

A Consensus Embedding Approach for Segmentation of High Resolution *In Vivo* Prostate Magnetic Resonance Imagery

Satish Viswanath^a, Mark Rosen^b, and Anant Madabhushi^a

^aDepartment of Biomedical Engineering, Rutgers University, 599 Taylor Road, Piscataway, NJ, USA 08854

^bDepartment of Radiology, University of Pennsylvania, 3400 Spruce Street, Philadelphia, PA, USA 19104

ABSTRACT

Current techniques for localization of prostatic adenocarcinoma (CaP) via blinded trans-rectal ultrasound biopsy are associated with a high false negative detection rate. While high resolution endorectal *in vivo* Magnetic Resonance (MR) prostate imaging has been shown to have improved contrast and resolution for CaP detection over ultrasound, similarity in intensity characteristics between benign and cancerous regions on MR images contribute to a high false positive detection rate. In this paper, we present a novel unsupervised segmentation method that employs manifold learning via consensus schemes for detection of cancerous regions from high resolution 1.5 Tesla (T) endorectal *in vivo* prostate MRI. A significant contribution of this paper is a method to combine multiple weak, lower-dimensional representations of high dimensional feature data in a way analogous to classifier ensemble schemes, and hence create a stable and accurate reduced dimensional representation. After correcting for MR image intensity artifacts, such as bias field inhomogeneity and intensity non-standardness, our algorithm extracts over 350 3D texture features at every spatial location in the MR scene at multiple scales and orientations. Non-linear dimensionality reduction schemes such as Locally Linear Embedding (LLE) and Graph Embedding (GE) are employed to create multiple low dimensional data representations of this high dimensional texture feature space. Our novel consensus embedding method is used to average object adjacencies from within the multiple low dimensional projections so that class relationships are preserved. Unsupervised consensus clustering is then used to partition the objects in this consensus embedding space into distinct classes. Quantitative evaluation on 18 1.5 T prostate MR data against corresponding histology obtained from the multi-site ACRIN trials show a sensitivity of 92.65% and a specificity of 82.06%, which suggests that our method is successfully able to detect suspicious regions in the prostate.

Keywords: segmentation, prostate cancer, computer-aided diagnosis, MRI, *in vivo*, 1.5 Tesla, manifold learning, consensus clustering, consensus embedding

1. INTRODUCTION

Prostatic adenocarcinoma (CaP) is the second leading cause of cancer related deaths among males in America, with an estimated 220,000 new cases every year (Source: *American Cancer Society*). Recently Magnetic Resonance Imaging (MRI) has emerged as a promising modality for possible early detection of CaP *in vivo* thus improving the accuracy of prostate biopsies.¹ The current standard for detection of CaP is transrectal ultrasound (TRUS) guided symmetrical needle biopsy which due to the poor image resolution of ultrasound is associated with a false negative rate of up to 30%.² While *in vivo* endorectal 1.5 Tesla (T) MR imaging of the prostate has allowed for greater discrimination between benign and cancerous prostatic structures as compared to TRUS, it has been unable to detect small foci of carcinoma contributing to a relatively low specificity.¹ Even at higher resolutions recognition of small tumors by experts is confounded by the fact that several benign features, such as atrophied gland and areas of stromal over-growth appear similar to tumor areas.

Recently several researchers have begun to explore the use of computer aided diagnosis (CAD) schemes for early detection and classification of CaP.³ In [4] Madabhushi et al. presented a novel supervised CAD

Contact: Anant Madabhushi, E-mail: anantm@rci.rutgers.edu, Telephone: 1 732 445 4500 x6213

scheme for detection of CaP from 4 T *ex vivo* prostate MRI. A weighted feature ensemble scheme was used to integrate multiple 3D texture features to generate a likelihood scene in which the intensity at every spatial location corresponded to the probability of cancer being present. Improvements to the method via use of non-linear dimensionality reduction (graph embedding) and multiple classifier systems were reported in [5] and [6] respectively. In [7] Chan et al. presented a multimodal statistical classifier which integrated texture features from multi-protocol 1.5 T *in vivo* MRI to generate a statistical probability map representing likelihoods of cancer for different regions within the prostate. Area under the Receiver-Operating Characteristic (ROC) curve (AUC) was used to estimate the classifier accuracy. A maximum AUC of 0.839 was reported.

In this paper we present a novel unsupervised scheme to segment different regions within 1.5 T endorectal *in vivo* MR prostate imagery. We first correct for MR related artifacts, bias field inhomogeneity⁸ and intensity non-standardness.⁹ This is followed by extraction of over 350 3D texture features at every spatial location within the MRI image. These features have been previously shown to be able to differentiate cancerous and non-cancerous regions.⁴ To avoid the *curse of dimensionality* associated with high dimensional feature spaces, the textural data at every spatial location is projected non-linearly into a lower dimensional space where the objects (MR voxels in this case) can be clustered into distinct classes. Due to inherent non-linearities in biomedical data, linear dimensionality reduction schemes such as Principal Component Analysis (PCA) have been shown to perform poorly¹⁰ compared to non-linear dimensionality reduction (NLDR) schemes such as Locally Linear Embedding¹¹ (LLE) and Graph Embedding¹² (GE) in unraveling object clusters while preserving class relationships.

While our group has had success in the use of NLDR for the automated classification of prostate magnetic resonance spectroscopy (MRS) data¹³ as well as protein and gene expression data,¹⁰ methods such as LLE and GE are sensitive to the choice of parameters. NLDR schemes attempt to preserve geodesic distances between objects from the high- to the low-dimensional spaces unlike PCA which preserves Euclidean distances. Methods such as LLE estimate object distances by assuming that within a small local neighborhood objects are linearly related. The geodesic estimate is thus a function of the size of the neighborhood within which local linearity is assumed. These NLDR schemes are also sensitive to the high dimensional feature space within which geodesic distances are computed since relative object adjacencies may change from one feature space to another. As an example, consider a feature vector $F_1(u)$ associated with each object u in a set C . Let the lower dimensional co-ordinates of three objects $u, v, w \in C$, based on F_1 , be given by $X_1(u), X_1(v), X_1(w)$, where X_1 is the principal Eigenvector obtained via application of NLDR to F_1 . Let us assume that of the 3 objects, u, v belong to class ω_1 while w belongs to class ω_2 . Assuming that the data has been properly projected into the lower dimensional space, then it should follow that $\|X_1(u) - X_1(v)\|_2 < \|X_1(u) - X_1(w)\|_2$, where $\|\cdot\|_2$ represents the Euclidean norm. Note that the above is true only if F_1 accurately encapsulates the class related information regarding $u, v, w \in C$. However, this may not hold for another feature set F_2 which on account of noisy or missing attributes may result in low dimensional projections in X_2 such that $\|X_2(u) - X_2(v)\|_2 > \|X_2(u) - X_2(w)\|_2$. In order to represent the true relationship between u, v, w , the adjacency between objects in these lower dimensional embedding spaces is then represented as a function of the distance between the objects along the lower-dimensional manifold. In this paper we propose a scheme wherein multiple such representations are combined to obtain a stable embedding representing the true class relationship between objects in high dimensional space. Analogous to classifier ensemble schemes for creating strong stable classifiers by combining multiple weak unstable classifiers with large bias and variance, the *consensus embedding* scheme will yield a more stable data embedding by reducing the variance in the individual embedding spaces. This is done by computing a consensus distance matrix W which reflects the averaged relative object adjacencies between $u, v, w \in C$ in multiple low dimensional data projections. Multidimensional scaling¹⁴ (MDS) is applied to W to obtain the final stable data embedding. Consensus clustering¹⁵ is then applied to segregate objects into distinct categories in an unsupervised fashion. Figure 1 illustrates the main steps in our system as a flowchart.

