Available online at www.sciencedirect.com
SCIENCE DIRECT®
Medical Image Analysis 8 (2004) 139–150
MEDICAL IMAGE ANALYSIS
www.elsevier.com/locate/media

Optimized prostate biopsy via a statistical atlas of cancer spatial distribution

Dinggang Shen [a,b,*], Zhiqiang Lao [a], Jianchao Zeng [c], Wei Zhang [d], Isabel A. Sesterhenn [d], Leon Sun [e], Judd W. Moul [e], Edward H. Herskovits [a], Gabor Fichtinger [b], Christos Davatzikos [a,b]

[a] *Department of Radiology, University of Pennsylvania, 3600 Market Street, Suite 380, Philadelphia, PA 19104-2644, USA*
[b] *Center for Computer-Integrated Surgical Systems and Technology, Johns Hopkins University, MD, USA*
[c] *Department of Electrical and Computer Engineering, Howard University, Washington, DC 20059, USA*
[d] *Armed Forces Institute of Pathology, Washington, DC 20306, USA*
[e] *DoD Center for Prostate Disease Research, Rockville, MD 20852, USA*

Received 23 November 2001; received in revised form 15 May 2002; accepted 3 November 2003

Abstract

A methodology is presented for constructing a statistical atlas of spatial distribution of prostate cancer from a large patient cohort, and it is used for optimizing needle biopsy. An adaptive-focus deformable model is used for the spatial normalization and registration of 100 prostate histological samples, which were provided by the Center for Prostate Disease Research of the US Department of Defense, resulting in a statistical atlas of spatial distribution of prostate cancer. Based on this atlas, a statistical predictive model was developed to optimize the needle biopsy sites, by maximizing the probability of detecting cancer. Experimental results using cross-validation show that the proposed method can detect cancer with a 99% success rate using seven needles, in these samples.

Keywords: Prostate cancer; Statistical atlas; Deformable registration; Image warping; Image normalization; Needle biopsy

1. Introduction

Prostate cancer is the second leading cause of cancer death for American men. When prostate cancer is diagnosed early, it is usually curable. Therefore, it is extremely important to detect prostate cancer at early stages. Currently, common prostate cancer screening methods are digital rectal exam (DRE) and prostate specific antigen (PSA). The combination of these two tests has been shown to be the most cost efficient screening method for prostate cancer (Littrup et al., 1993). When the PSA level is higher than normal or the DRE shows abnormal results, a needle biopsy, guided by transrectal ultrasound images, is typically recommended to help determine whether a tumor really exists and whether the tumor is benign or malignant. Notably, cancer is mostly undetectable in routinely used ultrasound images, because of low tissue contrast and low SNR of the images. Therefore, the transrectal ultrasound images are most often used only to determine the locations of the needles within the prostate. The biopsy sites are typically determined by the physician based on empirical rules. This inevitably leads to a significant number of prostate cancer cases remaining undetected in their initial biopsy. For example, the systematic sextant biopsy protocol (Hodge et al., 1989) is the most common biopsy protocol. However, studies have shown that this protocol results in a positive predictive value of only around 30% (Flanigan et al., 1994). Other clinical

* Corresponding author. Tel.: +1-215-349-8588; fax: +1-215-614-0266.
E-mail addresses: dgshen@rad.upenn.edu (D. Shen), zqlao@rad.upenn.edu (Z. Lao), jzeng@howard.edu (J. Zeng), ehh@rad.upenn.edu (E.H. Herskovits), gabor@cs.jhu.edu (G. Fichtinger), christos@rad.upenn.edu (C. Davatzikos).

1361-8415/$ - see front matter © 2003 Elsevier B.V. All rights reserved.
doi:10.1016/j.media.2003.11.002

studies have also suggested that the sextant technique may not be optimal and have investigated new biopsy protocols that might yield significantly better results (Eskew et al., 1997; Chen et al., 1997, 1999; Daneshgari et al., 1995). Different biopsy strategies have significantly different success rates (Chen et al., 1997), which is likely due to the fact that the spatial distribution of prostate cancer is inhomogeneous (Hodge et al., 1989; Flanigan et al., 1994; Eskew et al., 1997; Chen et al., 1997, 1999; Daneshgari et al., 1995). If the biopsy protocol can be optimized to increase the likelihood of detecting prostate cancer, according to certain objective and quantitative criteria, then significant improvement in diagnostic accuracy should be expected.

Researchers have begun to investigate the possibility of using a large number of patient histopathological images to determine prostate regions that are most likely to develop cancer, and therefore should be sampled during biopsy (Zeng et al., 2000; Chen et al., 1997). However, previously proposed strategies are limited in two respects. First, most of these strategies used very simple spatial normalization methods, such as matching of the prostate bounding boxes, to place the image data into a standardized coordinate system, a step called *spatial normalization*. Spatial normalization is a very important step in this procedure, because it accounts for inter-individual morphological variability and allows the direct superposition and comparison of images from a large group of patients, thereby enabling image statistics to be collected. Inaccuracies in spatial normalization result in significant overlap of *different* prostatic zones across individuals, and therefore introduce error to the calculation of cancer distribution within each zone. This leads to loss of statistical power. Second, most of the current strategies employ relatively simple statistical models of analysis, which are only suitable to solving a small-scale optimization problem. If applied to a large-scale optimization problem, such as optimization that requires fine resolution, those statistical models usually result in suboptimal solutions.

In this paper, we propose a methodology that addresses and partly overcomes both of these limitations. We build upon earlier work published in (Shen et al., 2001b). In order to overcome the problem of inter-individual variability and to render images from different patients directly comparable on a voxel-by-voxel basis, we use *elastically deformable anatomical models* (Shen and Davatzikos, 2000; Shen et al., 2001a), which spatially normalize the prostate images to a canonical coordinate system, thus allowing a voxel-by-voxel analysis. This eliminates the need for reduction of the data resolution via subdivision into relatively large anatomical partitions and improves accuracy of the estimated underlying spatial distribution of cancer. This distribution does not only reflect the probability of developing cancer at individual locations, but also at the spatial correlation of cancer incidence between different prostate regions. This is necessary, because regions between which cancer incidence is highly correlated need not be sampled simultaneously, as opposed to regions between which cancer occurrence is relatively independent. All of these models are used in an optimization framework for estimation of optimal needle biopsy strategy. After all of these processes are completed, we will obtain a statistical atlas, which includes the statistical shape model of the prostate, along with other information such as the locations of optimized biopsy needles and the map of the probability of prostate cancer distribution.

