Visual Positioning of Previously Defined ROIs on Microscopic Slides

Grigory Begelman, Michael Lifshits, Ehud Rivlin

Abstract— **In microscopy, regions of interest are usually much smaller than the whole slide area. Various microscopy related medical applications, such as telepathology and computer aided diagnosis, are liable to benefit greatly from microscope auto positioning on previously defined regions of interest. In this paper we present a method for image-based auto positioning on a microscope slide. The method is based on localization of a microscopic query image using a previously acquired slide map. It uses geometric hashing, a highly efficient technique drawn from the object recognition field. The algorithm exhibits high tolerance to possible variations in visual appearance due to slide rotations, scaling and illumination changes. Experimental results indicate high reliability of the algorithm.**

Index Terms— **Telepathology, virtual microscopy, geometric hashing, positioning.**

I. Introduction

MICROSCOPIC digital imagery is increasingly seen as a powerful tool in many medical related fields. In diagnostic pathology, for instance, a specimen is mounted on a glass slide for microscopical examination. Under high magnification a typical specimen can provide tens of thousands of distinct images. Because of the limited field of view of the microscope's eyepiece we only see a single image per view.

We refer to a partial image of a slide (e.g., the field of view of the microscope imaging system) as the "region of interest" (ROI). Various medical procedures that involve microscopic tissue examination and utilize an automated microscope system (e.g., telepathology) would obviously be more efficient with the ability to find previously defined ROIs on a slide. Our algorithm introduces the ability to find the location of a particular ROI on a slide in fractions of a second.

Several practical applications of the algorithm are obvious. The routine of pathological slide preparation includes slide pre-examination by a cytotechnician. The cytotechnician locates the ROIs on a slide and marks them with a pen for further examination by an expert pathologist. The marks can be inaccurate, and the procedure is invasive and time consuming. Moreover, it is impossible to mark features under high magnification due to the approach's very limited precision. The system described in this paper allows practitioners to put "virtual marks" on a slide, which are easily and automatically located at any magnification level. An illustration of this scenario is shown in Fig. 1.

Another possible application is the facilitation of computer aided diagnostics (CAD). Several CAD systems dealing with microscopical images and slides are described in [1], [2], [3], [4]). These systems provide diagnosis based on the slide content, which is subsequently approved by an expert pathologist. Usually the resulting diagnosis is determined based on the content of several relatively small key regions of the slide. A pathologist may wish to examine the vicinity of the key regions on the slide, previously marked by the CAD system as suspect. The ability to automatically locate their correct positions on the slide will greatly increase the efficiency of the pathologist's examination routine.

An auto positioning system may also be applied in the area of medical education. An educational slide is examined by an expert and numerous descriptions associated with certain places of interest on the slide are recorded in the database. Later, students reviewing this slide may, by providing a software system with an image of a particular spot on the slide, get a full description of that spot.

The problem of microscope positioning was addressed before. The existing solutions use specially manufactured slides. For example, Field Finder produced by Electron Microscopy Sciences, Hatfield, PA [5] has a precision rectangular-coordinate grid pattern, which is positioned on a special microscope slide. Each square is marked with a letter and a number. To measure the coordinates of a ROI, one should center the detail in the field of the microscope, replace the specimen slide with the field finder, and read the coordinates. CellFinder microscope slides produced by Microlab, Netherlands [6] use a similar approach for microscope positioning. CellFinder slide has a micro-pattern system consisting of squares, lines, dots and letters. It enables marking and later relocation of the ROI position. These solutions suffer from limited precision of positioning (the accuracy is about $1\mu m$). In addition, measurements and further positioning is performed manually, thus prone to mistakes.

The method proposed in this paper is based on localization of a microscopic query image using a previously acquired image map. The search is accomplished using geometric hashing, a well known object recognition technique. First, morphological features of the slide specimen are extracted. Then, a map is constructed from these features and is efficiently stored in a hash look-up table. Finally, accurate localization is achieved by non-sequential searching of the query ROI image on the map, which is accelerated by indexing.

Our algorithm achieves better positioning accuracy. The accuracy is not affected by slide rotation and displacement, which are likely to happen as it is impossible to place the slide exactly in the same manner under another microscope. Moreover, the algorithm overcomes partial degradation of slide