In this paper, a novel unsupervised classification algorithm is employed to partition MR images into distinct regions by means of consensus embedding and consensus clustering in 18 1.5 T *in vivo* prostate MR datasets. Ground truth estimates of cancer on partial whole mount histological sections were available which were used to define the spatial extent of CaP on MRI, which was then used for quantitative evaluation. The rest of this paper is organized as follows. In Section 2 we present the details of our experimental design. Results of quantitative and qualitative evaluation of our scheme are presented in Section 3, followed by concluding remarks in Section 4.

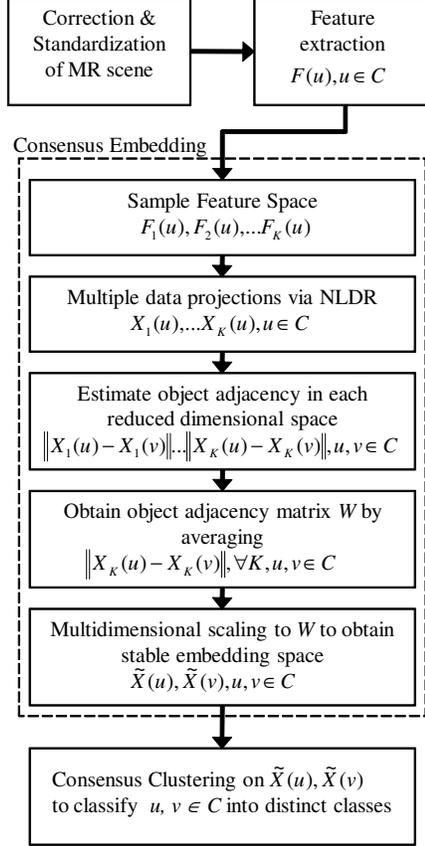


Figure 1: Flowchart showing different system components and overall organization.

2. EXPERIMENTAL DESIGN

2.1 Data Description and Notation

A total of 18 1.5 T *in vivo* endorectal MRI and MRS datasets were collected from the American College of Radiology Imaging Network (ACRIN) multi-site prostate trial*. For each patient, MRI data (T2 imaging protocol) was acquired prior to radical prostatectomy. Following resection, the gland was quartered and stained. These sections were then manually examined for CaP to constitute the ground truth on histology. These regions were then manually mapped onto the MRI images to estimate CaP presence and extent.

We define a 3D MRI scene $\mathcal{C} = (C, f)$ where C is a set of spatial locations $c_i \in C, i \in \{1, \dots, |C|\}$, $|C|$ is the cardinality of any set C and $f(c)$ is a function that assigns an intensity value to every $c \in C$. We define this 3D image at MRS metavoxel resolution as $\hat{\mathcal{C}} = (\hat{C}, \hat{f})$, where \hat{C} is a 3D grid of metavoxels at locations $\hat{c}_i \in \hat{C}, i \in \{1, \dots, |\hat{C}|\}$. Figure 2 shows the relationship between the MRS metavoxel \hat{c} and the MRI voxel c . It is important to note that the distance between any two adjacent metavoxels $\hat{c}_i, \hat{c}_j \in \hat{C}, \|\hat{c}_i - \hat{c}_j\|_2$, (where $\|\cdot\|_2$ denotes the L_2 norm) is roughly 16 times the distance between any two adjacent voxels $c_i, c_j \in C$. We accordingly define $\hat{f}(\hat{c}_i), \forall \hat{c}_i \in \hat{C}$.

2.2 Determination of Approximate Ground Truth for CaP on MRI

Partial ground truth for the 1.5 T MR datasets in the ACRIN database is available in the form of approximate sextant locations and sizes of cancer for each dataset as described previously. We have previously developed an algorithm for the registration of *ex vivo* MRI and whole-mount histological (WMH) images¹⁶ for accurate

*http://www.acrin.org/6659_protocol.html

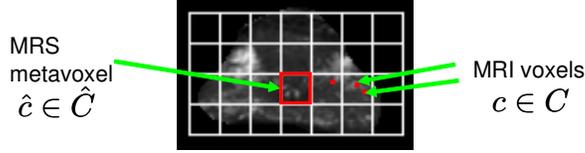


Figure 2: Figure showing the relationship between MRS metavoxels and MRI voxels. The spectral grid \hat{C} comprised of metavoxels \hat{c} has been overlaid on an MR slice and is shown in white. Note that the region outlined in red on \hat{C} corresponds to the area occupied by a metavoxel $\hat{c} \in \hat{C}$ but will contain multiple MRI voxels $c \in C$ within that region.

mapping of the spatial extent of CaP from WMH sections onto MRI. However most of the histology data in the ACRIN study are not WMH, but small sections of the gland which makes it difficult for them to be reconstituted into WMH sections. Hence the CaP ground truth estimate on the MRI sections is obtained in the following manner. The MR image of the prostate is visually divided into two lateral compartments: Left (L) and Right (R); and further divided into 3 regions longitudinally: Base (B), Midgland (M) and Apex (A). Presence of CaP (*potential cancer space*) has been previously determined in one or more of these six regions: Left Base (LB), Left Midgland (LM), Left Apex (LA), Right Base (RB), Right Midgland (RM) and Right Apex (RA) via manual mapping of CaP from histology onto the corresponding MRI sections. The maximum diameter of the tumor is also recorded in each of the 6 candidate locations and is denoted as *MaxDiameter*. The total number of possible cancer voxels $c \in C$ at the MR voxel resolution within the cancer space is given as: No. of candidate slices $\times \left\lceil \frac{MaxDiameter^2}{\Delta x \Delta y} \right\rceil$, where $\lceil \cdot \rceil$ refers to the ceiling operation and $\Delta x, \Delta y$ refer to the dimensions of the MR voxel c in the X and Y dimensions. Similarly, we calculated the number of possible cancer metavoxels $\hat{c} \in \hat{C}$ at the MRS metavoxel resolution as: No. of candidate slices $\times \left\lceil \frac{MaxDiameter^2}{\Delta \hat{x} \Delta \hat{y}} \right\rceil$, where $\Delta \hat{x}, \Delta \hat{y}$ refer to the dimensions of the MRS metavoxel \hat{c} in the \mathbf{X} and \mathbf{Y} dimensions. Note that the exact spatial location of CaP voxels on a particular slice is not available, only the size and sextant within which it occurs. This *potential cancer space* nonetheless serves as a basis to perform a semi-quantitative evaluation of our CAD scheme.¹³

2.3 Correcting Bias field and non-linear MR image intensity artifacts

Image intensity variations due to radio frequency (RF) field inhomogeneities may be caused due to a number of different factors including poor RF field uniformity, static field inhomogeneity and RF penetration.⁴ We used the ITK⁸ toolkit's BiasCorrector algorithm to correct the original 3D MR scenes for bias field inhomogeneity and an interactive version of the image intensity standardization algorithm previously presented⁹ to correct for non-linearity of image intensities. Figures 3(e), (f) shows the results of these pre-processing steps on a sample MR slice. Figure 3(e) shows the effect of bias field correction on Figure 3(d). Figure 3(f) then shows the effect of intensity standardization on the slice shown in Figure 3(e).

