A Hierarchical Unsupervised Spectral Clustering Scheme for Detection of Prostate Cancer from Magnetic Resonance Spectroscopy (MRS)

Pallavi Tiwari[1], Anant Madabhushi[1], and Mark Rosen[2]

[1] Department of Biomedical Engineering, Rutgers University, USA*
anantm@rci.rutgers.edu
[2] Department of Surgical Pathology, University of Pennsylvania, USA
Mark.Rosen@uphs.upenn.edu

Abstract. Magnetic Resonance Spectroscopy (MRS) along with MRI has emerged as a promising tool in diagnosis and potentially screening for prostate cancer. Surprisingly little work, however, has been done in the area of automated quantitative analysis of MRS data for identifying likely cancerous areas in the prostate. In this paper we present a novel approach that integrates a manifold learning scheme (spectral clustering) with an unsupervised hierarchical clustering algorithm to identify spectra corresponding to cancer on prostate MRS. Ground truth location for cancer on prostate was determined from the sextant location and maximum size of cancer available from the ACRIN database, from where a total of 14 MRS studies were obtained. The high dimensional information in the MR spectra is non linearly transformed to a low dimensional embedding space and via repeated clustering of the voxels in this space, non informative spectra are eliminated and only informative spectra retained. Our scheme successfully identified MRS cancer voxels with sensitivity of 77.8%, false positive rate of 28.92%, and false negative rate of 20.88% on a total of 14 prostate MRS studies. Qualitative results seem to suggest that our method has higher specificity compared to a popular scheme, z-score, routinely used for analysis of MRS data.

1 Introduction

Prostatic adenocarcinoma (CAP) is the second leading cause of cancer related deaths in America, with an estimated 220,000 new cases every year (Source: *American Cancer Society*). The current standard for detection of prostate cancer is transrectal ultrasound (TRUS) guided symmetrical needle biopsy which has a high false negative rate associated with it [1]. Over the past few years, Magnetic Resonance Spectroscopic Imaging (MRSI) has emerged as a useful complement to structural MR imaging for potential screening of prostate cancer. Kurhanewicz

* This work was supported by grants from the Coulter foundation, Busch Biomedical Award, Cancer Institute of New Jersey, New Jersey Commission on Cancer Research, and the National Cancer Institute (R21CA127186-01, R03CA128081-01).

N. Ayache, S. Ourselin, A. Maeder (Eds.): MICCAI 2007, Part II, LNCS 4792, pp. 278–286, 2007.

et al. [2] have suggested that the integration of MRS and MRI could potentially improve specificity and sensitivity for screening of CAP, compared to what might be obtainable from any individual modality. MRS is a non-invasive analytical technique for measuring the chemical content of living tissues which is used to detect changes in the concentrations of specific molecular markers in the prostate, such as citrate, creatine, and choline. Variations in the concentrations of these substances can detect the presence of CAP.

Most automated analysis work for MRS for cancer detection has focused on developing fitting techniques that yield peak areas or relative metabolic concentrations of different metabolites like choline, creatine and citrate as accurately as possible. The automated peak finding algorithms suffer from problems associated with the noisy data which worsens when a large baseline is present along with low signal to noise ratio. Mcnight, et al. [3] have looked at z-score (ratio of difference between population mean and individual score to the population standard deviation) analysis as an automated technique for quantitative assessment of 3D MRSI data for glioma. A predefined threshold value of the z score is used to classify spectra in two classes: malignant and benign. Kurhanewicz, et al. [4] have worked on the quantification of prostate MRSI by model based time fitting and frequency domain analysis. Some researchers have applied linear dimensionality reduction methods such as independent component analysis (ICA), principal component analysis (PCA) in conjunction with classifiers [5] [6] to separate different tissue classes from brain MRS. However, we have previously demonstrated that due to inherent non linearity in high dimensional biomedical studies, linear reduction methods are limited for purposes of classification [7].