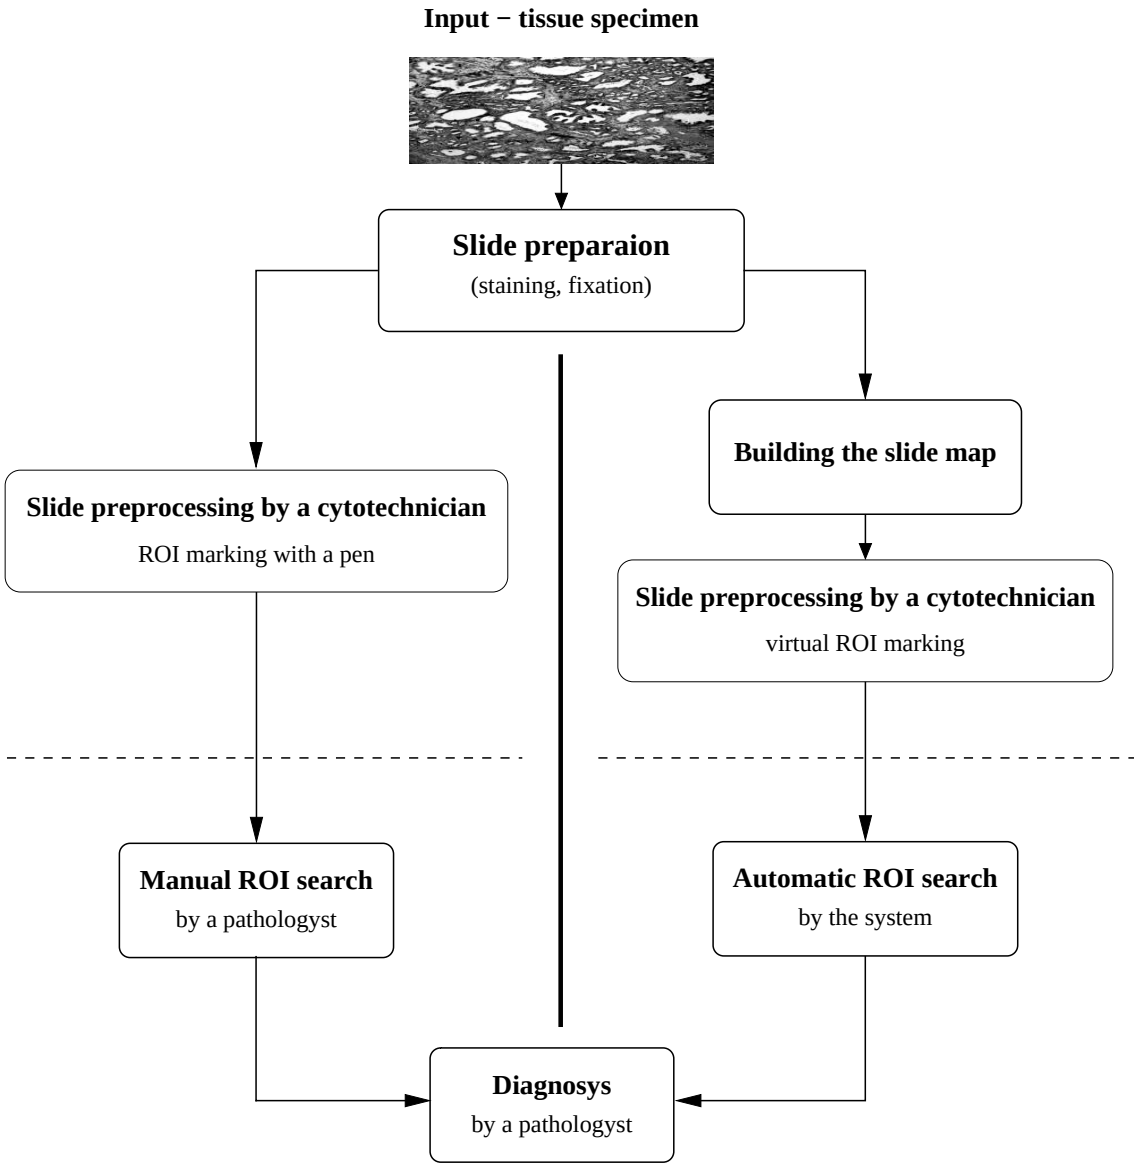

Fig. 1. The overview of a pathologist's examination routine. The tissue specimen is placed on a glass slide for examination under a microscope. The left side of the diagram describes the traditional approach, where the ROIs are marked with a pen. On the right side of the diagram the proposed automatic approach is outlined. The auto-positioning system allows placing "the virtual marks" on the slide during the preprocessing stage (above the dashed line) as a replacement for the traditional approach. The system uses a slide map for this purpose. Virtual marks are easily and automatically located during the pathologist's examination (below the dashed line).

quality. The algorithm is inherently robust to rotations, scale changes and partial pattern obliteration. Likewise, the features we employ are insensitive to a wide range of illumination variations.

The paper is organized as follows: Section II describes our approach to the positioning problem. In Section III the extraction of feature points is explained. Experimental results obtained using a slide positioning system based on the proposed algorithm are presented and analyzed in Section IV. Section V summarizes our work.

II. Proposed Approach for Positioning

Our approach for positioning on a slide is based on localization of small microscopic sub-images using a previously acquired slide map. The algorithm uses advanced image processing techniques and geometric hashing [7], [8], [9], [10], a highly efficient technique drawn from the object recognition field.

A. Localization Problem as an Object Recognition Task

Object recognition is a known problem in the computer vision field. Recognition is achieved by finding the correspondence between a given object and a set of predefined objects. In the model-based object recognition approach, the descriptions of previously known objects are prepared in terms of various properties, such as shape, color, etc. These descriptions are referred to as "models". A given query object will be matched to one of these models.

Localization on a slide is defined in the following manner: given an ROI on a slide, determine its exact coordinates on the slide map. Accordingly, map-based localization can be interpreted as model-based object recognition, as follows. First, the slide map is constructed from partial images captured by a microscope imaging system moving over the surface of the slide. The slide map can be divided into adjacent parts that will be identified during ROI localization. The map parts correspond to a model set in the object recognition framework and the ROI plays the role of query object. Matching the current ROI to one of the parts of the previously constructed slide map during localization is essentially the same as associating a query object to a known model in object recognition. The example of an ROI and a corresponding part on the slide map is shown in Fig. 2.

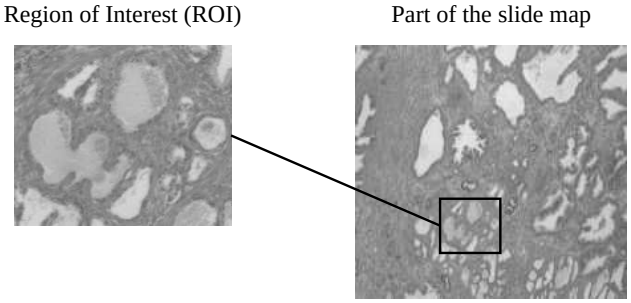

Fig. 2. Example of a real region of interest (ROI) on a slide map.

B. The Localization Algorithm

In order to handle the enormous amount of geometric shapes contained in biological images, we address the positioning problem using an efficient technique from the object recognition field, called geometric hashing. ROI localization is performed by applying object recognition in the following way. We assume a set of predefined geometric models $M_1, ..., M_n$, defining a slide map, and a query ROI image Q, formed from one of the models. The task is to find the model M_i, corresponding to the given query ROI Q. It is assumed that the models are defined by a set of geometric features (e.g., object skeleton junction points) and that the same features can be extracted from the query ROI image. A model can undergo similarity transformations to form the ROI: it can be rotated, translated and uniformly scaled. One way to make feature points invariant under this class of transformations is to represent them in the coordinate frame formed by the points themselves. For example, we may arbitrarily choose an ordered pair of model points to form a basis and describe the rest of the features in this coordinate frame. As there are multiple ways to choose a basis, we are faced with a combinatorial problem of finding the right one to match a model to the ROI.

The algorithm copes with this problem by shifting the computational burden to the off-line learning stage. Instead of going over all feasible ROI/model bases couples and trying to match them, all possible model representations are prepared in advance and stored in a hash table for efficient access. Thus, a query ROI projected onto an arbitrarily chosen basis has a matching model representation already stored in the hash table.