Image intensity standardization is a post-MR acquisition processing operation designed for correcting inter-acquisition signal intensity variations (non-standardness). Such grayscale intensities do not have a fixed tissue-specific meaning within the same imaging protocol, the same body region, or even within the same patient. When the histograms for different prostate studies (in different colors) are plotted together (as seen in Figure 3(a)), we see that the intensity distributions have different intensity ranges and are not in alignment. To lend meaning to corresponding gray values in the MR images we non-linearly map their intensity scales to a common scale, known as the *standard* scale. This is done by choosing one of the datasets to be the standard image; the intensity profile of this dataset is then used as the standard intensity scale. Landmarks are first manually identified on the intensity histograms for each of the studies (Figure 3(a)). These landmarks are then non-linearly mapped to corresponding landmarks manually identified on the standard intensity scale⁹ shown in black on Figure 3(c). Note that this process is different from simple linear scaling of image intensities which does not address the non-standardness issue (Figure 3(b)). Note that before intensity standardization as well as for the result of the linear technique, the intensity histograms are misaligned (Figures 3(a), (b)), but are correctly aligned (Figure 3(c)) after intensity standardization.

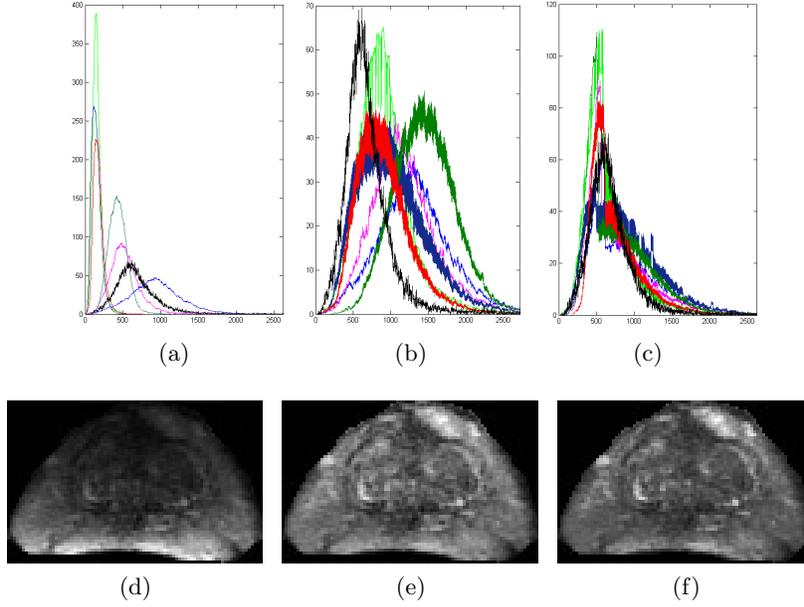


Figure 3: Image intensity histograms for 7 datasets (a) prior to standardization, (b) following linear scaling, and (c) following intensity standardization.⁹ Note the misalignment between the intensity histograms prior to standardization (a) which is not corrected for by linear scaling (b) but is by intensity standardization (c). The intensity histogram used for standardizing the other image intensity histograms is shown in black. The histograms for the other studies are shown with different colors. In Figures 3(d), (e), (f) are shown corresponding MRI slices from the original uncorrected 1.5 T *in vivo* endorectal MR scene, following bias field correction, and following intensity standardization respectively.

2.4 Feature Extraction

Over 350 3D texture feature scenes, corresponding to three different texture classes were extracted from each MRI scene. These feature representations were chosen since they have been demonstrated to be able to discriminate between the cancer and non-cancer classes.⁴ We calculated the feature scenes $\mathcal{F}_u = (C, f_u)$ for each C by applying the feature operators $\Phi_u, u \in \{1, \dots, 373\}$ within a local neighborhood associated with every $c \in C$. Hence $f_u(c)$ is the feature value associated with feature operator Φ_u at voxel c . We can hence define a feature vector associated with each $c \in C$ as $\mathbf{F}(c) = [f_u(c)|u \in \{1, \dots, 373\}]$. We define a κ -neighborhood centered on $c_i \in C$ as $\mathcal{N}_\kappa(c_i)$ where $\forall c_j \in \mathcal{N}_\kappa(c_i), \|c_j - c_i\| \leq \kappa, i, j \in \{1, \dots, |C|\}, c_i \notin \mathcal{N}_\kappa(c_i)$. We similarly define a κ -neighborhood $\mathcal{N}_\kappa(\hat{c}_i)$ for $\hat{c}_i \in \hat{C}$ where $\forall c_j \in \mathcal{N}_\kappa(\hat{c}_i), \|c_j - \hat{c}_i\| \leq \kappa, \hat{i} \in \{1, \dots, |\hat{C}|\}, j \in \{1, \dots, |C|\}, c_j \neq \hat{c}_i$. Based on our definition for a metavoxel in Section 2.1, we can similarly define a feature attribute for each metavoxel $\hat{c} \in \hat{C}$ as the median $\hat{f}_u(\hat{c}_i) = \text{MEDIAN}_{c_a \in \mathcal{N}_\kappa(\hat{c}_i)} [f_u(c_a)], a \in \{1, \dots, |C|\}, \hat{i} \in \{1, \dots, |\hat{C}|\}$. The corresponding feature vector is then given as $\hat{\mathbf{F}}(\hat{c}) = [\hat{f}_u(\hat{c})|u \in \{1, \dots, 373\}], \forall \hat{c} \in \hat{C}$ as earlier.

1. **Gradient Features:** Gradient features are calculated using steerable and non-steerable linear gradient operators. Eleven non-steerable gradient features were obtained using Sobel, Kirsch and standard derivative operations. Gabor gradient operators⁴ comprising the steerable class of gradient calculations were defined for every $c \in C$ where $c = (x, y, z)$,

$$f_u(c) = \frac{1}{2^{\frac{3}{2}} \sigma_X \sigma_Y \sigma_Z} e^{-\frac{1}{2} \left[\frac{x^2}{\sigma_X^2} + \frac{y^2}{\sigma_Y^2} + \frac{z^2}{\sigma_Z^2} \right]} \cos(2\pi\omega x), \quad (1)$$

where ω is the frequency of a sinusoidal plane wave along the X -axis, and σ_X, σ_Y , and σ_Z are the space constraints of the Gaussian envelope along the X, Y , and Z directions respectively. The orientation of the filter, θ , is affected by the coordinate transformations: $x' = r(x \cos \theta + y \sin \theta)$, $y' = r(-x \sin \theta + y \cos \theta)$