In this paper we present a novel automated approach for identification of cancer spectra on prostate MRS via the use of manifold learning and hierarchical clustering. Figure 1 illustrates the modules and the pathways comprising our automated quantitative analysis system for identifying cancer spectra on prostate MRS. In the first step, a manifold learning [8] scheme (spectral learning or graph embedding) is applied to embed the spectral data in a low dimensional space so that objects that are adjacent in the high dimensional ambient space are mapped to nearby points in the output embedding. Hierarchical unsupervised k-means clustering is applied to distinguish non-informative (zero-padded spectra and spectra lying outside the prostate) from informative spectra (within prostate). The objects in the dominant cluster, which correspond to the spectra lying outside the prostate, are pruned and eliminated from subsequent analysis. The recursive scheme alternates between computing the low dimensional manifold of all the spectra in the 3D MRS scene and the unsupervised clustering algorithm to identify and eliminate non-informative spectra. This scheme is recursed until the sub-clusters corresponding to cancer spectra are identified. The primary contributions and novel aspects of this work are,

- The use of non-linear dimensionality reduction methods (spectral clustering) to exploit the inherent non-linearity in the high dimensional spectral data and embed the data in a reduced dimensional linear subspace.

- A hierarchical clustering scheme to recursively distinguish informative from non-informative spectra in the lower dimensional embedding space.
- The cascaded scheme enables accurate identification of cancer spectra efficiently, and is also qualitatively shown to perform better compared to the popular z-score method.

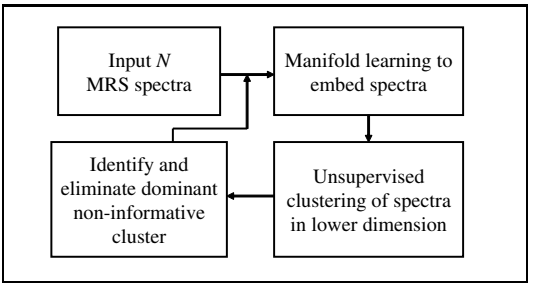

Fig. 1. Flowchart showing various system components and methodological overview

The rest of this paper is organized as follows. In Section 2 we present a detailed description of our methods while in Section 3 we describe our results. Concluding remarks and future directions are presented in Section 4.

2 Methods

2.1 Data Description

We represent the 3D prostate T2 weighted scene by $\mathcal{C} = (C, f)$ where C is a 3D grid of voxels $c \in C$ and f is a function that assigns an intensity value to every $c \in C$. We also define a spectral image $\mathcal{C}^s = (G, g)$ where G is also a 3D grid superposed on C and $G \subset C$. For every spatial location $u \in G$, there is an associated spectra $g(u)$. Hence while $f(c)$ is a scalar, $g(u)$ is a 256 dimensional vector valued function. Note that the size of the spatial grid locations $u \in G$ is equivalent to 8×8 voxels $c \in C$.

The spectral datasets used for the study were collected during the ACRIN multi-site trial of prostate MRS acquired with 1.5 Tesla GE Medical Systems through the PROSE(c) package (voxel width 0.4 x 0.4 x 3 mm). Datasets were obtained from 14 patients having CAP with different degrees of severity. The spectral grid was contained in DICOM images from which the 16×16 grid containing N=256 spectra was obtained using IDL. Of these N spectra, over half are zero-padded or non-informative lying outside the prostate. Figures 2 (a)-(b) show a spectral grid superimposed on a T2 MR image and the corresponding spectra. Figure 2(c) shows the magnified version of a single 256 dimensional MRS spectra from within G. Note that due to noise in the spectra, it is very difficult to identify individual peaks for creatine, choline and citrate from $g(u)$.