Let $\{\mathbf{q}_1, \ldots, \mathbf{q}_k\}$ be the feature points of the ROI Q and $M_i' = \{\mathbf{m}_1, \ldots, \mathbf{m}_k\} \subseteq M_i$ be the corresponding feature points of the matching model M_i stored in the hash table. Let us denote the transformation model M_i undergoes to form the ROI by T; then $Q = \mathrm{T}(M_i')$, or $\mathbf{q}_j = \mathrm{T}\mathbf{m}_j$ for $1 \le j \le k$. Consider an ordered points pair $(\mathbf{m}_1, \mathbf{m}_2)$ of M_i'. A vector $(\mathbf{m}_1 - \mathbf{m}_2)$ with another vector rotated by 90° form the basis of a 2-D coordinate frame. The coordinates (α, β) of any other point $m_j \in M_i'$ for $3 \le j \le k$, in this frame agree with

$$\mathbf{m}_j = \frac{\mathbf{m}_1 + \mathbf{m}_2}{2} + \alpha(\mathbf{m}_2 - \mathbf{m}_1) + \beta(\mathbf{m} - \mathbf{m}_1),$$

where we denote an end point of the rotated vector $\mathrm{Rot}_{90}(\mathbf{m}_1 - \mathbf{m}_2)$ by $\mathbf{m}$. These (α, β) coordinates will remain unchanged when any linear transformation is applied to the model points. Application of a linear transformation T on the model M_i transforms the point $\mathbf{m}_j$ to:

$$\mathrm{T}\mathbf{m}_j = \frac{\mathrm{T}\mathbf{m}_1 + \mathrm{T}\mathbf{m}_2}{2} + \alpha(\mathrm{T}\mathbf{m}_2 - \mathrm{T}\mathbf{m}_1) + \beta(\mathrm{T}\mathbf{m} - \mathrm{T}\mathbf{m}_2),$$

so the point $\mathrm{T}m_j$ has the same coordinates (α, β) in the frame formed by the ordered basis pair $(\mathrm{T}\mathbf{m}_1, \mathrm{T}\mathbf{m}_2)$. Thus, we refer to coordinates (α, β) as invariant coordinates.

Assuming that the model M_i contains N_i feature points, there are $\binom{N_i}{2}$ different bases for that model. To form a transformation-invariant model representation, the invariant coordinates (α, β) are computed using each one of these bases $B_{\mu\nu} = \{\mathbf{m}_\mu, \mathbf{m}_\nu\}$ for every other model point. The corresponding entry $(M_i, B_{\mu\nu})$ is stored in the hash table with index (α, β).

Positioning on a slide involves two stages. In the *off-line preprocessing stage* all of the time-consuming work needed to compute the transformation-invariant model representations is done. First, the slide map is constructed from small adjacent images captured by a microscope imaging system moving over the surface of the slide. The slide map is divided into parts that correspond to a model set. To allow unconstrained selection of ROIs by a pathologist, these parts should overlap each other by a section equal to the size of the ROI. As a result, any selected ROI is entirely contained in at least one model. Then, to find the feature points of each model from the set, the following procedures are applied. Each slide model is segmented and typical objects within the model are recognized. Feature points are extracted from these objects.

Model representation is prepared by taking every feasible basis b (composed of ordered features pair) and computing the invariant coordinates of all the remaining feature points in terms of the basis b. Then, this representation is stored in the database – the computed coordinates are used as an index to the hash table for recording the entry $(M; b)$. This stage is outlined in Fig. 3. The complexity of this stage is $O(n^3)$ per model, where n is the number of feature points contained in the model. However, since this stage is executed off-line, its complexity is of little significance.

The *on-line localization stage* uses the data prepared by the first stage to perform the matching and localization. Given the ROI image, the following steps are performed:

1) Segment the ROI image and detect typical objects;

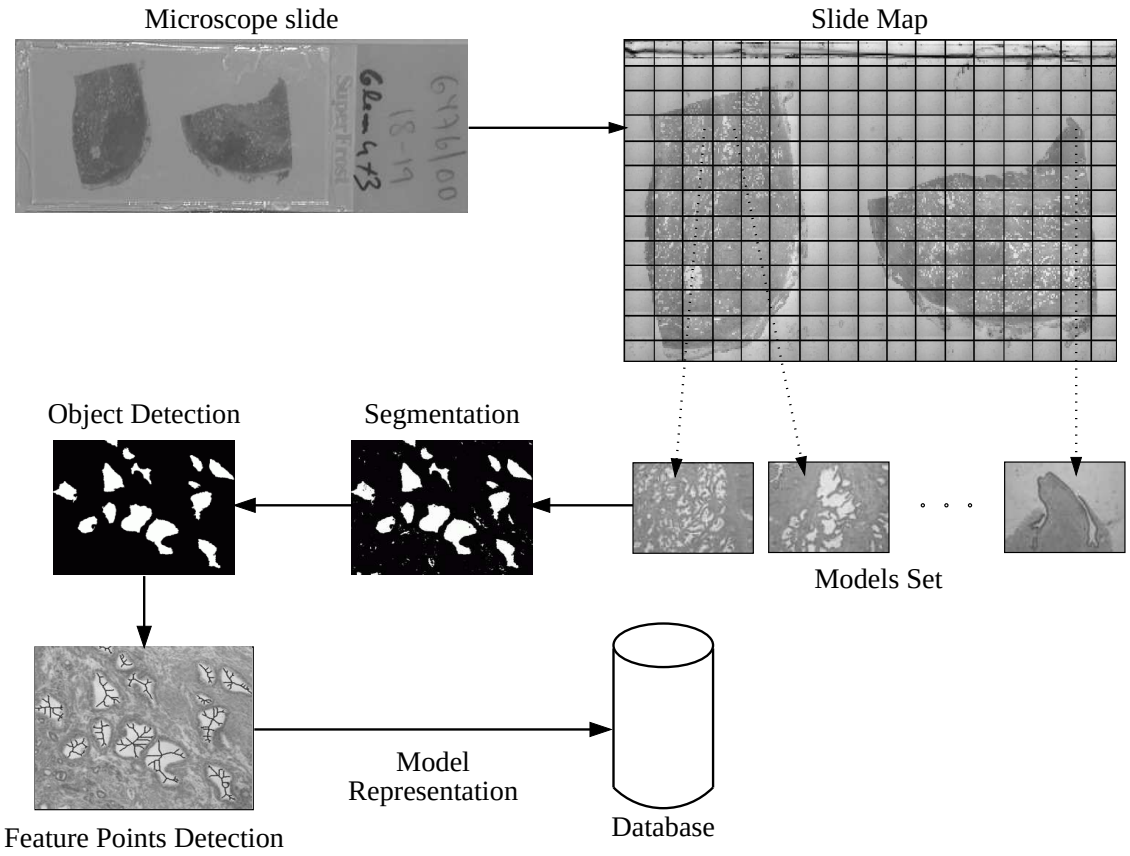

Fig. 3. Outline of the algorithm preprocessing stage. First, the slide map is constructed from small adjacent images captured by a microscope imaging system. Then the slide map is divided into parts corresponding to a model set. After that the feature points of each model are identified. Finally, the model representation is calculated and stored in the database.