This paper is organized as follows. In Section 2, we provide a brief description on reconstructing prostate samples, using surface modeling. In Section 3, we build a methodology for deformable registration and normalization of the prostate samples, with the goal of removing most of inter-individual morphological variability. In Section 4, we develop a probabilistic method for designing optimal biopsy strategy that best predicts the presence of prostate cancer in patients. In Section 5, we demonstrate the performance of our deformable model in registering prostate samples and creating a statistical atlas. Also, we validate the predictive power of our at-

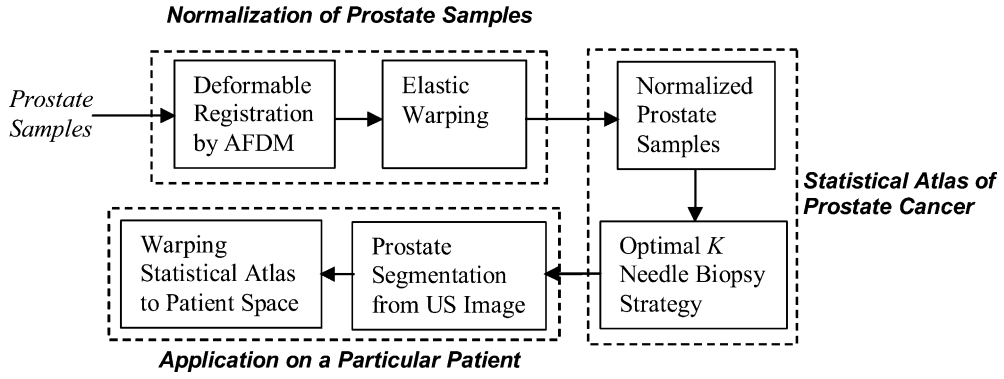

Fig. 1. A schematic representation of the proposed method used for prostate cancer detection.

las-based optimal biopsy strategy in cancer detection by using an existing prostate database, developed by the DoD Center for Prostate Disease Research. Fig. 1 shows the schematic representation of the methodology to be proposed in this paper.

2. Data description

We used a diverse group of male patients with prostate cancer. Each patient was recommended for biopsy based on his PSA level or abnormal digital rectal examination or both. The clinical stages of the detected cancer ranged from T1 to T2 and the PSA levels were below 15 ng/ml.

Three-dimensional (3-D) surfaces of prostate structures were reconstructed from digitized step-sectioned whole-mounted radical prostatectomy specimens with clinically localized cancers using deformable modeling techniques (Zeng et al., 2000). The prostatectomy specimens were first fixed in 10% formalin for 48 h before being step-sectioned into slices at 2.5-mm intervals in transverse planes, resulting in 10–15 slices for each specimen. Each slice was further embedded, sectioned into a 4-μm thickness, whole-mounted on a glass slide, and stained with hematoxylin and eosin before being examined by pathologists under a microscope. The pathologists identified key structures from each section and outlined their boundaries. Structures outlined include surgical margin, capsule, seminal vesicles, urethra, ejaculatory ducts as well as all cancerous tissue. For the purpose of orientation alignment, five external marks were placed, i.e., four at the four corners of each slide and one at the center of the top edge of the slide. In theory, three marks that do not lie on a straight line are enough to register any two slides. Extra marks were taken in this study to average out any potential minor errors. The outlined slides were then digitized at a resolution of 1500 dpi, and slides of a specimen were registered to each other with the help of five manual marks on each slide. Finally, prostate surface models were reconstructed based on the registered outline information. Open surfaces were used for the modeling of the prostate capsules, since the shapes in the prostate ends are not easy to reconstruct due to the use of large slice intervals (2.5 mm), compared to the actual size of the prostate. There exists a slight distortion or mis-registration between the slices used to reconstruct the prostate surface model, reflected as unsmooth prostate surfaces (cf. Fig. 4). This is probably due to the geometrical distortion during slicing, or due to the inconsistent placement of the five manual marks. These inaccuracies should be random, and will not significantly influence the quality of the statistical map of cancer distribution, since they are smoothed out by our elastic warping procedure.

3. Spatial normalization of the prostate samples

Spatial normalization is a very important step, because it accounts for inter-individual morphological variability and allows the direct superposition and comparison of images from a large group of patients, thereby enabling image statistics calculation. A major problem in developing a spatial normalization method is determining correspondences, i.e., developing a mechanism that establishes a map between anatomically corresponding points or regions. This correspondence problem can be solved via elastically deformable anatomical models (Kass et al., 1988), which have been explored extensively in several fields, including medical imaging (Pizer et al., 2001; Joshi et al., 2002; McInerney and Terzopoulos, 1996; Staib and Duncan, 1992; Davatzikos, 1997; Brejl and Sonka, 2000). In particular, deformable models have been widely used to register brain images across individuals and to construct statistical atlases. We have developed a deformable shape modeling framework, for segmentation and reconstruction of anatomical shapes, and for determining morphology-based correspondence across individuals, from tomographic images (Shen and Davatzikos, 2000; Shen et al., 2001a). This framework is based on an *adaptive focus deformable model* (AFDM). In AFDM, for a given set of structures, a shape model is first constructed to represent a typical shape of these structures. This shape model includes two kinds of information: information about the geometry of the structures and information about the statistical variation of these structures within a given population. Then in the application stage, the deformable shape model is placed in an image with the structures of interest and is subsequently let free to deform according to features extracted from the subject images, seeking objects that have similar geometry, but also objects that fall within the expected range of shape variation (2 SD from the shape mean in the direction of each eigenvector). AFDM is a fully automated procedure for establishing the correspondences between the points in the model and the points in an individual's anatomy. With these established correspondences, we can use an elastic warping technique (Davatzikos, 1997) to transform the subject's images into a *stereotaxic space*, i.e., the canonical coordinate system within which the model resides. This procedure is called spatial normalization, and it is the cornerstone of analysis of images across a population. It accounts for inter-individual anatomical variability and allows one to compare images from different patients point by point, in the stereotaxic space, thus enabling the full utilization of the image data resolution within the accuracy of the registration procedure.

In this section, we will use AFDM as a registration method to spatially normalize the external shape and internal structures of the prostate samples, such as the

capsule and the urethra. In spatial normalization, it is important to select an appropriate model that will not bias any of the training samples. By removing the linear differences among the subjects, a good model should be the average of the linearly aligned samples. In this way, the selected model will be at the statistical center of all samples, thereby making the deformation of model to samples much easier. In this paper, we selected one typical prostate as a model.

Many deformation algorithms have been proposed to propagate the deformations from voxel to voxel, such as viscoelastic models (Christensen et al., 1996). However, those models have two limitations relative to surface-based models. First, they are computationally very demanding, which often limits their use in surgical environments. Second, they do not include shape information (i.e., curvatures of boundaries at different scales). In this paper we use the surface-based model to warp the prostate samples to the space of the selected model. The warping procedure is performed in three stages. First, AFDM is used to reconstruct the shape of each structure and to determine point correspondences between the subjects and the model. Second, these point correspondences are interpolated elsewhere in the space of the prostate by using an elastic warping technique (Davatzikos, 1997). The tumor regions are warped according to the correspondences that are established between the subjects and the model in the whole volume. Although surface-based registration can potentially be less accurate than voxel-based registration, its accuracy depends on the number of surfaces used to drive the warping.

We now describe AFDM and its adaptation to this particular application.

