain NMF decompo-

sitions, giving us

$$\begin{aligned} \arg \min \|X_{\text{train}} - W_{\text{train}}H_{\text{train}}\|_F^2 \\ \arg \min \|X_{\text{test}} - W_{\text{test}}H_{\text{test}}\|_F^2 \end{aligned} \quad (4)$$

Now the prediction error can be obtained by solving the non-negative least-squares problem for

$$H_{\text{predict}} = \arg \min \|X_{\text{test}} - W_{\text{train}}H\|_F^2 \quad (5)$$

We then compare the maximum abundance values in $H_{\text{test}}, H_{\text{predict}}$ to check which endmember index contributed the most, and use this as a measure of prediction strength. This is because the H2NMF algorithm provides decomposition with one dominant endmember. This is a variant on the classical prediction strength measure (TW05). We demonstrate the results for the three hierarchical clustering algorithms. The prediction strength provides us a way to calculate the optimal number of clusters. Figure 4 demonstrates the prediction strength calculated over random samples of the training set X_{train} . This measure could also provide a reasonable guide in determining which clusters to split, aside the approximation error. We plot the maximum value of MRSA from equation (3) for different number of clusters in in figure 4. This demonstrates how far the pure-pixel in cluster in a given level of hierarchy from its rank-2 NMF approximation. The figure also shows the plot of the average approximation error in eq. 2 over different leaf sub-matrices for the three hierarchical clustering algorithms. Both of these measures provide us a tool to asses if the spectrum and sub-matrix at a leaf are well approximated by a pure-pixel. Finally, we evaluate the hierarchical clusters intrinsically by using the

gap statistic (TW05) which using the average pairwise dissimilarity per cluster to evaluate the number of clusters graphically, by looking for an elbow.

We now train the resulting clustering on the training set X_{train} . We use the Matlab's implementation of the random forest to train the Random forest(RF) with the weak learner set to axis-aligned-hyperplane. We first cut the subspace hierarchy by choosing the number of clusters r , and then use the clusters as labels to train the RF-classifier.

The predictions of the RF-classifier are shown in figure 5 when trained on two different levels of clustering. We remark here that the actual subspace hierarchy obtained by a sequence rank-2 NMF decompositions, is not completely approximated by the RF-classifier. Nevertheless we do see a reasonable segmentation. The rings shown in the in-vivo images correspond to physical markers that were placed and localize healthy and tumor tissues for evaluation purposes. We do not obtain a single cluster label that might correspond to tumor. To improve the classification it would be ideal to integrate the structure of the subspace hierarchy. This can be done by learning a random forest for each level in the hierarchy. This would of course be costly. A more principled solution would be to minimize an appropriate loss function that approximates the H2NMF subspace hierarchy.

The resulting segmentation from the classifier provides firstly the stable pure signatures of different materials in the ex-vivo tumors. The results are verified basically by checking the presence of unique label (or sets of labels) in a ring (1)) in the in-vivo images. The tissue samples within these rings have been verified posteriori by pathology. One ring is assured to surround a cancerous tissue sample, while the other a healthy one. In this setup there was estimate of the depth cancerous and its uniform exposure to the VNIR-camera. This is tough to ensure during critical surgeries. The color-map of the classification result in figure 5 correspond to the color-map of the unsupervised clustering in figure 3. The semantic meaning of the labels is temporary and the color-map further requires verification by a surgeon or pathology expert, to which finally a more sensible name shall be attached. But it is evident that we are able to identify plastic surgical materials, cotton.