2) Extract k feature points from the detected objects;
3) Choose a feasible pair of feature points as a basis b;
4) Compute the invariant coordinates of all the remaining feature points in terms of this basis b;
5) Use each computed invariant coordinate to index into the hash table and vote for all (M_i, b_j)'s retrieved from this bin;
6) Build a histogram for all (M_i, b_j)'s according to the number of received votes;
7) Establish a hypothesis of correspondence between the ROI and an instance of model M_i if (M_i, b_j), for some j peaks in the histogram with a sufficient number of votes;
8) Verify all the hypotheses established in Step 7 and repeat from Step 3 if all of them fail verification.

This stage is outlined in Fig. 4. The complexity of this stage is $O(n) + O(t)$ per probe, where n is the number of points contained in the ROI and t is the complexity of verifying one model. Note that this low on-line complexity determines the actual time for localization. It is independent of the number of models stored in the system, thereby allowing fast localization even on very large scale maps. The technique successfully deals with various possible visual transformations such as 2-D rotations, translations and uniform scale. An example of the localization process is shown in Fig. 6.

C. Verification

The localization algorithm is completed by verification. Given a set of candidate models that have accumulated the highest number of votes, one has to determine which best matches the query ROI. To form the ROI, the models have undergone a similarity transformation, which is a combination of translation, rotation and isotropic scaling. Thus, fitting a model to an ROI should be done by a similarity transformation estimation. The ROI is characterized in terms of a feature points set $\{\mathbf{x}'_i\}$ in $\mathbb{P}^2$, and each of the candidate matching models is likewise described by its feature points $\{\mathbf{x}_i\}$, where $\mathbb{P}^2$ is a projective space. It is essential to find all $\mathbf{x}_i \leftrightarrow \mathbf{x}'_i$ point correspondences to compute a similarity transformation $\mathrm{H_S}$, which transforms a model into the ROI: $\mathrm{H_S}\mathbf{x}_i = \mathbf{x}'_i$ for each i. Two correspondences are enough to fully constrain $\mathrm{H_S}$, as the total number of degrees of freedom for similarity is four (one for the rotation, two for the translation and one more for scaling) and every correspondence gives rise to two independent equations in the entries of $\mathrm{H_S}$. However, since the locations of points in the query ROI are not exact (due to noise), all of the correspondences should be used to determine the best transformation, given the data. Accordingly, $\mathrm{H_S}$ is calculated by finding the least-squares solution of the over-determined linear system.

An important issue is how to efficiently find all of the $\mathbf{x}_i \leftrightarrow \mathbf{x}'_i$ point correspondences. The voting stage of the

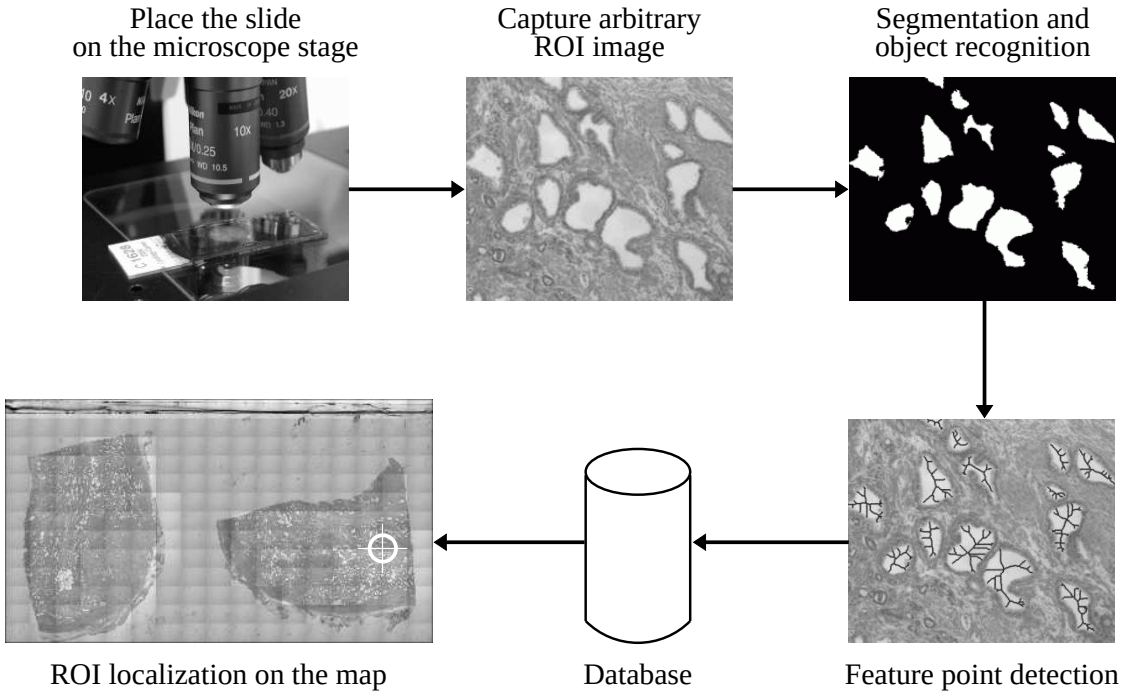

Fig. 4. Outline of the algorithm localization stage. The ROI image is segmented and the typical objects are identified. Then the feature points are extracted from the objects and ROI localization followed by position verification is performed.

algorithm provides one corresponding basis (two point-to-point correspondences) between the candidate model and the ROI. This allows us to approximate the desired transformation $\mathrm{H_S}$ by $\widehat{\mathrm{H_S}}$ and then, after applying $\widehat{\mathrm{H_S}}$ on the candidate model, every model point $\widehat{\mathrm{H_S}}\mathbf{x}_i$ will correspond to the closest ROI feature $\mathbf{x}'_i$. Formally,

$$\mathbf{x}'_i = \arg\min_k d(\mathbf{x}'_k, \widehat{\mathrm{H_S}}\mathbf{x}_i),$$

where subindex k indicates any ROI feature and $d(\mathbf{x}, \mathbf{y})$ is the Euclidean distance between two points $\mathbf{x}$ and $\mathbf{y}$.

