

Automatic coregistration of volumetric images based on implanted fiducial markers

Martin Koch^{a)}

Oncology Care Systems Group, Siemens Medical Solutions USA, Inc., 4040 Nelson Avenue, Concord, California 94520 and Chair of Pattern Recognition, Friedrich-Alexander-University Erlangen-Nuremberg, Martensstr. 3, 91058 Erlangen, Germany

Jonathan S. Maltz

Oncology Care Systems Group, Siemens Medical Solutions USA, Inc., 4040 Nelson Avenue, Concord, California 94520

Serge J. Belongie

Department of Computer Science and Engineering, University of California, San Diego, 9500 Gilman Drive, Mail Code 0404, La Jolla, California 92093

Bijumon Gangadharan, Supratik Bose, Himanshu Shukla, and Ali R. Bani-Hashemi

Oncology Care Systems Group, Siemens Medical Solutions USA, Inc., 4040 Nelson Avenue, Concord, California 94520

(Received 29 February 2008; revised 23 July 2008; accepted for publication 2 August 2008; published 16 September 2008)

The accurate delivery of external beam radiation therapy is often facilitated through the implantation of radio-opaque fiducial markers (gold seeds). Before the delivery of each treatment fraction, seed positions can be determined via low dose volumetric imaging. By registering these seed locations with the corresponding locations in the previously acquired treatment planning computed tomographic (CT) scan, it is possible to adjust the patient position so that seed displacement is accommodated. The authors present an unsupervised automatic algorithm that identifies seeds in both planning and pretreatment images and subsequently determines a rigid geometric transformation between the two sets. The algorithm is applied to the imaging series of ten prostate cancer patients. Each test series is comprised of a single multislice planning CT and multiple megavoltage conebeam (MVCB) images. Each MVCB dataset is obtained immediately prior to a subsequent treatment session. Seed locations were determined to within 1 mm with an accuracy of $97 \pm 6.1\%$ for datasets obtained by application of a mean imaging dose of 3.5 cGy per study. False positives occurred in three separate instances, but only when datasets were obtained at imaging doses too low to enable fiducial resolution by a human operator, or when the prostate gland had undergone large displacement or significant deformation. The registration procedure requires under nine seconds of computation time on a typical contemporary computer workstation. © 2008 American Association of Physicists in Medicine. [DOI: [10.1118/1.2975153](https://doi.org/10.1118/1.2975153)]

Key words: fiducial markers, point-based registration, image-guided radiation therapy

I. INTRODUCTION

Volumetric imaging immediately prior to radiation therapy treatment delivery facilitates accurate dose administration via improved patient positioning and/or treatment plan adaptation. By coregistering these volumetric datasets with the original planning computed tomography (CT) volume, it is possible to estimate the geometric transformation between anatomic elements.

Contemporary image-guided radiation therapy (IGRT) systems typically employ registration methods based on large scale gray level image matching. In scenarios such as pretreatment imaging of the prostate gland, intrinsically poor contrast between soft tissue structures can lead to inaccurate registration. Unsupervised matching algorithms are often disproportionately influenced by bony landmarks, the motion of which then dominates the determination of the resulting transformation parameters. This can lead to large inaccura-

cies, since target structures such as prostate can move by over a centimeter with respect to bony landmarks.¹

Radio-opaque fiducial markers are often implanted into mobile structures such as the prostate gland and lung in order to better identify the displacement and deformation of such organs.²⁻⁵ When such fiducials are available, it makes sense to employ these as landmarks in the coregistration procedure. This is because the presence of fiducials decreases reliance on the skill and anatomical knowledge of the operator, as well as the soft tissue contrast resolution of the imaging modality.

Since fiducial-based registration lends itself to unsupervised automation, algorithms were initially developed for the determination of seed displacement using a two-dimensional (2D) pretreatment projection x-ray image. This would be compared to a reference digitally reconstructed radiograph derived from the planning CT.⁶ The algorithm presented here