4 Issues

Spectral proximity between tissues : (SMP⁺¹³) study differentiation between benign and malignant tumors in canine mammary tumors, and detect tu-

mors using a imaging spectrometer. Ovarian tissue characterization and tumor differentiation (UBS⁺⁰¹), were performed using NIR spectroscopy. The results showed that both slope and intensity of reflectance spectra may have the ability to discriminate normal from abnormal conditions of the ovary. Slopes between 510 nm and 530 nm were most discriminatory for ovaries, while 630 nm to 900 nm in the study on kidneys (Pes07). The differences between malignant and benign tissues in human breast tissues, have been attributed to the metabolic difference due to the presence of more oxy-hemoglobin, lipids and water. These elements were used as the spectral markers considered for the additive construction of the spectrum. (KCT⁺¹⁰). This proximity in the spectral response of normal, benign and malignant tumor plays an crucial role in classification of tumors pixels. In our study, due the lack of information on the actual spectrum of tumor tissues we are unable to provide a good division of the subspace and the subsequent refinement of decision boundary between normal and possible cancerous tissue.

Specular reflections : Specular reflections are a prime problem in all our clustering and classification steps and contributes to mis-classification. Specular reflection occurs when the light source can be seen as a direct reflection on the surface of the tissue under study. The pixels affected by specular reflection contain mainly the light source and may or may not contain sufficient information on the reflectance of the tissue under study. These reflected rays may also be suspect of undergoing multiple reflections. An example is shown in figure 6.

5 Conclusion

In this paper we applied and studied the hierarchical rank-2 NMF clustering and endmember extraction to segment brain tissue structure in an unsupervised setting. Subsequently we have tried to train a random forest to learn the structure of the subspace clusters.

Subspace learning and Spatial features : Though the classifier can't discriminate between normal tissues and tumor spectra, it provides a good segmentation of the tissues and surgical materials in the scene. The H2NMF hierarchy provides a hierarchy of low rank sub-matrices that contain pure-pixels, and in this case correspond to different brain tissues. This cluster hierarchy can further be used by a surgeon in loop to search more pertinent subspaces and feed the results back to the classifier. We have seen the effect

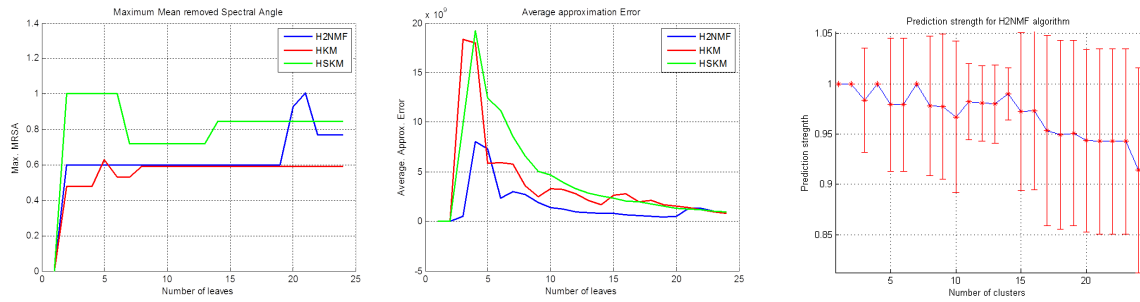


Figure 4: Left: the variation of the maximum MRSa (mean-removed spectral angle) to evaluate the homogeneity of clusters in terms of endmember signature. This is evaluated across the three hierarchical clustering algorithms: H2NMF, HSKM, HKM. We observe that all the three algorithms provide satisfactory results. Center: figure also shows the variation approximation error in equation (2) with increasing leaves for three hierarchical clustering algorithms. Right: Plot demonstrating the mean prediction strength calculated using the NMF decomposition. The red bars show the variance of the prediction strength. Optimal clusters are chosen based on the maximum of the prediction strength. Here for example it is 14.