3.1. Spatial connections for the surfaces in the prostate model

Adaptive focus deformable model comprises several inter-connected surfaces, each representing the anatomy of a structure of interest. For example, in the prostate model, there are two surfaces that represent the boundaries of the capsule and the urethra, i.e., S_C (for capsule) and S_U (for urethra). From the model given in Fig. 2(a), we can observe that these two surfaces, S_C and S_U, are not connected with each other, since they were reconstructed separately as two different surfaces. That is, no mesh connections exist between the points in S_C and the points in S_U. Therefore, any deformation on the capsule will not propagate to the urethra during segmentation, since the urethra is not connected with the capsule. A more robust segmentation method can be obtained if the capsule and the urethra interact with each other during deformation. In order to achieve this, we spatially connect those separated surface patches (as gray arrows in Fig. 2(b)), to obtain a single surface

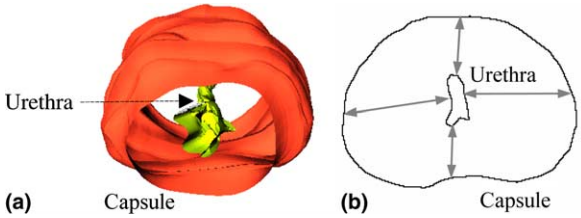

Fig. 2. A 3-D prostate model with two surfaces that represent the capsule and the urethral boundaries, respectively. (a) A 3-D display of the prostate surface model and (b) a cross-section display of the prostate surface model. The grey arrows denote that these two separate surfaces are connected and thus deformed together.

model. In particular, for every point in the urethra surface of the model, we search for its closest point in the capsule surface of the model and connect these two points as immediate neighbors to each other. Therefore, any deformation in one part of the model will propagate to all other parts of the surface model, and all prostate structures will deform jointly.

3.2. Affine-invariant attribute vector for each model point

In order to capture the geometry of anatomical structures in a hierarchical fashion, we introduced the concept of an attribute vector that is attached to each point of the prostate surface model and reflects the geometric structure of the model from a global and coarse scale to a local and fine scale.

Each attribute is defined as the volume of a tetrahedron (see Fig. 3(a)), formed by a model point V_i and any three points in its certain neighborhood layer. For each point V_i in the model, the first *neighborhood layer* includes its immediate neighbors. The second neighborhood layer includes the immediate neighbors of the immediate neighbors, and so forth. The neighborhood layers are constructed so that no point is repeated twice in the neighborhood of another point. The volume of the tetrahedron, formed by the immediate neighbors, reflects the local structure of the surface around the point V_i. The volumes of larger tetrahedra represent the more global properties of the surface around the point V_i. Accordingly, the attribute vector corresponding to, say, a high-curvature region is completely different from attribute vectors of flat segments of the surface. More importantly, even points of similar curvatures might have very different attribute vectors, depending on the number of neighborhood layers used, or equivalently, the number of components of the attribute vector. The attribute vectors are an important aspect of AFDM, since they provide a means for finding correspondences across individuals by examining the similarity of the underlying attribute vectors. Fig. 3(b) gives an example of the selected correspondences that were established between the model and a subject during the spatial

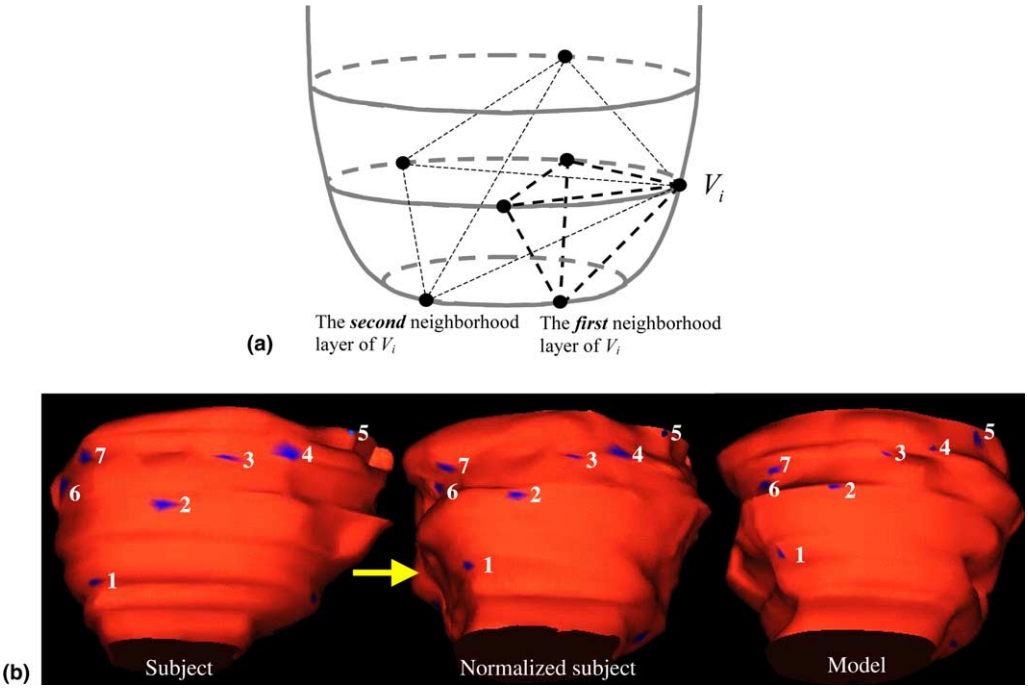

Fig. 3. Definition of the attribute vector and the use of attribute vectors for the correspondence detection during spatial normalization. (a) The attributes of the point V_i are defined as the volumes of the tetrahedra in the different neighborhood layers around V_i, reflecting the geometric structure of the prostate from a global and coarse scale to a local and fine scale. (b) A set of selected correspondences, which were automatically determined between the model and one subject during the spatial normalization procedure. The numbers in (b) are used to represent the corresponding points.

normalization procedure. The corresponding points are assigned the same numbers in the three images.

Attribute vector $F(V_i)$ of a model point V_i contains the volumes of tetrahedra that are calculated in the different neighborhood layers. Thereby, it captures different levels of shape information for the segment of surface around the point V_i. The number of the attributes and, equally, the number of neighborhood layers, can be adaptively changed with the progress of the deformation. We use a large number of attributes, i.e. 30, for the initial deformation stages and then decrease this number gradually. In our case, the initial 30 neighborhood layers cover about 1/3 of the model surface. Notably, the attribute vector $F(V_i)$ can be further made affine-invariant by normalizing it over the entire surface of the model, i.e.

$$\hat{F}(V_i) = \frac{F(V_i)}{\sum_i \|F(V_i)\| / \sum_i 1.0},$$

where $\|F(V_i)\|$ is the magnitude of the vector $F(V_i)$ and $\hat{F}(V_i)$ is the affine-invariant attribute vector for the ith model point V_i.

3.3. Energy function of the prostate surface model

Deformable models integrate model-driven and data-driven analysis through the deployment of an energy function and a set of regularization parameters. The energy that our deformable model minimizes is defined as follows:

$$E = \sum_i \omega_i E_i = \sum_i \omega_i \left(E_i^{\text{model}} + E_i^{\text{data}}\right). \tag{1}$$

The weighting parameter ω_i determines the relative weight given to the local energy term E_i, which is defined for the ith model point V_i in the surface model. The energy term E_i is composed of two terms, E_i^{model} and E_i^{data}.