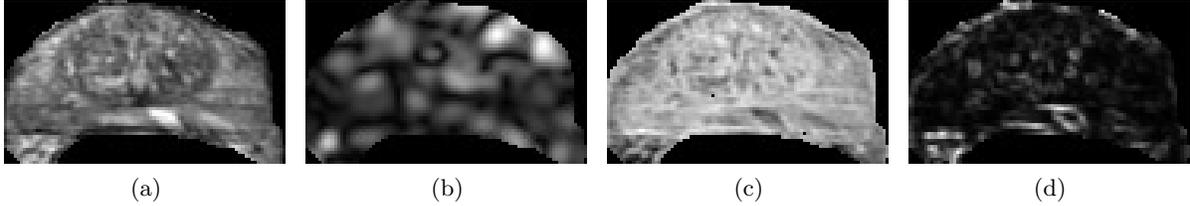


Figure 4: (a) A 2D section from \mathcal{C} following bias field correction and intensity standardization, and corresponding 2D sections from feature scenes \mathcal{F}_u for, (b) Gabor ($\theta = \frac{\pi}{3}$, $\lambda = -1$, $\kappa = 5$) (c) first order statistical (range, $\kappa = 3$), and (d) second order statistical (Haralick energy, $\kappa = 3$, $G = 64$, $d = 1$).

and $z' = r(z)$, where r is the scaling factor. These were computed within the sliding window neighborhood \mathcal{N}_κ . Gabor gradient features were calculated at 13 scales ($r \in \{-\frac{\pi}{16}, -\frac{\pi}{8\sqrt{2}}, \dots, \frac{\pi}{16}\}$), 6 orientations ($\theta \in \{0, \frac{\pi}{6}, \frac{\pi}{3}, \frac{\pi}{2}, \frac{2\pi}{3}, \frac{5\pi}{6}\}$) and 3 window sizes ($\kappa \in \{3, 5, 7\}$). Figure 4(b) shows a feature image extracted by applying a Gabor operator ($\theta = \frac{\pi}{3}$, $r = -\frac{\pi}{\sqrt{2}}$, $\kappa = 5$) on a 2D slice from \mathcal{C} (Figure 4(a)).

2. **First Order Statistical Features:** Four first order statistical features for 3 different window sizes were calculated. They included the mean, median, standard deviation, and range for the gray values of voxels within the sliding window neighborhood \mathcal{N}_κ , $\kappa \in \{3, 5, 7\}$. Figure 4(c) shows a feature image obtained via application of a first order statistical operator (range, $\kappa = 3$) for the MRI section in Figure 4(a).
3. **Second Order Statistical Features:** To calculate the second order statistical (Haralick) feature scenes, we compute a $G \times G$ co-occurrence matrix $P_{d,c,\kappa}$ associated with $\mathcal{N}_\kappa(c)$, where G is the maximum gray scale intensity in \mathcal{C} . The value at any location $[g_1, g_2]$ in $P_{d,c,\kappa}$, where $g_1, g_2 \in \{1, \dots, G\}$, represents the frequency with which two distinct voxels $c_i, c_j \in \mathcal{N}_\kappa(c)$, $i, j \in \{1, \dots, |C|\}$ with associated image intensities $f(c_i) = g_1$, $f(c_j) = g_2$ are separated by distance d . A total of 13 Haralick features including energy, entropy, inertia, contrast, correlation, sum average, sum variance, sum entropy, difference average, difference variance, difference entropy, local homogeneity and average deviation were extracted at every voxel $c \in \mathcal{C}$, based on $P_{d,c,\kappa}$, for $\kappa \in \{3, 5, 7\}$, $d = 1$ and $G \in \{64, 128, 256\}$. Figure 4(d) shows a feature image (energy) from the co-occurrence matrix ($\kappa = 3$, $G = 64$, $d = 1$).

2.5 Non Linear Dimensionality Reduction

2.5.1 Graph Embedding

The aim of graph embedding¹² is to find an embedding vector $X_{GE}(c), \forall c \in \mathcal{C}$ such that the relative ordering of the distances between objects in the high dimensional feature space is preserved in lower dimensional space. Thus, if $c_i, c_j \in \mathcal{C}$, $i, j \in \{1, \dots, |C|\}$ are close in high dimensional feature space, then $\|X_{GE}(c_i) - X_{GE}(c_j)\|_2$ should be small, where $\|\cdot\|_2$ represents the Euclidean norm. This will only be true if the distances between all $c_i, c_j \in \mathcal{C}$ are preserved in the low dimensional mapping of the data. To compute the optimal embedding, we first define a matrix $W_{GE} \in \mathbb{R}^{|C| \times |C|}$, representing the adjacency between all objects $c \in \mathcal{C}$ in high-dimensional feature space. For all $c_i, c_j \in \mathcal{C}$, W_{GE} is defined as

$$W_{GE}(i, j) = e^{-\|\mathbf{F}(c_i) - \mathbf{F}(c_j)\|_2}, \forall c_i, c_j \in \mathcal{C}, i, j \in \{1, \dots, |C|\}. \quad (2)$$

$X_{GE}(c)$ is then obtained from the maximization of the function:

$$E(\mathcal{X}_{GE}) = 2\gamma \times \text{tr} \left[\frac{\mathcal{X}_{GE}(D - W_{GE})\mathcal{X}_{GE}^T}{\mathcal{X}_{GE}D\mathcal{X}_{GE}^T} \right], \quad (3)$$

where tr is the trace operator, $\mathcal{X}_{GE} = [X_{GE}(c_1), X_{GE}(c_2), \dots, X_{GE}(c_n)]$, $n = |C|$ and $\gamma = n - 1$. Additionally, D is a diagonal matrix where for all $c \in \mathcal{C}$, the diagonal element is defined as $D(i, i) = \sum_j W_{GE}(i, j)$. The embedding space is defined by the Eigenvectors corresponding to the smallest β Eigenvalues of $(D - W_{GE})$ $\mathcal{X}_{GE} = \lambda D W_{GE}$. The matrix $\mathcal{X}_{GE} \in \mathbb{R}^{|C| \times \beta}$ of the first β Eigenvectors is constructed, and $\forall c_i \in \mathcal{C}$, $X_{GE}(c_i)$ is defined as row i of \mathcal{X}_{GE} . $X_{GE}(c_i)$ is therefore a vector consisting of element number i from each of the first β Eigenvectors, which represents the β -dimensional Cartesian coordinates.