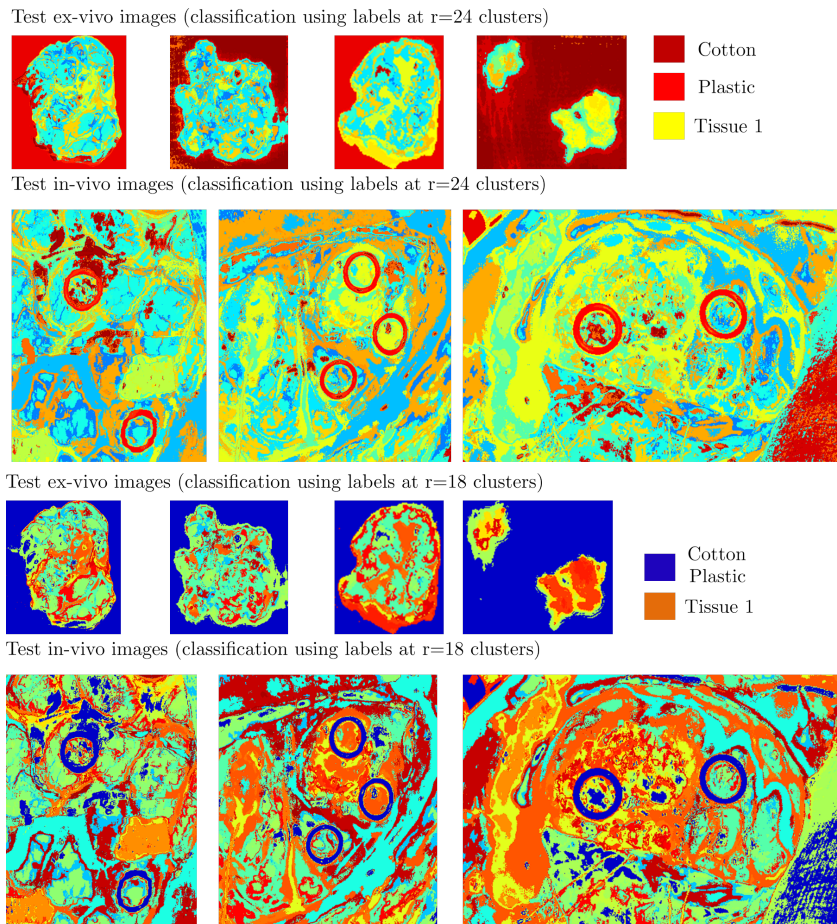


Figure 5: Results of the classification on the test set which consists of Ex-vivo images and In-vivo images test-set. Two sets of classification maps are shown, the top image tile corresponding to learning the labels by cutting the subspace tree at $r = 18$ clusters and the bottom tile at $r = 24$. We observe that there are different tissue structures and materials found in the original clustering in figure 3. It would be ideal to have a cluster/endmember corresponding to the tumor, though we remark that it is not obvious to extract them or check for their existence. We can also notice that there are sets of labels co-occurring in the marker rings though no trivial or consistent co-occurrence that can ensure a tumor detection. We can also notice in the figure that certain materials have been approximated by the same subspace cluster: the cotton and the plastic. This can also be visually confirmed by look at the subspace tree. Discriminating materials in this subspace tree structure is a crucial analysis tool. As already mentioned, multiple classifiers are needed to approximate well the subspace clustering hierarchy. This would merge the two levels of classification in this image into a more precise classification map.




Figure 6: An in-vivo image and its classification map with the specular reflection marked in dotted red ellipse. This phenomena is not restricted to the optical range.

of learning classifiers on different levels of the hierarchical clustering. Though this produces reasonable segmentations, it does not approximate the subspace hierarchy exactly. We will study how to define loss functions to encode this structure. Given that the difference in spectrum between tumor tissues and normal tissues are small, and prone to noise and variations across patients, building a robust spatial structure descriptor is important. One of the results of our study points to the fact that the tissue structures surrounding a tumor is a key feature. There is already evidence that in the micro-scale these normal and cancerous tissues have a different topological arrangement (Dvo03). This structural information of tumor is useful to obtain a better detector for cancerous tissues in hyperspectral images, since spectrum alone is not sufficient to classify them robustly. Spatio-spectral features have been used in (LHW⁺14), though our aim in the future is to use the spatial features extracted by scattering-transform (BM13) on the different abundance maps extracted by the hierarchical clustering algorithm. This enables us to perform a principled search for features across various scales.