Thus, to compute all of the point correspondences we only need to check the distance of each point $\mathbf{x}'_i$ to every transformed model point $\widehat{\mathrm{H}}\mathbf{x}_i$. If the model contains m points and the ROI contains n points, these inter-set distances are computed in $O(m\,n)$ time. This computation can be accelerated by employing a *Voronoi tessellation* [11] for segmentation of the ROI image. Voronoi tessellation partitions of a plane with n points into n convex polygons such that each polygon contains exactly one point and every point in a given polygon is closer to its central point than to any other. We start the verification by constructing the Voronoi tessellation from the points in the query ROI, which is done in $O(n\log(n))$ time [11] (see Fig. 5). This allows us to find the corresponding point of $\mathbf{x}_i$ in $O(log(n))$ time by checking which polygon within the Voronoi tessellation contains the transformed point $\widehat{\mathrm{H}}\mathbf{x}_i$ and choosing its center point. This task is performed using a search data structure based on the Delaunay triangulation, which is dual to the Voronoi tessellation (see [11] for details). It follows that the time needed for point correspondences' calculation is reduced from $O(m\,n)$ to $O(m\,log(n))$.

In practice, the situation is complicated by the fact that some ROI feature $\{\mathbf{x}'_i\}$ might be mistakenly reported and will not match any model point. The unmatched points, outliers, can severely disturb the estimated transformation, and consequently, should be identified. In order to make the verification robust to outliers, one has to obtain a big enough set of inliers from the presented correspondences so that the transformation can be re-estimated in an optimal manner. This is done by the RANSAC algorithm [12].

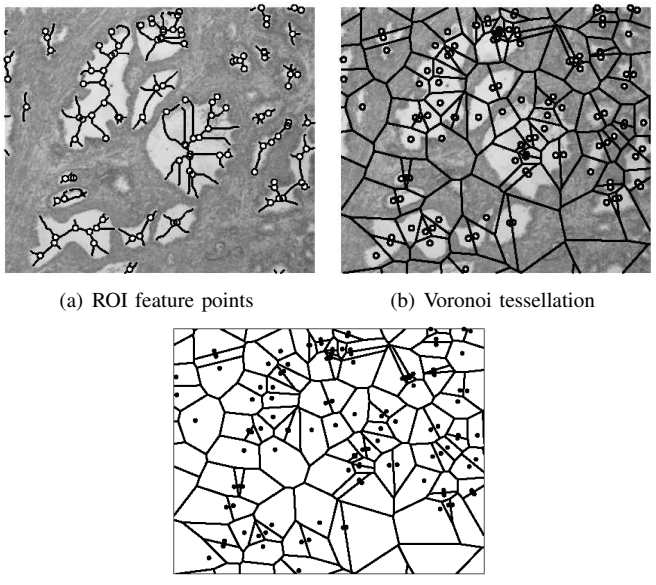

(a) ROI feature points (b) Voronoi tessellation

(c) Voronoi diagram

Fig. 5. The process of constructing the Voronoi tessellation of the ROI for verification acceleration. The junction points of glands object skeletons are marked by white dots in (a).

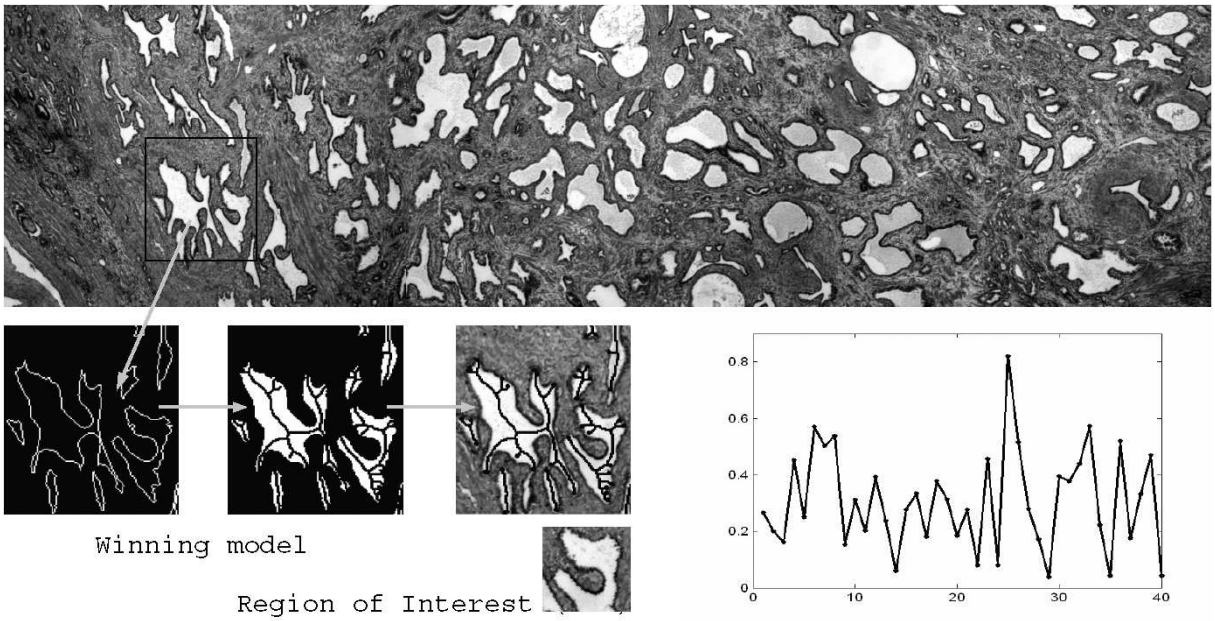

Fig. 6. Outline of the localization process. A part of the slide map that is constructed from 40 model images is shown on the top. A histogram of the voting results for the query ROI (shown on the bottom in the middle) are plotted on the right. The enlarged image of the winning model (25) and the process of its feature points formation is presented on the left. First, glands are found with the robust segmentation algorithm and skeletons are constructed from their shape. Then skeleton junction points are used as the feature points for localization.