2.5.2 Locally Linear Embedding (LLE)

LLE¹¹ operates by assuming that objects in a neighborhood of a feature space are locally linear. Consider the set of feature vectors $\mathcal{F} = \{\mathbf{F}(c_1), \mathbf{F}(c_2), \dots, \mathbf{F}(c_n)\}$, $n = |C|$. We wish to map the set \mathcal{F} to the set $\mathcal{X} = \{X_{LLE}(c_1), X_{LLE}(c_2), \dots, X_{LLE}(c_n)\}$ of embedding co-ordinates. For all objects $c \in C$, LLE maps the feature vector $\mathbf{F}(c)$ to the embedding vector $X_{LLE}(c)$. Let $\{c_{\eta_i(1)}, \dots, c_{\eta_i(k)}\}$ be the k nearest neighbors of c_i where $\eta_i(k)$ is the index of the k^{th} neighbor of c_i in C . The feature vector $\mathbf{F}(c_i)$ and its k nearest neighbors (k NN), $\{\mathbf{F}(c_{\eta_i(1)}), \mathbf{F}(c_{\eta_i(2)}), \dots, \mathbf{F}(c_{\eta_i(k)})\}$ are assumed to lie on a patch of the manifold that is local linearly, allowing us to use Euclidean distances between the neighbors. Each $\mathbf{F}(c_i)$ can then be approximated by a weighted sum of its k NN. The optimal reconstruction weights are given by the sparse matrix W_{LLE} (subject to the constraint $\sum_j W_{LLE}(i, j) = 1$) that minimizes

$$S_1(W_{LLE}) = \sum_{i=1}^n \left\| \mathbf{F}(c_i) - \sum_{j=1}^k W_{LLE}(i, \eta_i(j)) \mathbf{F}(c_{\eta_i(j)}) \right\|_2. \quad (4)$$

Having determined the weighting matrix W_{LLE} , the next step is to find a low-dimensional representation of the points in \mathcal{F} that preserves this weighting. Thus, for each $\mathbf{F}(c_i)$ approximated as the weighted combination of its k NN, its projection $X_{LLE}(c_i)$ will be the weighted combination of the projections of these same k NN. The optimal \mathcal{X}_{LLE} in the least squares sense minimizes

$$S_2(\mathcal{X}_{LLE}) = \sum_{i=1}^n \left\| X_{LLE}(c_i) - \sum_{j=1}^n W_{LLE}(i, j) X_{LLE}(c_j) \right\|_2 = \text{tr}(\mathcal{X}_{LLE} L \mathcal{X}_{LLE}^T), \quad (5)$$

where tr is the trace operator, $\mathcal{X}_{LLE} = [X_{LLE}(c_1), X_{LLE}(c_2), \dots, X_{LLE}(c_n)]$, $L = (I - W_{LLE})(I - W_{LLE}^T)$ and I is the identity matrix. The minimization of (5) subject to the constraint $\mathcal{X}_{LLE} \mathcal{X}_{LLE}^T = I$ (a normalization constraint that prevents the solution $\mathcal{X}_{LLE} \equiv \mathbf{0}$) is an Eigenvalue problem whose solutions are the Eigenvectors of the Laplacian matrix L . Since the rank of L is $n-1$ the first Eigenvector is ignored and the second smallest Eigenvector represents the best one-dimensional projection of all the samples. The best two-dimensional projection is given by the eigenvectors with the second and third smallest eigenvalues, and so forth.

2.6 Consensus Embedding to obtain stable low dimensional data representation

We require a lower dimensional embedding that models the true nature of the underlying manifold that is described in high dimensional space. Varying the feature subspaces of the high dimensional manifold and the parameters (e.g. the number of k nearest neighbors in LLE) associated with NLDR methods achieves multiple embeddings which individually model relationships between objects. We propose a novel method to obtain this representation by generating multiple lower dimensional embeddings of feature subspaces and capturing the adjacencies between the voxels in the lower dimensional spaces. These adjacencies can then be combined to yield a more stable representative embedding. We generate multiple embeddings $X_{\phi, \alpha}(c)$ for $c \in C$, based on feature subspaces $F_\alpha(c) \subseteq \mathbf{F}(c)$, $\alpha \in \{1, \dots, B\}$ using the NLDR schemes $\phi \in \{GE, LLE\}$ described earlier. Each embedding $X_{\phi, \alpha}$ will hence represent adjacencies between voxels $c_i, c_j \in C$ based on the feature subspace F_α . Thus $\|X_{\phi, \alpha}(c_i) - X_{\phi, \alpha}(c_j)\|_2$ will vary as a function of F_α . To represent the true adjacency and class relationship between $c_i, c_j \in C$ we need to combine the multiple embeddings $X_{\phi, \alpha}$. A confusion matrix $W_{\phi, \alpha} \in \mathbb{R}^{|C| \times |C|}$ based on representing the adjacency between any two voxels $c_i, c_j \in C$ in the lower dimensional embedding representation $X_{\phi, \alpha}$ is first calculated as:

$$W_{\phi, \alpha}(i, j) = \|X_{\phi, \alpha}(c_i) - X_{\phi, \alpha}(c_j)\|_2, \quad (6)$$

where $c_i, c_j \in C, i, j \in \{1, \dots, |C|\}, \phi \in \{GE, LLE\}, \alpha \in \{1, \dots, B\}$. The confusion matrix $W_{\phi, \alpha}$ will hence represent the relationships between the voxels in each of the B embedding spaces $X_{\phi, \alpha}$, obtained via $F_\alpha, \alpha \in \{1, \dots, B\}$. We can average these voxel adjacencies as

$$\tilde{W}_\phi(i, j) = \frac{1}{B} \sum_{\alpha} W_{\phi, \alpha}(i, j), \forall i, j \in \{1, \dots, |C|\}, \quad (7)$$

where $\tilde{W}_\phi(i, j)$ represents the average distance in the reduced dimensional space over B feature sets F_α between the voxels $c_i, c_j \in C$. The idea is that not every F_α will represent the true class relationship between $c_i, c_j \in C$; hence $\tilde{W}_\phi(i, j)$ is a more reliable estimate of the true embedding distance between c_i, c_j . We then apply multidimensional scaling¹⁴ (MDS) to this \tilde{W}_ϕ to achieve the final combined embedding \tilde{X}_ϕ . MDS is implemented as a linear method that preserves the Euclidean geometry between each pair of voxels $c_i, c_j \in C$. This is done by finding optimal positions for the data points c_i, c_j in lower-dimensional space through minimization of the least squares error in the input pairwise Euclidean distances in \tilde{W}_ϕ . The complete algorithm for the consensus embedding scheme is described below:

Algorithm *ManifoldConsensusEmbed*

Input: $F_\alpha(c) \subseteq \mathbf{F}(c)$, for $\alpha \in \{1, \dots, B\}$, for all $c \in C$, B

Output: $\tilde{X}_\phi(c), \phi \in \{GE, LLE\}$

begin

0. *for* $\alpha = 1$ *to* B *do*

1. Sample $F_\alpha(c)$ from $\mathbf{F}(c)$ for all $c \in C$;

2. Use method ϕ to calculate $X_{\phi,\alpha}(c)$ for all $c \in C$, $\phi \in \{GE, LLE\}$;

3. Calculate $W_{\phi,\alpha}(i, j) = \|X_{\phi,\alpha}(c_i) - X_{\phi,\alpha}(c_j)\|_2$ for all $c_i, c_j \in C$;

4. *endfor*

5. Obtain $\tilde{W}_\phi(i, j) = \frac{1}{B} \sum_\alpha W_{\phi,\alpha}(i, j), \forall i, j \in \{1, \dots, |C|\}$;

6. Apply MDS to \tilde{W}_ϕ to obtain final combined embedding \tilde{X}_ϕ ;

end

2.7 Consensus Clustering on the consensus embedding space

To overcome the instability associated with centroid based clustering algorithms, we generate multiple weak clusterings $V_{\phi,t}^1, V_{\phi,t}^2, V_{\phi,t}^3, t \in \{0, \dots, T\}$ by repeated application of k -means clustering on the combined low dimensional manifold $\tilde{X}_\phi(c)$, for all $c \in C$ and $\phi \in \{GE, LLE\}$. Each cluster $V_{\phi,t}$ is a set of objects which has been assigned the same class label by the k -means clustering algorithm. As the number of elements in each cluster tends to change for each such iteration of k -means, we calculate a co-association matrix H_ϕ with the underlying assumption that voxels belonging to a *natural* cluster are very likely to be co-located in the same cluster for each iteration. Co-occurrences of pairs of voxels c_i, c_j in the same cluster $V_{\phi,t}$ are hence taken as votes for their *association*. $H_\phi(i, j)$ thus represents the number of times $c_i, c_j \in C$ were found in the same cluster over T iterations. If $H_\phi(i, j) = T$ then there is a high probability that c_i, c_j do indeed belong to the same cluster. We apply MDS¹⁴ to H_ϕ followed by a final unsupervised classification using k -means, to obtain the final stable clusters $\tilde{V}_\phi^1, \tilde{V}_\phi^2, \tilde{V}_\phi^3$. The algorithm¹⁷ is described below:

Algorithm *ManifoldConsensusClust*

Input: $\tilde{X}_\phi(c)$ for all $c \in C$, T

Output: $\tilde{V}_\phi^1, \tilde{V}_\phi^2, \tilde{V}_\phi^3$

begin

0. Initialize co-association matrix $H_\phi \in \mathbb{R}^{|C| \times |C|}, \phi \in \{GE, LLE\}$ to zero;

1. *for* $t = 0$ *to* T *do*

2. Use k -means on $\tilde{X}_\phi(c)$ to cluster all $c \in C$ into clusters $V_{\phi,t}^1, V_{\phi,t}^2, V_{\phi,t}^3$;

3. *if* $c_i, c_j \in C$ belong to the same cluster

4. $H_\phi(i, j) = H_\phi(i, j) + 1$;

5. *endif*

6. *endfor*

7. Apply MDS to H_ϕ followed by k -means to obtain final clusters $\tilde{V}_\phi^1, \tilde{V}_\phi^2, \tilde{V}_\phi^3$;

end

3. RESULTS AND DISCUSSION

3.1 Qualitative Results

Our scheme was applied to 18 1.5 T datasets. The CAD analysis was done at both voxel and metavoxel resolutions (\mathcal{C} and $\hat{\mathcal{C}}$).

3.1.1 1.5 T MR data at metavoxel resolution

Figure 5 shows the results of applying the algorithms *ManifoldConsensusEmbed* and *ManifoldConsensusClust* to the feature data in $\hat{\mathbf{F}}(\hat{c}), \forall \hat{c} \in \hat{\mathcal{C}}$. The different class labels obtained via clustering are represented via different colors as shown in the figure. The first column in Figure 5 corresponds to a 2D section from \mathcal{C} overlaid with the corresponding metavoxel grid $\hat{\mathcal{C}}$ from three different *in vivo* prostate MR studies (Figures 5 (a), (e), (i)). Figures 5 (b), (f), (j) show the *potential cancer space* area shaded with a translucent red on the overlaid grid $\hat{\mathcal{C}}$. Figures 5 (c), (g), (k) show the labels for objects clustered into $\hat{V}_{GE}^1, \hat{V}_{GE}^2, \hat{V}_{GE}^3$ plotted back onto the MRI slice (in red, green and blue) generated via Graph Embedding (GE). Figures 5 (d), (h), (l) show the labels for objects clustered into $\hat{V}_{LLE}^1, \hat{V}_{LLE}^2, \hat{V}_{LLE}^3$ plotted back onto the MRI slice obtained via LLE. It can be seen that the labels shaded in red correspond closely with the potential cancer space for both LLE and GE.

3.1.2 1.5 T MR data at voxel resolution

Figure 6 shows the results of our algorithm on 1.5 T prostate MRI studies at the voxel resolution; each row corresponding to a different dataset. The leftmost image in each row in Figure 6 shows the corrected and standardized slice from \mathcal{C} prior to analysis from three different *in vivo* prostate MR studies (Figures 6 (a), (e),

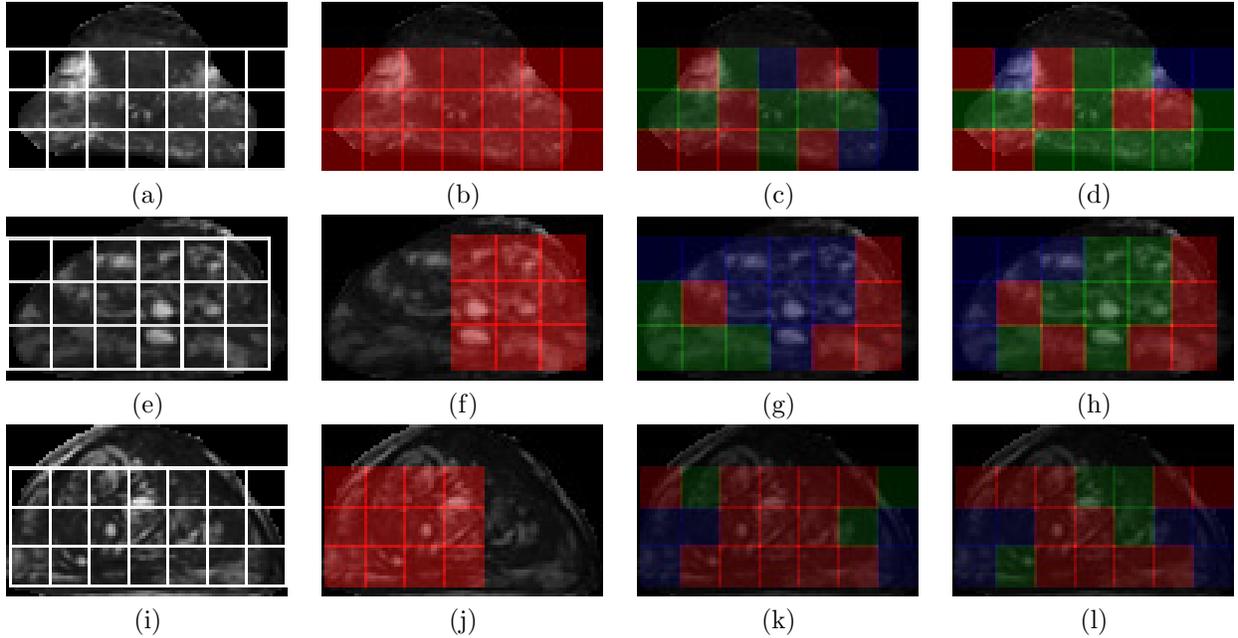


Figure 5: 2D sections from 3 different endorectal *in vivo* 1.5 T prostate studies, (a), (e), and (i) with the metavoxel grid $\hat{\mathcal{C}}$ superposed on them; (b), (f), and (j) showing the location of *potential cancer space* (at metavoxel resolution) on the corresponding slices (a), (e), (i) shaded in a translucent red; (c), (g), and (k) showing the result of plotting the labels $\hat{V}_{GE}^1, \hat{V}_{GE}^2, \hat{V}_{GE}^3$ in different colors (red, green and blue) back onto the original MR slice at the metavoxel resolution; (d), (h), and (l) showing the result of highlighting the labels $\hat{V}_{LLE}^1, \hat{V}_{LLE}^2, \hat{V}_{LLE}^3$ in different colors (red, green and blue) on the original MR slice at the metavoxel resolution. Note that (c), (g), and (k) were obtained by applying GE within *ManifoldConsensusEmbed*, while (d), (h), and (l) were obtained with LLE within *ManifoldConsensusEmbed*. In each case the cluster shaded red corresponds closely to the potential cancer space.