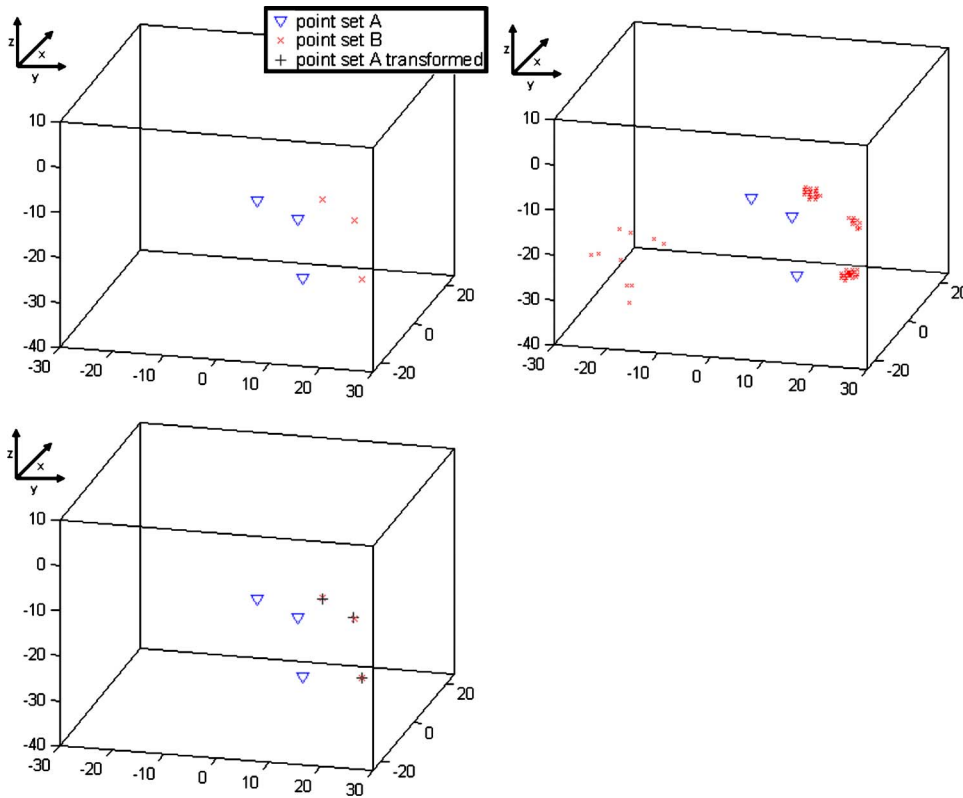


FIG. 1. Illustrative example of the application of a point matching algorithm based on a rigid geometric transformation. Upper left: The locations of reference points from dataset A, and the true locations of these points in dataset B are plotted in a single coordinate system. Upper right: After the application of feature extraction techniques to dataset B, a large number of candidate point matches are possible. The point matching algorithm must find correspondences that are consistent with a constrained rigid geometrical transformation of the reference points from dataset A. Lower left: The point matching algorithm correctly identifies the point correspondences between datasets A and B. In the present application, the transformation associated with these points may then be used to adjust the patient position or the treatment plan accordingly.

achieves the same objective using three-dimensional (3D) datasets. The major advantage of a 3D-to-3D registration is that information from all angular projections is used to determine the seed positions. Consequently, greater accuracy is possible, especially in situations where the presence of more than three markers leads to identification ambiguity.

A second potential advantage relates to the distribution of imaging radiation dose to the patient. As compared to 2D methods based on paired orthogonal projections, tomographic acquisition distributes total imaging dose in a different manner throughout the patient. For example, an anterior arc may not only redistribute dose away from the rectal wall, but can also be routinely considered as part of the treatment dose distribution for overall optimization.⁷ We apply our algorithm to images obtained over a wide dose range in order to illustrate that the method is effective in registering datasets acquired at a doses as low as 1.0 cGy, which is significantly less than that used for typical megavoltage (MV) orthogonal portal image-based IGRT.

II. METHODS

II.A. Seed coregistration as a point matching problem

The registration of 3D datasets based on implanted fiducial markers is a classical point matching problem. The location of a known point set taken from volumetric dataset A must be found within volumetric dataset B. The latter dataset will generally contain, in addition to the fiducials present in A, an indeterminate number of additional points that confound the identification of the desired point set.

A useful point matching algorithm must be able to accurately detect the corresponding points in dataset B, while in the process rejecting the spurious features present in B. Chui *et al.* present a robust point matching algorithm that estimates the correspondence between point sets by computing a nonrigid transformation between the two sets.⁸ For IGRT applications that do not involve treatment replanning, it is prudent to accept only translations and rotations of the target structures. A full nonrigid registration approach allows for tissue deformation and this cannot be compensated for by simply readjusting patient position. We therefore pursue a rigid body transformation approach to point matching. While such an approach is able to tolerate some degree of nonrigid deformation, such transformations are associated with a higher “cost” in terms of a goodness-of-fit metric. Unacceptably high cost function values indicate that it may not be appropriate to deliver a plan to a patient on a particular day, or indeed, that replanning may be necessary.

The rigid point matching approach we develop here is inspired by the random sample consensus (RANSAC) algorithm,⁹ which is widely used for the robust fitting of models in the presence of many data outliers. An example application of this algorithm to fiducial seed data appears in Fig. 1. In the upper left panel of this figure, we see the reference point set from dataset A plotted alongside the true locations of these points in dataset B. The upper right panel illustrates, in addition to the reference points in dataset A, the collection of potential matching points after feature extraction has been applied to dataset B. While many points are clustered around the true locations of the fiducials, several

TABLE I. Table of mathematical symbols used.