III. Feature Points Extraction

The microscopic biological images of prostate and lung tissues we experimented with contain rich color, texture and shape information. Due to possible variations in the illumination conditions and orientation of the slides, we are limited in selection of features that are invariant to these changes. For the current zoom level we use shape information about the glands for positioning.

In order to extract the shape of the glands from the microscopic images of prostate tissue we perform segmentation. This is a challenging problem for the data set we use due to nonuniform illumination, sensor noise and various artifacts. Moreover, there is a color variation between slides as a result of color fading in the stained tissues.

The image segmentation is performed by means of unsupervised clustering of image pixels. The only *a priori* information used for clustering is the existence of a certain number of clusters to be learned (three main clusters in our case: nuclei, glands, and stroma area). The choice of these clusters is biologically driven, and we expect Gaussian distribution of the pixels' color features for each cluster. In the following section we present the Gaussian Mixture Model (GMM) and the related Expectation Maximization (EM) algorithm for unsupervised estimation of GMM parameters.

A. Data Representation with GMM

Gaussian mixture models are widely used as probabilistic models because the number of model parameters is adjustable depending on the complexity of the data distribution. Gaussian mixture models are a semi-parametric alternative to non-parametric histograms [13] (which can also be used as probability densities) and provide greater flexibility and precision in modeling the underlying statistics of sample data. Given a data set, we estimate the model parameters by the EM algorithm [14], [15], and then compute the conditional probabilities for each color pixel [16].

The GMM, a linear combination of Gaussian distributions, can be written in the form:

$$P(x|\Theta) = \sum_{i=1}^{k} \alpha_i n_i(x; \mu_i, \Sigma_i),$$

where $n_i(x, \mu_i, \Sigma_i)$ is a multivariate Gaussian distribution with mean vector μ_i, and covariance matrix Σ_i; k denotes the number of mixed components, α_i denotes the weight of each component, and $\Theta = (\alpha_i, \mu_i, \Sigma_i)_{i=1}^{k}$ represents the collection of model parameters. The EM algorithm is generally used to determine the model parameters Θ from X.

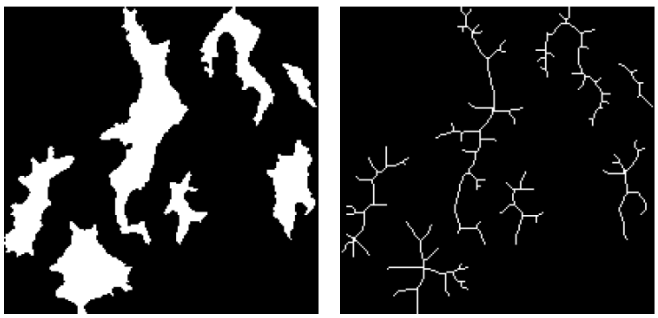

(a) Segmented glands (b) Calculated skeletons

Fig. 7. Examples of skeletons calculated for glands.

Fig. 8. Screenshot of the localization application. The virtual marks corresponding to the previously defined ROIs are shown in the slide overview section (each ROI is marked by two concentric circles). The list of the virtual marks along with their description is on the right side. Current microscope field of view and the best matching candidate on the slide map are displayed on the bottom of the application screen. The coordinates of each candidate and the received matching score are presented in the bottom right corner. A model corresponding to the matching candidate is indicated by a black rectangle in the slide overview section.

B. GMM-based Clustering

Having estimated the model's parameters, we assign each pixel to one cluster according to the estimate of the point's membership function based on its color features. From Bayes rule for each pixel, we have

$$P(cluster\#(x_j) = i) = \frac{\alpha_i n_i(x_j|\mu_i, \Sigma_i)}{\sum_{i=1}^{k} \alpha_i n_i(x_j|\mu_i, \Sigma_i)}.$$

Finally, each pixel is associated with the most probable cluster

$$cluster\#(x_j) = \arg\max_i P(cluster\#(x_j) = i).$$

At this point we apply a morphological cleaning algorithm to the segmented image. As a result we get three types of areas: stroma, glands and nuclei (an example of segmentation is shown in Figs. 3 and 4).

We utilize prior information about object colors in order to identify the components of the Gaussian mixture. This information is based on the hematoxylin and eosin slide staining process: the stroma areas are colored pink, the nuclei are colored blue, and the glands remain unstained (transparent).

C. Shape Information Extraction

The next step is to extract shape information from the detected glands. There are many approaches to extracting shape representation (see [17] for a comprehensive survey). Important and essential criteria for shape representation are consistency and conciseness of the objects' shape description. Topology skeletons constitute an intuitively appealing homotopic and thin version of the shape [18]. Therefore, we use topology skeletons (the result of the medial axis transform) as the glands' shape descriptors.

Medial axis transform is a process of reducing foreground regions in a binary image to skeletal remnants that largely preserve the extent and connectivity of the original regions while throwing away most of the original foreground pixels. Given an object boundary $\delta\Omega$, the skeletons are defined as the set of centers of maximal balls contained in $\delta\Omega$, or the locus of points at equal distance from at least two boundary points:

$$S(\delta\Omega) = \{p \in \Omega | \exists q, r \in \delta\Omega, q \neq r : dist(p, q) = dist(p, r)\}$$

The distance metric is usually the Euclidean one: $dist(p, q) = \|p - q\|_2$. The distance transform (DT) of the $\delta\Omega$ may be used in producing skeletons. It is defined as

$$DT(p) = min_{q \in \delta\Omega} dist(p, q),$$

for all points $p \in \Omega$. The equation above assigns to all the points p the distance to the closest boundary point q. All the pixels lying along the singularities (i.e., curvature discontinuities) of the distance transform function constitute the topology skeletons. Examples of the calculated skeletons for prostate glands are shown in Fig. 7.

The medial axis is an attractive shape feature for our method, as it describes the shape of the typical objects concisely. However, it is sensitive to boundary noise. In order to overcome the sensitivity problem, we regularize the shape boundary and prune the resulting skeletons. We apply

curvature flow smoothing [18] for regularization of the shape boundary. Then we use an area-based significance measure [19] to prune the insignificant skeleton parts.

We select the junction points of the computed skeletons as feature points for our experiments (see Fig. 5 for examples of segmented glands, skeletons of the glands and their extracted feature points).

IV. Experiments

This section presents a number of experiments demonstrating both quantitative and qualitative aspects of the proposed positioning algorithm. We performed tests on both cytological and histological microscopic image sets.