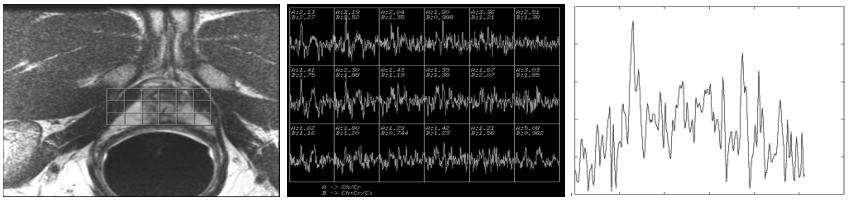

Fig. 2. (a) Slice of T2 weighted MR image with overlay of MRS grid (G), (b) Individual MRS spectra acquired from each $u \in G$, and (c) a single 256 dimensional MRS spectra ($g(u)$). Note that the prostate is normally contained in a 3×6 or a 3×7 grid which varies in different studies based on the prostate size. Radiologists look at relative peak heights of creatine, choline and citrate within $g(u)$ to identify possible cancer presence.

2.2 Determining Ground Truth for Cancer Extent on Prostate MRS

Since whole mount histological sections corresponding to the MR/MRS studies were not available for the ACRIN database, we were only able to determine approximate spatial extent of cancer within G [9]. During the MRSI ACRIN study, *arbitrary divisions* were established by the radiologists to obtain a rough estimate of the location of cancer. The prostate was first divided into two regions: Left(L) and Right(R) and the slices were then further divided into three regions: Base(B), Midgland(M) and Apex(A). Thus a total of 6 potential cancer locations were defined: Left Base(LB), Left Midgland(LM), Left Apex(LA), Right Base(RB), Right Midgland(RM) and Right Apex(RA). Presence or absence of cancer in each of these 6 candidate locations, determined via needle biopsy, was recorded. The maximum diameter of the cancer was also recorded in each of the 6 candidate locations.

For a MRS scene $\mathcal{C}^s$, with known cancer in left midgland (LM), the prostate contained in a 3×6 grid and prostate midgland extending over 2 contiguous slices, we define a potential cancer space $G^P \subset G$, within which the cancer is present. If we separate G into two equal right and left halves of 3×3, the total number of voxels $u \in G^P$ is 18 ($3 \times 3 \times 2$). The total number of actual cancer voxels within the cancer space, G^P, is obtained by knowledge of maximum diameter of cancer and given as: No of candidate slices $\times \left\lceil \frac{(MaxDiameter)^2}{\Delta x \Delta y} \right\rceil$, where $\lceil\rceil$ refers to the ceiling operation and $\Delta x, \Delta y$ refer to the size of voxel u in the x and y dimensions. Hence for a study with a cancer with maximum diameter of 13.75 mm in LM, 8 voxels within G^P correspond to cancer. Note that the cancer ground truth we determine, does not convey information regarding the precise spatial location of cancer voxels within G^P, only the number.

2.3 Manifold Learning Via Spectral Clustering

The spectra $g(u)$, for $u \in G$ lies in a 256 dimensional space. Hence, our aim is to find a embedding vector $\hat{X}(u)$ for each voxel $u \in G$, and its associated class ω

such that the distance between $g(u)$ and ω is monotonically related to G in the lower dimensional space. Hence if voxels $u, v \in G$ both belong to class ω, then $[\hat{X}(u) - \hat{X}(v)]^2$ should be small. To compute the optimal embedding, we first define a matrix W representing the similarity between any two objects $u, v \in G$ in high dimensional feature space.

$$W(u,v) = e^{||g(u)-g(v)||} \in R^{|G|\times|G|}, \tag{1}$$

where $|G|$ is the cardinality of set G. The embedding vector $\widehat{X}$ is obtained from the maximization of the function:

$$E_W(\hat{X}) = 2\gamma \frac{\widehat{X}^T(D-W)\widehat{X}}{\widehat{X}^T D\widehat{X}}, \tag{2}$$

where $D(u,u)=\sum_v W(u,v)$ and $\gamma = |G| - 1$. The embedding space is defined by the eigenvectors corresponding to the smallest A eigenvalues of $(D - W)$ $\widehat{X} = \lambda DW$ for every $u \in G$, the embedding $\hat{X}(u)$ contains the coordinates of u in the embedding space and is given as, $\hat{X}(u) = [\hat{e}_A(u) \| A \in \{1, 2, \cdots, \beta\}]$ where $\hat{e}_A(u)$, is a A dimensional vector of eigen values associated with u.