The model energy term E_i^{model} is defined to allow AFDM to determine correspondences, in addition to segmenting structures of interest. In particular, the model energy term E_i^{model} is defined as the difference between the attribute vectors of the model and its deformed configuration at the point V_i, and it is given by

$$E_i^{\text{model}} = D\left(\hat{F}^{\text{Def}}(V_i), \hat{F}^{\text{Mdl}}(V_i)\right),$$

where $\hat{F}^{\text{Def}}(V_i)$ and $\hat{F}^{\text{Mdl}}(\mathrm{V}_i)$ are, respectively, the normalized attribute vectors of the deformed model configuration and the model at the point V_i. $D(\cdot)$ is a similarity definition. In this paper, we use a simple definition for the difference between the two attribute vectors, that is,

$$D\left(\hat{F}^{\text{Def}}(V_i), \hat{F}^{\text{Mdl}}(V_i)\right) = \sum_l \delta_l \left(\hat{F}_l^{\text{Def}}(V_i) - \hat{F}_l^{\text{Mdl}}(V_i)\right)^2,$$

where $\hat{F}_l^{\text{Def}}(V_i)$ and $\hat{F}_l^{\text{Mdl}}(V_i)$ are the lth component of the attribute vectors $\hat{F}^{\text{Def}}(V_i)$ and $\hat{F}^{\text{Mdl}}(V_i)$, respectively. The parameter δ_l denotes the weighting for the lth attribute $\hat{F}_l^{\text{Def}}(V_i)$. We use small weights for the large attributes

that are calculated from the higher neighborhood layers, in order to avoid large attributes dominating the calculation of the difference of the attribute vectors.

The data energy term E_i^{data} is defined for the model point V_i, and it is designed to move the deformable model towards an object boundary. Since our deformation mechanism *deforms a segment* of *surface* around each model point V_i at a time, we design a data energy term E_i^{data} that reflects how each segment of surface, rather than a single model point, fits with the subject boundaries. This deformation mechanism provides some degree of robustness to spurious or disconnected edges. Suppose that H_i is a segment of surface around the model point V_i. The size for the segment of surface H_i is large in the initial deformation stages, and it is designed to decrease gradually with the progress of deformation. For every point $\mathbf{x}$ on H_i, we calculate its shortest distance $d(\mathbf{x})$ to the corresponding surface in the subject. Notably, all the prostate boundaries, such as capsule and urethra, have been reconstructed and labeled as separate surfaces, using the procedures described in Section 2. In this way, for each model point such as a point in the capsule, we are able to estimate its shortest distance to the corresponding surface of subject, such as the capsule of subject. The data energy term E_i^{data} is defined as follows:

$$E_i^{\text{data}} = \sum_{\mathbf{x}\in H_i} \eta_{\mathbf{x}} d(\mathbf{x}),$$

where $\mathbf{x}$ is a point on the segment of surface H_i and $\eta_{\mathbf{x}}$ is a weighting parameter.

3.4. Adaptive-focus deformation strategies

In our previous work we have determined that model adaptivity is very important for robust segmentation and correspondence estimation. We explore the utility of adaptive modeling further in the registration of prostate samples (see Fig. 4). Two types of *prior knowledge* about prostate shapes are available. First, the mesh structures in the model points along the two open surface boundaries are very different from the mesh structures in other internal model points, such as points in the middle cross-section of the prostate model. This is because the boundary points have unpaired mesh edges, while all other points have the paired mesh edges. With this prior knowledge, we can focus on the two surface ends of the prostate shape model in the initial deformation stages, which leads to a rough match between the prostate model and the individual prostate. Second, the prostate capsule is more accurately and reliably outlined in the histopathological samples, while the urethra, which does not include soft tissue, is very deformed and more difficult to segment. By using this priori knowledge of reliability, AFDM approach initially focuses on deforming the capsule surface of the model, a relatively more reliable anatomical part, to the capsule of the individual, while the urethral surface of the model simply follows this deformation by the forces from the capsule surface and also the constraints from shape statistics. After several iterative deformations, both capsule and urethral surfaces of the model are close to their counterparts in the individual. At that point, the urethral surface of the model, a relatively less reliable anatomical part, starts to actively search for its counterpart in the individual's images. The use of these adaptive strategies provides some degree of robustness of our approach in registering prostate samples.

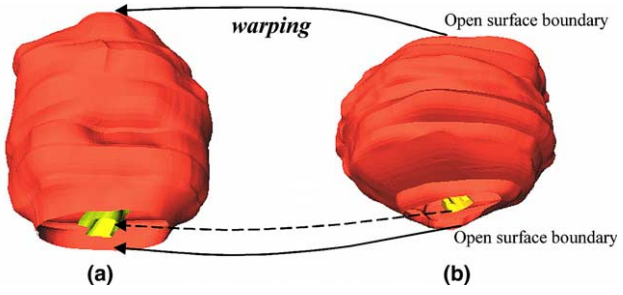

Fig. 4. The procedure of warping an individual's prostate to the prostate model. (a) A side view of the prostate model that is given in Fig. 2; (b) a side view of the prostate of an individual, whose front view is displayed in Fig. 5(b1).

4. Optimal needle biopsy strategy

After warping and normalizing all prostate subjects into the space of the prostate model by AFDM that is described in Section 3, we construct the 3-D statistical atlas of spatial distribution of prostate cancer and further design optimal needle biopsy strategy for diagnosing prostate cancer. Our optimal biopsy strategy reflects the statistical properties of the spatial distribution of cancer. In the following, we will describe an algorithm to design optimal K-biopsy (biopsy with a number K of needles) strategy by minimizing the probability of missing the cancer.

4.1. Method description

A K-biopsy strategy can be optimized by maximizing the probability that at least one needle detects cancer or equivalently by minimizing the probability that none of the K needles detects cancer. The latter probability is defined by

$$P(B(x_i) = \text{NC},\ i = 1, \ldots, K), \tag{2}$$

where $B(x_i)$ is the biopsy outcome of the ith needle at the location x_i and NC denotes a negative cancer detection result. The location x_i can be anywhere inside of a 3-D prostate model. We can find the configuration that

minimizes this probability by using standard optimization methods. Since the probability in (2) is likely to have many local minima, we use the simulated annealing technique to find globally optimal values for $x_1, \ldots, x_k$. We start with an initial guess for the coordinates of the K needles, and then we iteratively change these values in a direction that decreases the probability function in (2). Initially, changes in the direction that increases the probability in (2) are allowed, but these steps are progressively discouraged more and more as the algorithm proceeds, as customary in random optimization methods. In the following section, we will give a fast heuristic method for determining a good initial guess of the coordinates of the K needles.

4.2. A fast heuristic method

The principle of the heuristic optimization method is to place the needles in the regions where the cancer developments are almost independent each other. In this way, for the same cancer detection rate, the number of the needles that are used for biopsy can be minimized. In the following, we describe a very fast heuristic method. The result of this method will be used as an initial guess of the coordinates of the K needles, which will be iteratively optimized later by the simulated annealing technique.