We created an application implementing our algorithm. Its main window is shown in Fig. 8. The application allows us to define virtual marks corresponding to certain ROIs on the slide, which are represented by the partial slide images. As a result, one can search for the location of any of these ROIs on the slide. In addition, the application allows to estimate the exact coordinates of the image seen in the eyepiece of a microscope (despite the slide's rotation and movement) by searching for this image on the slide map.

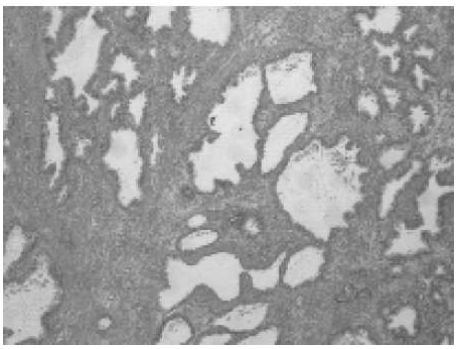

(a) A sample of prostate tissue

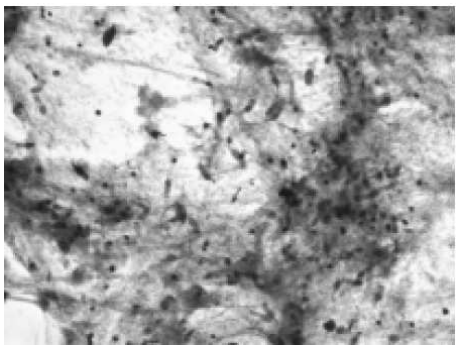

(b) A sample of lung tissue

Fig. 9. Examples of histological (a) and cytological (b) tissue images we used in our experiments.

A. Test Data Sets

We experimented with microscopic histological and cytological slides. To create data sets we constructed two model collections (slide maps), one collection per slide. Both maps covered an area of $10\mathrm{mm} \times 4\mathrm{mm}$ on the slides' surfaces. Every model had about 200 feature points. The model processing time during the preprocessing stage was about 2 minutes on Pentium 4, 2.4 GHz. Examples of the models are shown in Fig. 9. The whole map of the histological slide appears in Fig. 3. To create the slide map we used adjacent microscopic images obtained on a Nikon Eclipse E600 microscope at zoom $\times 20$. The microscope is equipped with a Point Grey CCD camera producing color images with a resolution of 1024×768 pixels. The microscope is driven by a stepping motor control MCL-2 (Lang, Hüttenberg, Germany), operated by a software program (Wincommander; Märzhäuser-Wetzlar, Wetzlar, Germany [20]) via an RS-232 interface.

B. Overall System Test

In this section we demonstrate the capabilities of the proposed localization algorithm and provide an evaluation of its performance. In order to index the hash table during localization and vote for the correct model, an invariant description of the query ROI is calculated. This description is based on a pair of features that form a basis, as illustrated in Section II-B. In practice, it is possible that one of the points used to form a basis will be reported by mistake and, as a result, not match any model point. Moreover, possible inaccuracies in the basis point locations (induced by noise) may have a similar effect. In order to avoid this, multiple attempts should be made using different bases (e.g., different descriptions) to ensure, with sufficiently high probability, that at least one of them is free of extremely noisy points or outliers.

To obtain a statistically meaningful measure of the algorithm's performance we tested it on a total of 10^4 different ROI localization tasks. We varied the number of different ROI feature bases used in voting, and evaluated the localization performance with different levels of added Gaussian noise (see Fig. 10). Each time we selected a random ROI and then, if the correct location on the slide map was reported by the algorithm (ground truth was available due to the nature of data set formation), the result was considered to be true positive (TP). The result was correspondingly considered a miss if an incorrect or no location was detected. The summary of the obtained results is presented in Fig. 10. The hit rate $\left(\frac{\sharp \mathrm{TP}}{\sharp \mathrm{Tests}}\right)$ was above 90% for a moderately noisy level ($\sigma < 1$) when fifty bases are used.

We formulated the accuracy of the localization result as follows. Assuming the ROI features are measured with a Gaussian error of standard deviation σ, it can be shown that [21] the root mean square (RMS) distance of the estimated point location from its true value is $\sigma(2/n)^{1/2}$, where n is the number of correspondences used. Therefore, for 50 sample points the algorithm estimation error is 0.2 pixels, taking $\sigma = 1$. As a result, if the ROI image of size 200×200 pixels captures an area of 40×40 micron, our algorithm provides a positioning accuracy of $0.04\ \mu\mathrm{m}$.

C. Algorithm Robustness to Image Variations

The aim of this test is to provide an experimental analysis of the algorithm behavior when the slide images are distorted. We developed a procedure that systematically distorts the original slide images to simulate possible visual changes. We tested

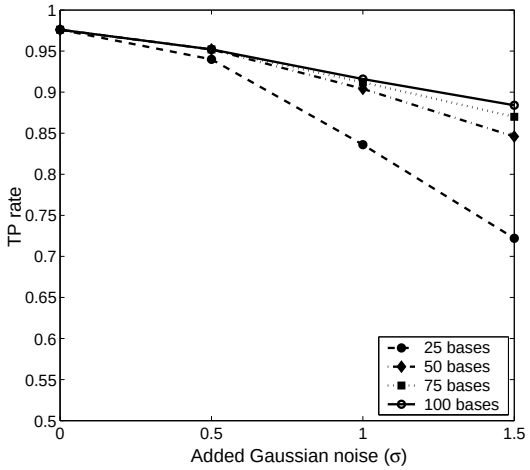

Fig. 10. System behavior with a different number of bases used in voting and different levels of added Gaussian noise.