2.4 Hierarchical Cascade to Prune Non-informative Spectra

At each iteration t, for a subset of voxels u, $\tilde{G}_t \subset G$ is obtained by eliminating the non-informative spectra $g(u)$. The voxels $u \in \tilde{G}_t$ are aggregated into clusters V_T^1, V_T^2, V_T^3 by applying k-means clustering to all $u \in G$ in the low dimensional embedding $\widehat{X}(u)$. The number of clusters $k = 3$ was chosen empirically to correspond to cancer, benign and classes whose attributes are intermediate to normal tissue and cancer (e.g. benign hyperplasia (BPH), high-grade prostatic intraepithelial neoplasia (HGPIN)). Initially, most of the locations $u \in G$ correspond to zero padded or non informative spectra and hence the scheme proceeds by eliminating the dominant cluster. Clusters corresponding to cancer and areas within the prostate only become resolvable at higher levels of the cascade scheme after elimination of the dominant non informative spectra. The algorithm below describes precisely how our methodology works.

Algorithm *HierarclustMRSprostate*
Input: $g(u)$ for all $u \in G$, T, G
Output: $\tilde{G}_T$, V_T^1, V_T^2, V_T^3
begin
0. Initialize $\tilde{G}_0 = G$;
1. for $t = 0$ to T do
2. Apply Graph Embedding to $g(u)$, for all $u \in \tilde{G}_t$ to obtain $\hat{X}_t(u)$;
3. Apply k-means clustering to obtain clusters V_t^1, V_t^2, V_t^3;
4. Identify largest cluster V_t^{max};
5. Create set $\tilde{G}_{t+1} \subset \tilde{G}_t$ by eliminating all $u \in V_t^{\max}$ from $\tilde{G}_t$;
6. *endfor*
end

Note that since we employ an unsupervised learning approach, it is not clear which of V_T^1, V_T^2, V_T^3 actually represents the cancer cluster. The motivation behind the $HierarclustMRSprostate$ algorithm however is to obtain clusters V_T^1, V_T^2, V_T^3 which represent, to the extent possible, distinct tissue classes.

3 Results

3.1 Qualitative Results

The images in Figure 3 demonstrate qualitative results of our hierarchical cascade scheme for distinguishing informative from non-informative spectra. Figure 3(a) represents a spatial map of $\tilde{G}_0$ (16×16) superimposed on the corresponding T2 weighted scene and every $u \in \tilde{G}_0$ is assigned one of three colors, in turn corresponding to one of the three classes determined based on the embedding $\widehat{X}(u)$. Note that the dominant cluster(spatial locations in blue in Figure 3(a)) has been eliminated in $\tilde{G}_1$ (16×8) (Figure 3(b)) and from $\tilde{G}_2$ (8×4). The lowest level in the cascade ($\tilde{G}_3$ in 3(d)) is obtained from $\tilde{G}_2$ after the spectra on the periphery of the prostate (blue locations) have been removed. The cancer spectra are visible as a distinct cluster (blue cells) on Figure 3(d). The cascade scheme permits the resolvability of the 3 tissue classes (one of which is cancer) which were collapsed together within the informative cluster at the higher levels in the cascade $(\tilde{G}_0, \tilde{G}_1, \tilde{G}_2)$.