The probability of missing the cancer can be expressed as a product of conditional probabilities,

$$\begin{aligned} &P(B(x_i) = \mathrm{NC},\ i = 1, \ldots, K) \\ &\quad = P(B(x_1) = \mathrm{NC}) \\ &\qquad \times P(B(x_2) = \mathrm{NC} | B(x_1) = \mathrm{NC}) \times \cdots \\ &\qquad \times P(B(x_K) = \mathrm{NC} | B(x_i) = \mathrm{NC},\ i = 1, \ldots, K-1). \end{aligned} \tag{3}$$

Our heuristic method sequentially minimizes a series of the conditional probabilities of missing cancer, each of them being one of the terms in (3). Suppose that there are N different locations in the prostate model and M prostate samples in the training set. To minimize (3), one of K biopsies, let us say the first biopsy $B(x_1)$, will be taken from the location x_1 where the likelihood of cancer is the highest, that is, the first term in (3), $P(B(x_1) = \mathrm{NC})$, is the lowest. Knowing the location of the first biopsy x_1, in order to calculate the conditional probability $P(B(x_2) = \mathrm{NC} | B(x_1) = \mathrm{NC})$, we remove those prostate samples that have cancer at location x_1, since those do not satisfy the condition $B(x_1) = \mathrm{NC}$, and recalculate the probability at each location. It is important to note that if the incidence of cancer at a location other than x_1 is strongly related with the incidence of cancer at location x_1, then the cancer occurrence probability of this location will become very low in the conditional probability $P(B(x_2) = \mathrm{NC} | B(x_1) = \mathrm{NC})$. This is because all subjects with cancer at x_1 have been excluded in calculating the conditional probability $P(B(x_2) = \mathrm{NC} | B(x_1) = \mathrm{NC})$. This implies that the second biopsy will most likely not be placed at the cancer locations strongly related with the location x_1. With this new conditional probability, the biopsy location x_2 can be determined by selecting the location where the likelihood of cancer is highest, i.e., the conditional probability $P(B(x_2) = \mathrm{NC} | B(x_1) = \mathrm{NC})$ is lowest. Using the same procedure, the locations of other biopsy sites can similarly be determined. Effectively, this procedure minimizes each of the terms in (3) sequentially, rather than operating at the full joint distribution, and is therefore extremely fast. The result of this heuristic optimization method is used as an initial guess for the coordinates of the K needles, and then the simulated annealing technique is used to find globally optimal values for $x_1, \ldots, x_k$.

4.3. Spatial smoothing of statistical atlas

In surgical planning, the prostate model that is used to define a standard coordinate system for reporting cancer locations is adapted to the individual morphology of a patient, thereby transferring the statistical atlas of cancer locations to the patient's images native space. This adaptation leads to a patient-specific atlas for image-guided needle biopsy (Grimm et al., 1994). Ultrasound images are the popular image modality for this kind of adaptation. So far, we have developed a statistical shape model for the segmentation of the prostate boundaries from the 2-D ultrasound images (Shen et al., 2003). This method is currently being extended to the 3-D case.

The performance of the actual needle biopsy in the individual patients is related to at least three errors, i.e., the errors in reconstructing atlas and prostate surface models, the unavoidable errors in registration and adaptation of the model to an individual patient's ultrasound volume, and the errors in placing needles inherent in the tracking devices. In order to make our biopsy strategy robust to such errors, each cancer probability (or conditional cancer probability) function is spatially smoothed prior to finding optimal needle locations. This is also designed to avoid the cases that optimal needle locations are determined to be on isolated, high-probability regions, since those regions might be difficult to accurately sample in practice due to unavoidable errors in the model registration and the needle placements. With this formulation, our biopsy strategy becomes more robust, since it places needles on regions with both high cancer probability and relatively wider spatial extent. Therefore, even if a needle is placed in the neighborhood of (but not on) the expected location, due to instrumentation errors or registration errors, our method would still have a high probability of detecting

cancer. In this paper, we used an isotropic Gaussian filter of 1 mm, to smooth the statistical atlas of prostate cancer distribution.

5. Experiments

Two groups of experiments are described in this section. The first set of experiments is provided to demonstrate the performance of our deformable registration. The second set of experiments is used to validate the predictive power of our atlas-based optimal biopsy strategy in detecting prostate cancer. These two sets of experiments were performed on 100 prostate subjects from the CPDR database.

5.1. Performance of deformable registration

Several experiments are provided here to demonstrate the performance of our 3-D deformable registration and warping algorithm. All the results shown below are done in 3-D. Fig. 5 demonstrates a procedure of registering and warping a representative prostate specimen image (Fig. 5(b1)) to the space of the prostate model (Fig. 5(a1)). The side views of the model and the subject are shown in Fig. 4. Fig. 5(a2) shows a middle cross-section of the model prostate. The corresponding cross-section of the subject's prostate is shown in Fig. 5(b2). Notable are the shape differences between these two cross-sections. After using the warping algorithm, we obtain the subject's warped image in Fig. 5(c2), whose shape is very similar to that of the model prostate in Fig. 5(a2). In Fig. 6, we provide the side views of 10 typical training subjects, before and after AFDM-based deformable registration. The pictures in the first row of Figs. 6(a) and (b) are the original prostate subjects, while the pictures in the second row of Figs. 6(a) and (b) are the spatially normalized versions of the prostate subjects. Here, we can easily observe that the variations across the prostate subjects are very large, before deformable normalization. Therefore, the procedure of deformable registration is important, by observing the now-very-similar normalized shapes in the second rows of Figs. 6(a) and (b). In Fig. 7, we provide the middle cross-sections of those 10 subjects, before and after the deformable normalization. Also, we can observe that the slices are very similar when finally deformed.

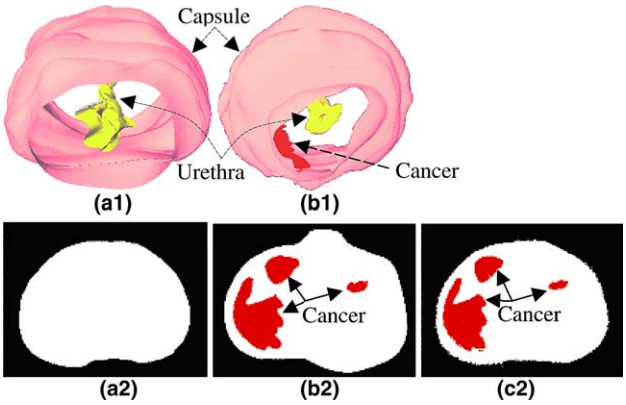

Fig. 5. Results of the deformable registration technique on prostate subjects. (a1) A model of the prostate to which all subjects' images are warped, (b1) a selected subject. (a2) and (b2) show the representative cross-sections, respectively, taken from (a1) and (b1). (c2) is the normalized version, roughly corresponding to (b2).

5.2. Predictability of our statistical atlas

Using the registration and warping algorithm, we can eliminate the overall shape differences across individuals. In this way, we can find the spatial distribution of cancer within the space of the prostate model, which can be used to determine the needle biopsy strategy. We tested our needle optimization method on 100 subjects ($M = 100$). All $N = 256 \times 256 \times 124$ voxels were considered to be candidate biopsy locations. In Fig. 8, the optimal biopsy sites are shown as white spheres and the prostate capsule is shown as red. The underlying spatial statistical distribution of cancer inside of prostate capsule is shown as green. Brighter green indicates higher likelihood in finding cancer in that location. It can be observed that tumors are not symmetrically distributed, as indicated in (Donohue and Miller, 1991). Seven needles were adequate to detect the tumor 100% in those 100 subjects. In particular, the first five needles can detect the tumor 97%. An important implication is that the optimized needle placement is not necessarily in regions that have a high likelihood of cancer. As we can see from Fig. 8, only first *three* white needles were placed in brighter green (high likelihood) regions. The remaining *four* were placed in regions that were almost statistically independent from the first three. Furthermore, Fig. 9 gives more information about the locations of the seven needles, with the needles shown overlaid on cross-sections of the cancer distribution atlas. In Fig. 9, the white spheres/points denote the locations of the needles, while the white curves are the capsule boundaries of the prostate model. It can be easily observed that two needles are placed near the apex of the prostate, three needles in the middle of the prostate, and two needles near the base of the prostate.