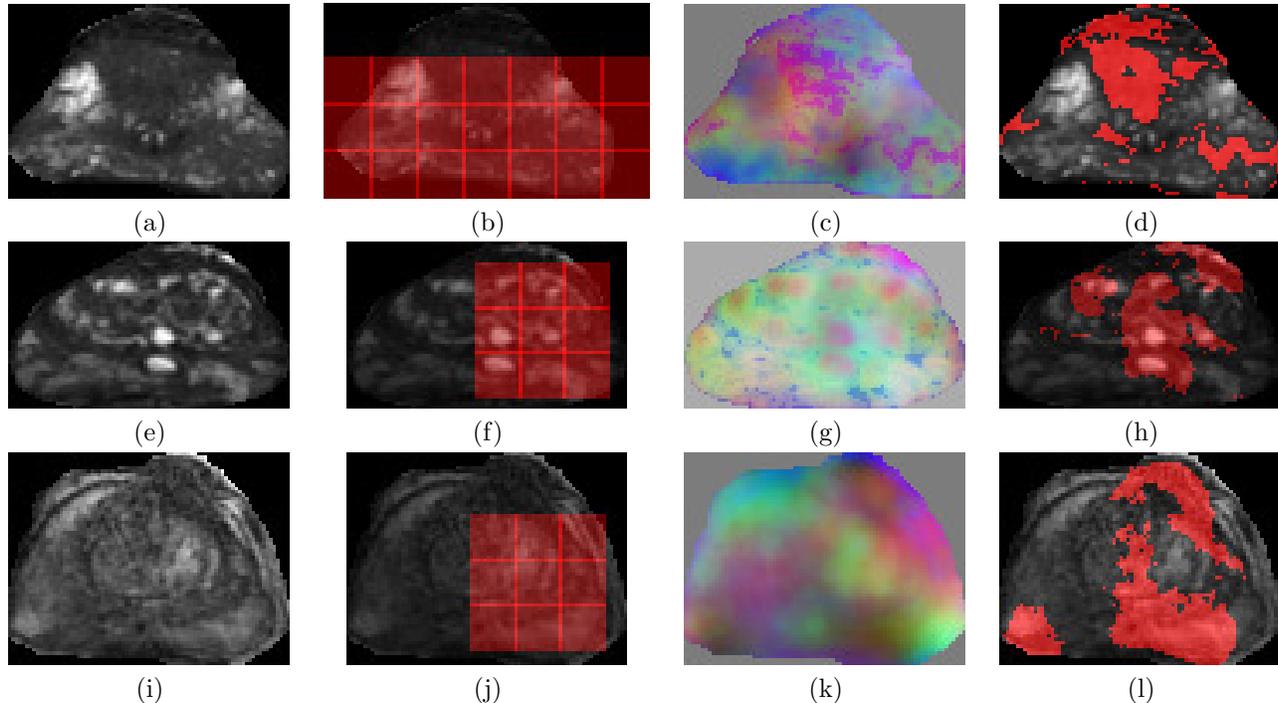


Figure 6: (a), (e), and (i) 2D sections from 3 different endorectal *in vivo* 1.5 T prostate studies; (b), (f), and (j) showing the location of *potential cancer space* (in red) on the corresponding slices (a), (e), (i); (c), (g), and (k) show the spatial mapping of $X_{GE}(c)$ via a RGB colormap onto (a), (e), and (i) respectively; (d), (h), and (l) show the most likely segmentation of cancerous areas (in red) as a result of plotting the label from the cluster that best corresponds to the potential cancer space back onto the slice. Note the correspondence of the red regions in (d), (h), and (l) with the cancer space((b), (f), and (j)).

(i)). Figures 6 (b), (f), (j) show the *potential cancer space* overlaid as a translucent red shading on the slice. The third column (from left to right) shows a novel way of visualizing the low dimensional embedding coordinates of each $c \in C$. Each spatial location $c \in C$ in Figures 6(c), (g), (k) is represented by the 3 principal Eigenvalues in $X_{GE}(c)$. The Eigenvalues have been appropriately scaled so that they can be displayed as RGB values. The results shown in Figures 6 (c), (g), and (k) reveal at least 3 distinct tissue regions via our manifold visualization scheme. Using the algorithm *ManifoldConsensusClust*, we then obtain the final clusters $\hat{V}_{GE}^1, \hat{V}_{GE}^2, \hat{V}_{GE}^3$. The labels corresponding to the cluster with the greatest overlap with the cancer space were plotted back onto the slice. The CaP segmentations (in red) thus obtained are shown at voxel resolution of MR data on Figures 6 (d), (h), and (l). Note that in each case there is significant overlap between the region determined to be cancer via *ManifoldConsensusClust* (Figures 6 (d), (h), and (l)) and the corresponding cancer space (Figures 6 (b), (f), and (j)).

3.2 Quantitative Evaluation

3.2.1 1.5 T MR data at metavoxel resolution

We have already determined the counts of metavoxels lying in the potential cancer space for each dataset (as described in Section 2.2). The counts of the metavoxels in each of the clusters $\hat{V}_{\phi}^1, \hat{V}_{\phi}^2, \hat{V}_{\phi}^3, \phi \in \{GE, LLE\}$ are compared with this ground truth value. Sensitivity, specificity and positive predictive values (PPV) for each of $\hat{V}_{\phi}^1, \hat{V}_{\phi}^2, \hat{V}_{\phi}^3$ are obtained. The cluster with the highest sensitivity, specificity and PPV is then determined as the cancer class. These results are then averaged over 18 datasets and are summarized in Table 1. The results suggest that GE yields a marginally better sensitivity compared to LLE, but LLE has higher specificity and PPV. Note that with both GE and LLE the average detection sensitivity and specificity are 80% or higher.

Resolution	NLDR Method	Sensitivity	Specificity	PPV
Metavoxel	Graph Embedding	84.50	79.40	69.09
Metavoxel	LLE	82.87	80.70	72.10
Voxel	Graph Embedding	92.65	82.06	57.75

Table 1: CaP detection sensitivity, specificity and PPV at the metavoxel resolution averaged over 18 1.5 T MRI datasets using LLE and GE in conjunction with consensus embedding and consensus clustering. Note that on account of memory constraints the LLE algorithm was unable to yield any useful results at the voxel resolution.

3.2.2 1.5 T MR data at voxel resolution

Similar to the manner in which we evaluate the metavoxel resolution results, we have already determined the counts of voxels lying in the potential cancer space for each dataset (as described in Section 2.2). The counts of the voxels in each of the clusters $\tilde{V}_\phi^1, \tilde{V}_\phi^2, \tilde{V}_\phi^3, \phi \in \{GE\}$ are compared with this ground truth value. Sensitivity, specificity and positive predictive values (PPV) for each of $\tilde{V}_\phi^1, \tilde{V}_\phi^2, \tilde{V}_\phi^3$ are obtained. The cluster with the highest sensitivity, specificity and PPV is then determined as the cancer class. These results are then averaged over 12 datasets and are summarized in Table 1. Note that due to the much higher spatial resolution of the data analyzed the performance measures at voxel resolution are higher compared to those at metavoxel resolution.