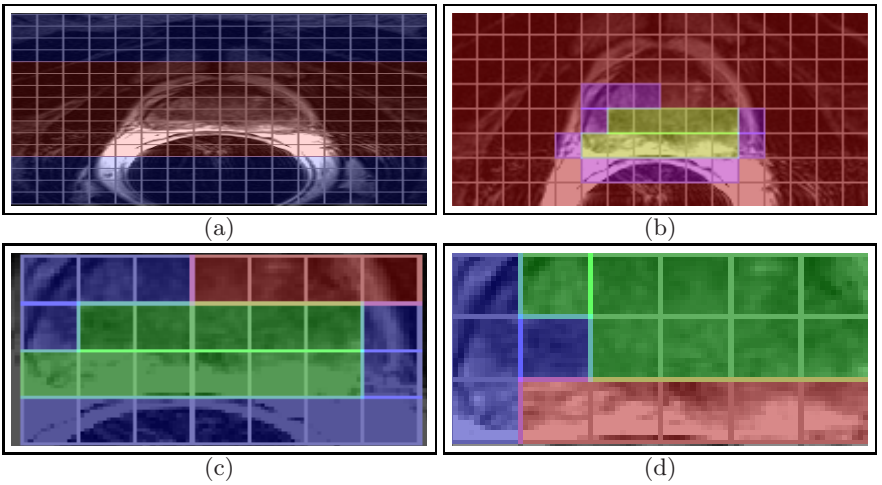

Fig. 3. Spectral grids for a single slice within C^S are shown superimposed on T2 for (a) $\tilde{G}_0$, (b) $\tilde{G}_1$, (c) $\tilde{G}_2$, and (d) $\tilde{G}_3$. Note that the size of the grid reduces from 16×16 (a) to 6×3 (d) by elimination of non-informative spectra on the 16×16 grid on T2. In 3 (d) cancer class could clearly be discerned as the blue class, since the cancer is located in the right MG slice. Note that the right/left conventions in radiology are reversed.

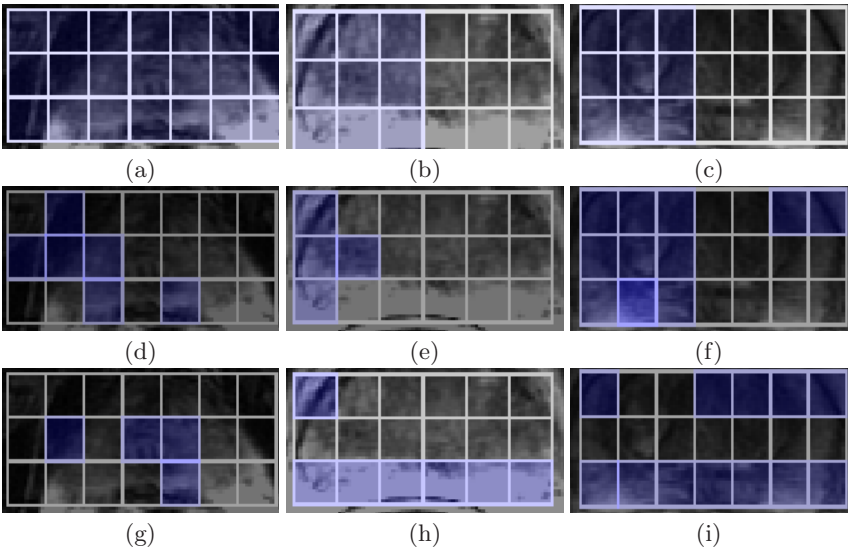

Fig. 4. (a)-(c) represents the potential cancer locations in blue. Figures 4(d)-(f) shows our classification results for 3 different studies with the cancer voxels shown in blue, while Figures 4(g)-(i) shows the z-score results. Note that for the last study in 4(f), our method has 100% sensitivity with just 2 FP voxels.

Figures 4(a)-(c) shows the potential cancer location volume within G on a single 2D slice of a T2 weighted prostate MR scene for 3 different studies and Figures 3 (d)-(f)show the corresponding result of our hierarchical scheme. In all 3 studies, the cancer cluster can be appreciated as a distinct class on the grid, corresponding to the location specified in the pathology reports. In Figures 4 (d)-(i) we compare the grid maps obtained via our cascaded clustering scheme (4(d)-(f)) with the corresponding plots obtained using the z-score (4(g)-(i)), a popular MRS analysis scheme. In this method, the z-score value $z(u)$, is assigned to each $u \in G$. The spectra $g(u)$ is then assigned to one of the two classes based on whether $z(u)$ is lower or greater than a pre determined threshold. It is apparent from the plots in 4(g)-(i) that none of the classes appear to represent the cancer class. These results clearly suggest that our cascaded scheme has higher specificity compared to the z-score method.