We validated the predictive power of our statistical atlas in several ways. *First*, we used the leave-one-out method to validate it. For each time, we selected one subject from our 100 prostate samples and we regenerated the statistical atlas from the remaining 99 samples. We then determined again the optimal biopsy sites, and we applied them to this left-out subject. We repeated this procedure for 100 times, equal to the number of the prostate samples. This way, we measured the probability of missing the cancer. For the 7-biopsy strategy, the

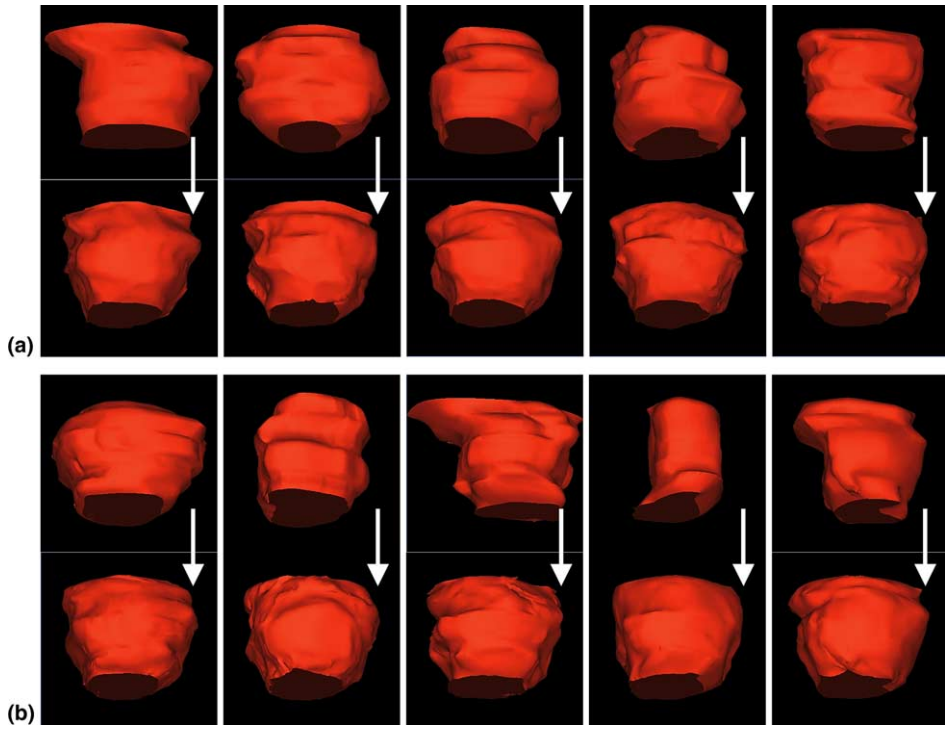

Fig. 6. Side views of 10 selected prostate subjects, before and after deformable registration. The images in the first row of (a) and (b) are the original prostate subjects, while the images in the second row of (a) and (b) are the spatially normalized prostate subjects. The arrows here indicate the procedure of deformable registration.

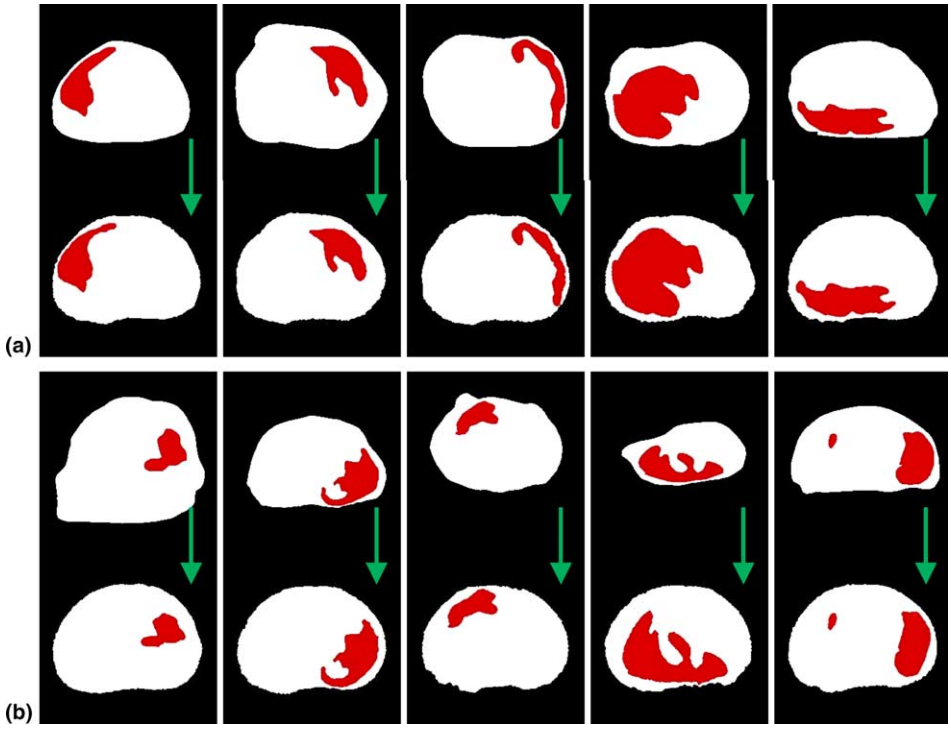

Fig. 7. The middle cross-sections of 10 prostate subjects, before and after normalization. These cross-sections are extracted from those 10 subjects in Fig. 6 and are displayed in the same order.

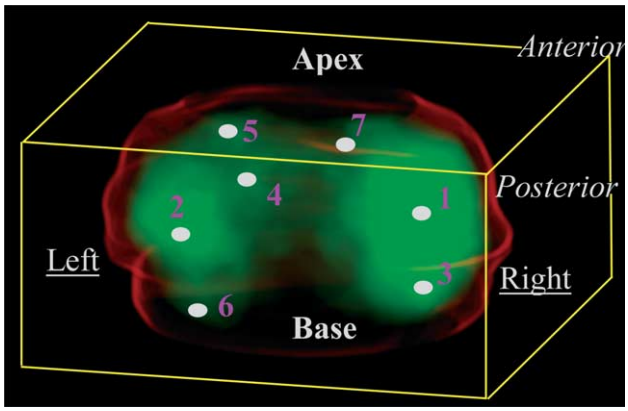

Fig. 8. Optimal biopsy strategy determined by 100 samples. The seven biopsy positions are shown as white spheres, with the statistical atlas of cancer shown as green and prostate capsule shown as red.