4. CONCLUDING REMARKS

In this paper we have presented a powerful clinical application of computer-aided diagnosis scheme for automatically detecting CaP from within 1.5 Tesla *in vivo* prostate MR imagery. From a methodology perspective the novel contribution of this paper is a consensus embedding scheme that aims to achieve a stable low dimensional representation of high dimensional data by averaging object adjacencies from across multiple low dimensional data projections. A total of 18 1.5 T *in vivo* MRI datasets were considered in this study. Over 350 3D texture features were extracted and analyzed to distinguish between different prostate regions including malignant and benign tumor areas on these datasets. Further, we qualitatively and quantitatively evaluated our schemes at two different spatial resolutions for the 18 1.5 T datasets. We have also examined the performance of 2 different non-linear dimensionality reduction methods, graph embedding and locally linear embedding, within the framework of our scheme. Our scheme yielded an average sensitivity of 83% and specificity of 79% at the metavoxel resolution, and a 92.65% sensitivity and 82.06% specificity at voxel resolution. Table 2 summarizes the computation times for the different modules of our scheme, estimated from time measurements over 18 3D datasets of average size $256 \times 256 \times 10$ at voxel resolution and average size $7 \times 3 \times 10$ at metavoxel resolution. All programs were developed using MATLAB 7.3 (Mathworks Inc.) and IDL 6.4 (ITT Visual Information Systems) on a Ubuntu Linux platform and executed on a 32 GB RAM, 2 dual core 2.33Ghz 64-bit Intel Core 2 processor machine. We are currently working on extensively evaluating our system on a larger cohort of datasets. A prerequisite to performing a more robust quantitative evaluation of our methodology is the knowledge of precise CaP extent on the MRI studies. We are currently developing a registration methodology to map CaP extent accurately from histology onto MRI for the ACRIN database.

ACKNOWLEDGMENTS

This work was possible due to grants from the Coulter Foundation (WHCF4-29349, WHCF 4-29368), the Busch Biomedical Award, the Cancer Institute of New Jersey, the New Jersey Commission on Cancer Research and

Spatial resolution	Pre-processing	Feature Extraction	Consensus Embedding	Consensus Clustering
Metavoxel	20	20	1.52	0.17
Voxel	20	20	-	240

Table 2: Computation times (in minutes) for each module of the segmentation scheme on 18 3D MR datasets of average size $256 \times 256 \times 10$ at voxel resolution and $7 \times 3 \times 10$ at metavoxel resolution, with the extraction of 373 texture features.

the National Cancer Institute (R21CA127186-01, R03CA128081-01). The authors would like to acknowledge the ACRIN trial for the use of their MRI data. They would also like to thank Dr. James Monaco for useful discussions held on the subject.

REFERENCES

1. M. Schiebler, M. Schnall, H. Pollack, R. Lenkinski, J. Tomaszewski, A. Wein, R. Whittington, W. Rauschnig, and H. Kressel, "Current role of MR Imaging in Staging of Adenocarcinoma of the Prostate," *Radiology* **189**(2), pp. 339–352, 1993.
2. P. Borboroglu, S. Comer, R. Riffenburgh, and C. Amling, "Extensive Repeat Transrectal Ultrasound Guided Prostate Biopsy in Patients with Previous Benign Sextant Biopsies," *The Journal of Urology* **163**(1), pp. 158–162, 2000.
3. Y. Zhu, S. Williams, and R. Zwiggelaar, "Computer Technology in Detection and Staging of Prostate Carcinoma: a Review," *Medical Image Analysis* **10**(2), pp. 178–199, 2006.
4. A. Madabhushi, M. Feldman, D. Metaxas, J. Tomaszewski, and D. Chute, "Automated Detection of Prostatic Adenocarcinoma from High-Resolution Ex Vivo MRI," *IEEE Transactions on Medical Imaging* **24**(12), pp. 1611–1625, 2005.
5. A. Madabhushi, J. Shi, M. Rosen, J. Tomaszewski, and M. Feldman, "Graph Embedding to Improve Supervised Classification and Novel Class Detection: Application to Prostate Cancer," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI 2005)*, **3749**, pp. 729–737, 2005.
6. A. Madabhushi, J. Shi, M. Feldman, M. Rosen, and J. Tomaszewski, "Comparing Classification Performance of Feature Ensembles: Detecting Prostate Cancer from High Resolution MRI," in *Workshop on Computer Vision Methods in Medical Image Analysis (CVAMIA 2006)*, **4241**, pp. 25–36, 2006.
7. I. Chan, W. Wells, R. Mulkern, S. Haker, J. Zhang, K. Zou, S. Maier, and C. Tempany, "Detection of Prostate Cancer by Integration of Line-scan Diffusion, T2-mapping and T2-weighted Magnetic Resonance Imaging; a Multichannel Statistical Classifier," *Medical Physics* **30**(9), pp. 2390–2398, 2003.
8. L. Ibanez, W. Schroeder, L. Ng, and J. Cates, *The ITK Software Guide*. Kitware, Inc. ISBN 1-930934-15-7, second ed., 2005.
9. A. Madabhushi and J. Udupa, "New Methods of MR Image Intensity Standardization via Generalized Scale," *Medical Physics* **33**(9), pp. 3426–34, 2006.
10. G. Lee, C. Rodriguez, and A. Madabhushi, "An Empirical Comparison of Dimensionality Reduction Methods for Classifying Gene and Protein Expression Datasets," in *Bioinformatics Research and Applications (ISBRA 2007)*, **4463**, pp. 170–181, 2007.
11. S. Roweis and L. Saul, "Nonlinear Dimensionality Reduction by Locally Linear Embedding," *Science* **290**(5500), pp. 2323–2326, 2000.
12. J. Shi and J. Malik, "Normalized Cuts and Image Segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **22**(8), pp. 888–905, 2000.
13. P. Tiwari, A. Madabhushi, and M. Rosen, "Manifold Learning and Hierarchical Unsupervised Clustering for Detection of Prostate Cancer from Magnetic Resonance Spectroscopy (MRS)," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI 2007)*, **4792**, pp. 278–286, 2007.
14. J. Venna and S. Kaski, "Local multidimensional scaling," *Neural Networks* **19**(6), pp. 889–899, 2006.
15. A. Strehl and J. Ghosh, "Cluster Ensembles – A Knowledge Reuse Framework for Combining Multiple Partitions," *Journal on Machine Learning Research* **3**, pp. 583–617, 2002.
16. J. Chappelow, A. Madabhushi, M. Rosen, J. Tomaszewski, and M. Feldman, "Multimodal Image Registration of ex vivo 4 Tesla Prostate MRI with Whole Mount Histology for Cancer Detection," in *Proceedings of SPIE Medical Imaging (SPIE 2007)*, **6512**, pp. S1–S12, 2007.
17. A. Fred and A. Jain, "Combining Multiple Clusterings Using Evidence Accumulation," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **27**(6), pp. 835–850, 2005.