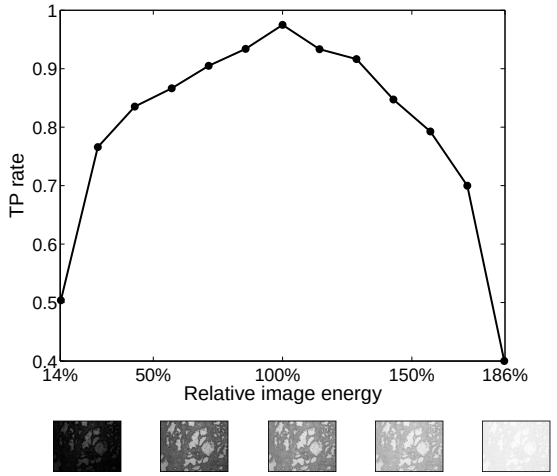

Fig. 11. System behavior under simulated varying illumination. The images under the graph are the examples corresponding to various illumination values.

our algorithm on these systematically distorted images. We concentrated on the image distortions that are most likely to happen in the real world: brightness changes (due to varying illumination) and image rotations (due to inaccurate slide placement under a microscope). We transformed the original image brightness according to the following colormap transformation. The values of the colormap are raised to the power of γ, where γ is

$$\gamma = \begin{cases} 1-\beta, & \beta > 0 \\ \frac{1}{1+\beta}, & \beta \leq 0 \end{cases} \qquad \beta \in [-0.9 \ldots 0.9].$$

This ad-hoc simulation does not model all relevant illumination changes but produces images coherent with a number of real-world examples that we analyzed. In addition, we rotated the images up to 12°; this range corresponds to probable inaccuracy in manual slide placement. The average results of positioning performed on randomly chosen distorted ROI images are shown in Figs. 11 and 12 along with some examples of these images. The localization results were interpreted as in the overall system test (see Subsection IV-B). Note that in both cases the hit rate (the rate of correctly matching the ROI) does not drop below 0.8. In the former case this hit rate was achieved even when the relative energy of the ROI image is less than 30% or greater than 150% (see Fig. 11). In these cases the images' contrast is very low and most image details are indistinguishable, even by the human eye.

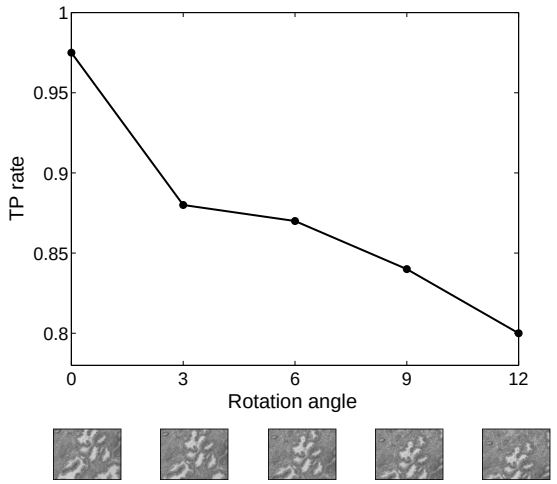

Fig. 12. System behavior under simulated slide rotation. The images under the graph are the examples corresponding to rotations by 0-12 degrees, respectively.

V. Conclusions

We presented a method for image-based positioning on microscopic slides, based on geometric hashing. Precise positioning is achieved using invariant features extracted from the specimen representative objects. The method was verified on cytological and histological microscopic slides. It is inherently insensitive to slide rotation and displacement. Extensive experimental analysis demonstrated high tolerance to a broad range of illumination changes and robustness to noise.

This approach can greatly increase the efficiency of the pathologist's examination routine and facilitate computer-aided diagnostics systems. Likewise, it can be applied to a wide range of medical tasks involving examination of biological tissues under a microscope.

References

[1] K. Jafari-Khouzani and H. Soltanian-Zadeh, "Multiwavelet grading of pathological images of prostate," *IEEE Transactions on Biomedical Engineering*, vol. 50, no. 6, pp. 697–704, June 2003.

[2] D. Comaniciu, B. Georgescu, P. Meer, W. Chen, and D. Foran, "Decision support system for multiuser remote microscopy in telepathology," in *Proceedings of 12th IEEE Symposium on Computer-Based Medical Systems*, 1999, pp. 150–155.

[3] D. Foran, D. Comaniciu, P. Meer, and L. Goodell, "Computer-assisted discrimination among malignant lymphomas and leukemia using immunophenotyping, intelligent image repositories, and telemicroscopy," *IEEE Transactions on Information Technology in Biomedicine*, vol. 4, no. 4, pp. 265–273, 2000.

[4] R. Allan and W. Kinsner, "A study of microscopic images of human breast disease using competitive neural networks," in *Canadian Conference on Electrical and Computer Engineering*, vol. 1, 2001, pp. 13–16.

[5] "Electron microscopy sciences," http://www.emsdiasum.com.

[6] "Microlab," http://www.antenna.nl/microlab/index-uk.html.

[7] J. T. Schwartz and M. Sharir, "Identification of partially obscured objects in two and three dimensions by matching noisy characteristic," *Int. J. Rob. Res.*, vol. 6, no. 2, pp. 29–44, 1987.

[8] Y. Lamdan and H. Wolfson, "Geometric hashing: A general and efficient model-based recognition scheme," *ICCV*, vol. 88, pp. 238–249, 1988.

[9] Y. Lamdan, J. Schwartz, and H. Wolfson, "Affine invariant model-based object recognition," *IEEE Transactions on Robotics and Automation*, vol. 6, no. 5, pp. 578 – 589, 1990.

[10] H. J. Wolfson and I. Rigoutsos, "Geometric hashing: An overview," *IEEE Computational Science and Engineering*, vol. 13, pp. 10–21, 1997.

[11] M. de Berg, O. Schwarzkopf, M. van Kreveld, and M. Overmars, *Computational geometry : algorithms and applications*, 2nd ed. Springer-Verlag, 2000.

[12] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Commun. ACM*, vol. 24, no. 6, pp. 381–395, 1981.

[13] C. Bishop, *Neural Networks for Pattern Recognition*. Oxford University Press, 1995.

[14] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. Royal Statistical Soc.*, vol. 39, pp. 1–38, 1977.

[15] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*, 2nd ed. John Wiley & Sons, Inc., 2001.

[16] Y. Raja, S. McKenna, and S. Gong, "Segmentation and tracking using colour mixture models," in *Asian Conference on Computer Vision*, 1998.

[17] S. Loncaric, "A survey of shape analisys techniques," *Pattern Recognition*, vol. 31, no. 8, pp. 983–1001, 1998.

[18] F. Mokhtarian and A. K. Mackworth, "A theory of multiscale, curvature-based shape representation for planar curves," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, pp. 789–805, 1992.

[19] D. Attali, G. Sanniti di Baja, and E. Thiel, "Pruning discrete and semicontinuous skeletons," in *Proc. of 8th ICIAP*, 1995.

[20] "Märzhäuser - micropositioning systems," http://www.marzhauser.com.

[21] R. Hartley and A. Zisserman, *Multiple view geometry in computer vision*. Cambridge University Press, 2000.