leave-one-out method showed that the rate of success of locating cancer was 99%. *Second*, we evaluated the predictability of our statistical atlas as the way described in Table 1. For example, case 1 used 20 samples for training and the other 80 samples for testing. In this case, five needles were adequate to detect the tumor 100% in those 20 training samples, while these five optimal biopsy sites successfully detected cancer in 76 out of 80 testing samples. Case 2 used 64 samples for the training purpose and the other 36 samples for testing. In this case, six needles were adequate to detect cancer 100% in those 64 training samples, while these six optimal biopsy sites successfully detected cancer in 35 out of 36 testing samples. The positions of the optimal biopsy sites, determined by 20, 64 and 100 samples, respectively, are displayed in Fig. 10. From this figure, we can observe that the locations of the six needle sites determined by 64 samples are very similar to those determined by 100 samples. This indicates that, after a sufficient number of prostate samples have been used for

Table 1
Testing the predictability of the statistical atlas by using different number of training samples

	No. of training samples	No. of testing samples	No. of optimal biopsy sites	Success rate
Case 1	20	80	5	76/80
Case 2	64	36	6	35/36

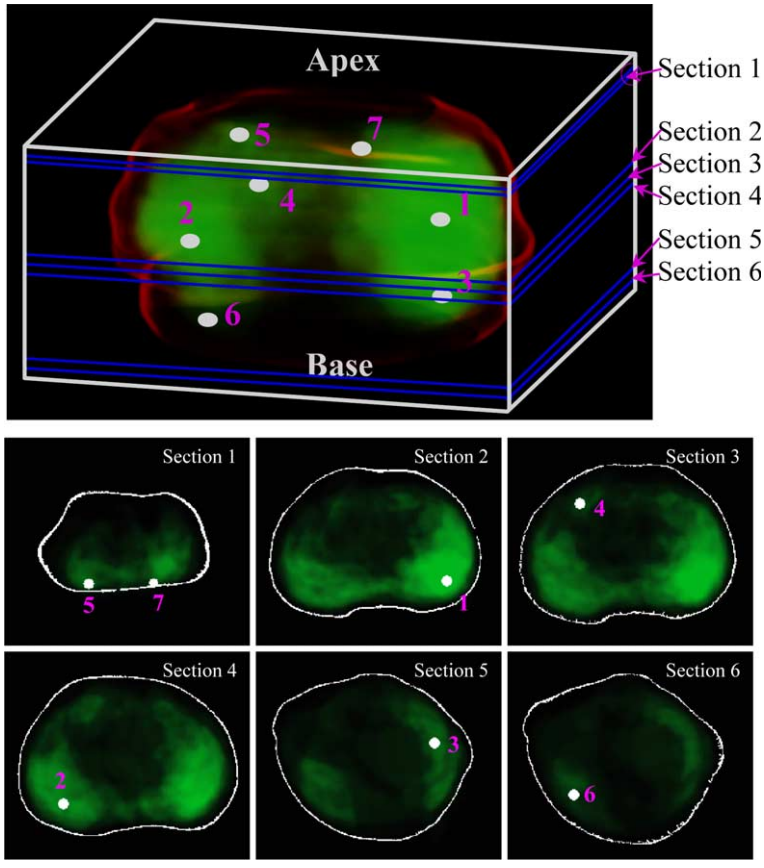

Fig. 9. Optimal needle positions overlaid on cross-sections of the probability of cancer, from the apex to the base of prostate model. The top image displays the statistical atlas and the capsule of the prostate model in the 3-D. There are two needles near the apex, three needles in the middle, and two needles near the base. Notice that the posterior part of the prostate is displayed in the front. Here, white spheres/points denote the locations of the needles and white curves are the capsule boundaries of the prostate model. Notably, these optimal biopsy sites are determined by 100 samples.

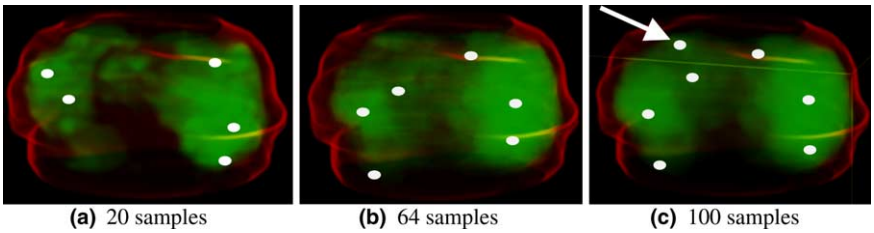

Fig. 10. The probability distribution of cancer development and optimal needle biopsy sites, changing with the number of samples used to construct the atlas. The optimal biopsy sites that are determined for detecting cancer 100% are displayed as white. (a) 20 training samples, requiring five biopsy sites for 100% cancer detection rate; (b) 64 training samples, requiring six optimal biopsy sites; (c) 100 training samples, requiring seven optimal biopsy sites, which are the exactly same as those in Figs. 8 and 9. Notably, the probability of cancer distribution changes less after enough samples are used. For example, the locations of six optimal needles determined by 64 samples are very similar to those determined by 100 samples, except adding one additional needle that is indicated by an arrow in (c). This also shows the robustness of the optimal biopsy sites after sufficient samples have been used.

training, the calculated statistical atlas of the cancer distribution becomes relatively reliable.

We also evaluated the success rates of cancer detection, using the tissue volume that the needle actually extracts from the expected locations. For example, in Fig. 11 the black region represents the tissue volume that a certain needle is expected to sample. In reality, due to the errors in needle placement and image registration, only the volumes represented by the grey region in Fig. 11 are sampled. That is, only a part of the target volume is extracted. Suppose that the volume extracted by a needle is V. In Fig. 12, we provide the rates of success, as a function of the number of the needles (K) and the partial tissue volume that the needle actually extracts from the expected locations. If each needle actually extracts 80–100% of V from its optimized location, then identical curves of success rates are obtained, shown as the curve with '●'. Seven needles are adequate to detect the tumor 100%, while five needles can detect the tumor in 97 out of 100 cases, and six needles can detect the tumor in 99 out of 100 cases. For other percentages of actual vs. theoretical volume extraction, i.e., 30∼70%, 20% and 10%, the rates of success with the number of the needles are shown as curves with '■', '▲' and '◆'. Importantly, this numerical figure of success rate can be valuable for clinicians using our method, since they will be able to quantitatively evaluate the trade-off between biopsy success rate and patient discomfort.

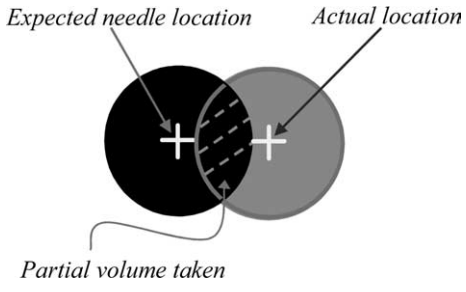

Fig. 11. Schematic description of the actual tissue volume that is taken from the expected locations. The black region represents the expected tissue volume that a certain needle should take. The grey region is the actual tissue volume that is taken in reality from the prostate. That means that only a partial volume of the necessary tissue, as shown by the overlapped region, is taken from the expected location.