Acknowledgment : This work has been supported in part by the European Commission through the FP7 FET Open programme ICT-2011.9.2, by the European Project HELICoiD “HypErspectral Imaging Cancer Detection” under Grant Agreement 618080.

REFERENCES

- Joan Bruna and Stephane Mallat. Invariant scattering convolution networks. *IEEE Trans. Pattern Anal. Mach. Intell.*, 35(8):1872–1886, Aug 2013.
- Harold F Dvorak. How tumors make bad blood vessels and stroma. *The American journal of pathology*, 162(6):1747–1757, 2003.
- Nicolas Gillis, Da Kuang, and Haesun Park. Hierarchical clustering of hyperspectral images using rank-two nonnegative matrix factorization. *IEEE T. Geoscience and Remote Sensing*, 53(4):2066–2078, 2015.
- Nicolas Gillis and Robert Luce. Robust near-separable nonnegative matrix factorization using linear optimization. *Journal of Machine Learning Research*, 15:1249–1280, 2014.
- Shwayta Kukreti, Albert E Cerussi, Wendy Tanamai, David Hsiang, Bruce J Tromberg, and Enrico Gratton. Characterization of metabolic differences between benign and malignant tumors: High-spectral-resolution diffuse optical spectroscopy. *Radiology*, 254(1):277–284, 2010.
- Da Kuang. *Nonnegative matrix factorization for clustering*. PhD thesis, Georgia Institute of Technology, 2014.
- Guolan Lu and Baowei Fei. Medical hyperspectral imaging: a review. *Journal of biomedical optics*, 19(1):010901–010901, 2014.
- Guolan Lu, Luma Halig, Dongsheng Wang, Xulei Qin, Zhuo Georgia Chen, and Baowei Fei. Spectral-spatial classification for noninvasive cancer detection using hyperspectral imaging. *Journal of biomedical optics*, 19(10):106004–106004, 2014.
- Zhi Liu, Hongjun Wang, and Qingli Li. Tongue tumor detection in medical hyperspectral images. *Sensors*, 12(1):162–174, 2011.
- W.-K. Ma, J.M. Bioucas-Dias, Tsung-Han Chan, N. Gillis, P. Gader, A.J. Plaza, A. Ambikapathi, and Chong-Yung Chi. A signal processing perspective on hyperspectral unmixing: Insights from remote sensing. *Signal Processing Magazine, IEEE*, 31(1):67–81, Jan 2014.
- Disha Laxman Peswani. *Detection of Positive Cancer Margins Intra-operatively During Nephrectomy and Prostatectomy Using Optical Reflectance Spectroscopy*. ProQuest, 2007.
- SV Panasyuk, S Yang, DV Faller, D Ngo, RA Lew, JE Freeman, and AE. Rogers. Medical hyperspectral imaging to facilitate residual tumor identification during surgery. *Cancer Biol Ther*, pages 439–46, March 2007.
- Michael Schulder and Peter W Carmel. Intraoperative magnetic resonance imaging: impact on brain tumor surgery. *Cancer Control*, 10(2):115–124, 2003.
- Amrita Sahu, Cushla McGoverin, Nancy Pleshko, Karin Sorenmo, and Chang-Hee Won. Hyperspectral imaging system to discern malignant and benign canine mammary tumors. In *SPIE Defense, Security, and Sensing*, pages 87190W–87190W. International Society for Optics and Photonics, 2013.
- Robert Tibshirani and Guenther Walther. Cluster validation by prediction strength. *Journal of Computational and Graphical Statistics*, 14(3):511–528, 2005.
- Urs Utzinger, Molly Brewer, Elvio Silva, David Gershenson, Robert C Blast, Michele Follen, and Rebecca Richards-Kortum. Reflectance spectroscopy for in vivo characterization of ovarian tissue. *Lasers in surgery and medicine*, 28(1):56–66, 2001.