3.2 Quantitative Results

Table 1 shows the quantitative results for 14 different studies. True positive (TP), False positive (FP) and False negative (FN) fractions for every dataset were obtained by comparing the automated results with the ground truth voxels for all the 3 classes obtained. The class which corresponded to maximum TP and minimum FP and FN rates was identified as the cancer class and the respective TP, FP and FN values were reported for that particular class. TP,

FP and FN percentage values for each of the dataset were then calculated by dividing the TP, FP and FN fraction by the total number of ground truth voxels determined as described in Section 2.2. Average results over 14 studies have been reported. Clearly our scheme appears to have high cancer detection sensitivity and specificity.

Table 1. Table showing the average percentage values of TP, FP and FN for our automated detection scheme, averaged over a total of 14 MRS studies

Average TP	Average FP	Average FN
77.80	28.97	20.92

4 Concluding Remarks

In this paper we have presented a novel application of manifold learning and hierarchical clustering for the automated identification of cancer spectra on prostate MRS. Main contributions of our work are:

- The integration of unsupervised clustering with a manifold learning scheme to identify and eliminate non informative spectra.
- The hierarchical cascade detection scheme to efficiently and accurately identify prostate cancer spectra.
- Comparison of our scheme against a popular current MRS analysis scheme (z-score) with respect to an approximately detected ground truth suggests that our hierarchical detection algorithm has comparable sensitivity and higher specificity.

In future work we intend to extend our scheme for prostate MRS analysis to incorporate corresponding structural MRI data in order to develop better predictors for identifying CAP.

References

1. Catalona, W., et al.: Measurement of Prostate-Specific Antigen in serum as a Screening Test for Prostate Cancer. J.Med. 324(17), 1156–1161 (1991)
2. Kurhanewicz, J., et al.: Analysis of a Complex of Statistical Variables into Principal Components. Magnetic Resonance in Medicine 50, 944–954 (2003)
3. McKnight, T., et al.: An Automated Technique for the Quantitative Assessment of 3D-MRSI Data from Patients with Glioma. Journal of Magnetic Resonance Imaging 13, 167–177 (2001)
4. Pels, P., Ozturk-Isik, et al.: Quantification of Prostate MRSI Data by Model-Based Time Domain fitting and Frequency Domain Analysis. NMR in Biomedicine 19, 188–197 (2006)
5. Ma, J., Sun, J.: MRS Classification based on Independent Component Analysis and Support Vector Machines. In: IEEE Intl. Conf. on Hybrid Intel. Syst., pp. 81–84. IEEE Computer Society Press, Los Alamitos (2005)

6. Simonetti, A., Melssen, W., Edelenyi, F., et al.: Combination of Feature-Reduced MR Spectroscopic and MR Imaging Data for Improved Brain Tumor Classification. NMR in Biomedicine 18, 34–43 (2005)
7. Lee, G., Madabhushi, A., Rodriguez, C.: An Empirical Comparison of Dimensionality Reduction Methods for Classifying Gene and Protein Expression Datasets. ISBRA (2007)
8. Anant Madabhushi, J.S., et al.: Graph Embedding to Improved Supervised Classification and Novel Class Detection: Application to Prostate Cancer. In: Duncan, J.S., Gerig, G. (eds.) MICCAI 2005. LNCS, vol. 3749, pp. 729–737. Springer, Heidelberg (2005)
9. Madabhushi, A., Feldman, M., et al.: Automated Detection of Prostatic Adenocarcinoma from High-Resolution Ex Vivo MRI. IEEE Transactions on Medical Imaging 24(12) (2005)