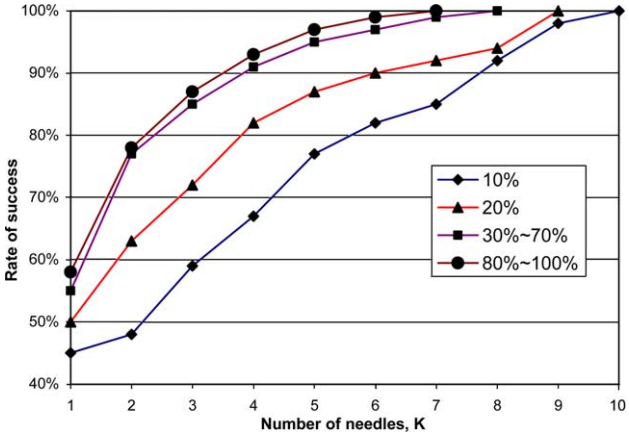

Fig. 12. The rates of success, as a function of the number of needles K, and different percentages of tissue volume that the needle actually extracts from the expected locations.

We also compared our results with the method proposed in (Zeng et al., 2000). That method emphasizes a representation that a physician is familiar with and requires left–right symmetry in biopsy locations. The major limitation of this method is that the registration method is less precise because it uses the bounding box, and the needle optimization is performed only on a resolution of prostate zones. Our method is based on voxel-to-voxel registration and analysis techniques. It eliminates the need to reduce of the data resolution, by subdividing the segments of the prostate images into relatively large anatomical partitions. Therefore, the accuracy of prostate registration in our method is higher relative to that in (Zeng et al., 2000). Also, the success rate of cancer detection is higher. Tested on the same datasets, except more of them, the method in (Zeng et al., 2000) had detection rate 79.3% using six needles,

82.9% using eight needles, and 85.5% using 10 needles (Zeng et al., 2000). Importantly, these results reflected detection rates in the training samples, since the leave-one-out method was not applied in (Zeng et al., 2000).

6. Summary and conclusion

We presented a methodology for constructing a statistical atlas of prostate cancer and for using this atlas to determine an optimal needle biopsy strategy. We used 100 prostate samples provided by the DoD Center for Prostate Disease Research. Using cross-validation, we demonstrated that 99% of the cases are detected with seven needles, which represents a higher detection rate than previous techniques for needle placement optimization.

Several future extensions of this methodology are possible. In particular, we are working toward adding more prostate structures to the prostate model and adding more prostate subjects to the predictive model. Currently, there are over 300 prostate samples available in the same database, which we will use in the future to augment our current statistical model. Moreover, our current and future work involves techniques for deformable registration of the prostate cancer atlas with ultrasound images of the patients obtained during biopsy, which will allow us to adapt the optimal needle locations to the individual patient's anatomy. Finally, extensions of our work will incorporate physical constraints in the optimizer, related to the possible entry points and angling of the biopsy needle. These constraints are imposed by practical issues of the clinical application of the method.

Acknowledgements

This work was supported in part by a grant from the National Science Foundation to the Engineering Research Center for Computer Integrated Surgical Systems and Technology. J. Zeng was supported in part by The Whitaker Foundation.

References

Brejl, M., Sonka, M., 2000. Object localization and border detection criteria design in edge-based image segmentation: automated learning from examples. IEEE Trans. Med. Imag. 19 (10), 973–985.

Chen, M.E., Troncoso, P., Johnston, D.A., Tang, K., Babaian, R.J., 1997. Optimization of prostate biopsy strategy using computer based analysis. J. Urol. 158 (6), 2168–2175.

Chen, M.E., Troncoso, P., Tang, K., Babaian, R.J., Johnston, D., 1999. Comparison of prostate biopsy schemes by computer simulation. Urology 53 (5), 951–960.

Christensen, G., Rabbitt, R.D., Miller, M.I., 1996. Deformable templates using large deformation kinematics. IEEE Trans. Image Process. 5 (10), 1435–1447.

Daneshgari, F., Taylor, G.D., Miller, G.J., Crawford, E.D., 1995. Computer-simulation of the probability of detecting low-volume carcinoma of the prostate with 6 random systematic core biopsies. Urology 45 (4), 604–609.

Davatzikos, C., 1997. Spatial transformation and registration of brain images using elastically deformable models. Comput. Vis. Image Understanding 66 (2), 207–222.

Donohue, R.E., Miller, G.J., 1991. Adenocarcinoma of the prostate: biopsy to whole mount. Denver VA experience. Urol. Clin. North. Am. 18, 449–452.

Eskew, A.L., Bare, R.L., McCullough, D.L., 1997. Systematic 5-region prostate biopsy is superior to sextant method for diagnosing carcinoma of the prostate. J. Urol. 157, 199–202.

Flanigan, R.C., Catalona, W.J., Richie, J.P., et al., 1994. Accuracy of digital rectal examination and transrectal ultrasonography in localizing prostate cancer. J. Urol. 152, 1506–1509.

Grimm, P.D., Balsko, J.C., Ragde, H., 1994. Ultrasound guided transperineal implantation of iodine 125 and palladium 103 for the treatment of early stage prostate cancer. Atlas Urol. Clin. North Am. 2, 113–125.

Hodge, K.K., Mcneal, J.E., Terris, M.K., Stamey, T.A., 1989. Random systematic versus directed ultrasound guided transrectal core biopsies of the prostate. J. Urol. 142, 71–74.

Joshi, S., Pizer, S., Fletcher, P.T., Yushkevish, P., Thall, A., Marron, J.S., 2002. Multi-scale deformable model segmentation and statistical shape analysis using medial descriptions. IEEE Trans. Med. Imaging 21 (5), 538–550.

Kass, M., Witkin, A., Terzopoulos, D., 1988. Snakes: active contour models. Int. J. Comput. Vis. 1, 321–331.

Littrup, P., Goodman, A., Mettlin, C., 1993. The benefit and cost of prostatic cancer early detection. CA Cancer J. Clin. 43, 134–149.

McInerney, T., Terzopoulos, D., 1996. Deformable models in medical image analysis: a survey. Med. Image Anal. 1 (2), 91–108.

Pizer, S.M., Joshi, S.C., Thomas Fletcher, P., Styner, M., Tracton, G., Chen, J.Z., 2001. Segmentation of single-figure objects by deformable M-reps. MICCAI, 862–871.

Shen, D., Davatzikos, C., 2000. An adaptive-focus deformable model using statistical and geometric information. IEEE Trans. PAMI 22 (8), 906–913.

Shen, D., Herskovits, E.H., Davatzikos, C., 2001a. An adaptive-focus statistical shape model for segmentation and shape modeling of 3D brain structures. IEEE Trans. Med. Imaging 20 (4), 257–270.

Shen, D., Lao, Z., Zeng, J., Herskovits, E.H., Fichtinger, G., Davatzikos C., 2001b. A statistical atlas of prostate cancer for optimal biopsy, Medical Image Computing and Computer-Assisted Intervention (MICCAI), Utrecht, The Netherlands, 14–17 October 2001, pp. 416–424.

Shen, D., Zhan, Y., Davatzikos, C., 2003. Segmentation of prostate boundaries from ultrasound images using statistical shape model. IEEE Trans. Med. Imaging 22 (4), 539–551.

Staib, L.H., Duncan, J.S., 1992. Boundary finding with parametrically deformable models. IEEE Trans. PAMI 14 (11), 1061–1075.

Zeng, J., Bauer, J.J., Sofer, A., Yao, X., Opell, B., Zhang, W., Sesterhenn, I.A., Moul, J.W., Lynch, J., Mun, S.K., 2000. Distribution of prostate cancer for optimized biopsy protocols. MICCAI, Pittsburgh.