Image R	reference image
Image n	n th pretreatment image
VOI_n	volume-of-interest of n th image
$f[q, r, s]$	intensity value of the voxel at discrete position $[q, r, s]$
I_{\max}	maximum intensity value in the VOI
γ	threshold parameter $0 \leq \gamma \leq 1$
P	number of fiducial markers
\mathbf{x}	world coordinates of a single point
\mathbf{x}_p^R	world coordinates of markers in reference image
k	index for thresholding iteration
i	index for matching iteration
j	index for grouping iteration
S_i^k	set of points at iteration k , i
g_l	set of points within close mutual proximity, called group
L	number of groups
G_{ij}	set of all groups at iteration i , j
C	temporary set of points in grouping process
D	set of points or groups that are fed to the matching algorithm
r	proximity radius parameter
UB	upper bound used in the matching process
a_v	size of the VOI
V_{\max}	maximum volume of feasible features
ε_i	transformation error at iteration i
$\varepsilon_{\text{final}}$	error obtained upon final application of the matching algorithm
ε_+	threshold for early termination
ε_{\max}	maximum allowed transformation error
h	transformation function
λ	parameter vector that determines the geometric transform
$\Delta\lambda$	changes in parameter vector
$\mathbf{b}^l, \mathbf{b}^h$	maximum permitted rotation angles in each direction

remote outliers are also present. Application of the point matching algorithm finds the best rigid transformation between reference set A and all of the potential points by exhaustively searching all possible correspondences between point sets. The final match is illustrated in the lower left panel. This match is not perfect, indicating in this case that some minor elastic deformation of the anatomy has occurred.

II.B. Detailed algorithm description

We now describe the proposed algorithm in the context of the imaging and treatment workflow. While we define all mathematical symbols as these are introduced in the follow-

TABLE II. Table of mathematical operators used.

\cup	union of sets
\subset	subsets of a set
$ A $	cardinality=number of elements in set S
$A \setminus B$	set exclusion (select only those elements of set A that are not present in set B)
$n!$	factorial= $1 \times 2 \times \dots \times (n-1) \times n$
$\binom{n}{k}$	number of possible ways of selecting k elements from a set of n elements without replacement $=n!/[k!(n-k)!]$

ing exposition, for the convenience of the reader we list all symbols and operators used in Tables I and II, respectively.

- (1) At the start of the treatment planning process, the patient is imaged using a 3D modality such as x-ray CT or magnetic resonance imaging. We denote this reference image as image R .
- (2) The locations of the fiducial markers are determined manually and saved for later use. We denote the world coordinates of the P markers as $\mathbf{x}_p^R = [x_p, y_p, z_p]^T$, $p = 1, 2, \dots, P$.
- (3) Before delivery of the n th treatment fraction, the patient is imaged using a 3D modality to yield image n .
- (4) The proposed algorithm is used to automatically detect the positions of the fiducial markers within image n and compute the transformation of these positions with respect to the reference image. The \mathbf{x}_p^R are given as input.

The registration process proceeds by iterating over three nested loops. The variables k , i , and j represent the indices for the thresholding, matching, and grouping loops, respectively. The algorithm proceeds as follows for image n :

- (a) A volume-of-interest (VOI_n) is defined that constrains the unknown fiducial positions to a feasible volume within the image. This is normally predefined for a specific type of study and is not adjusted on a case-by-case or patient-by-patient basis.
- (b) A thresholding operation is performed to identify voxels in VOI_n that contain highly attenuating material. These voxel indices are assigned to the set

$$S_0^k = \{[q, r, s] | f_n[q, r, s] > \gamma \times I_{\max}[q, r, s] \in \text{VOI}_n\}, \quad (1)$$

where $f_n[q, r, s]$ is the intensity value of the voxel of image n at discrete position $[q, r, s]$, I_{\max} is the maximum intensity value in the VOI, and γ is a threshold parameter chosen by the operator.

- (c) Connected component analysis^{10,11} (CCA) is applied to group adjacent points into features. This is performed only for the purpose of removing from the candidate point set the contributions of contiguous features that are larger than seeds. Features whose volume exceeds V_{\max} , the maximum expected apparent volume of a seed, are removed from S_0^k .
- (d) In order to obtain a reduction in dimensionality and a resulting decrease in computational burden, voxels within close mutual proximity are grouped into the L possibly overlapping sets g_l , $l=0, 1, \dots, L$. These sets of voxels are collected so as to become subsets of the possibly larger set

$$G_{ij} = \{g_l | g_l \subset S_i^k\},$$

$$\cup G_{ij} = S_i^k. \quad (2)$$

This step differs from the CCA analysis applied in the previous step in that mutual proximity rather than contiguity is the criterion for grouping. Also, the intention behind this step is dimensionality reduction rather than large-feature exclusion.

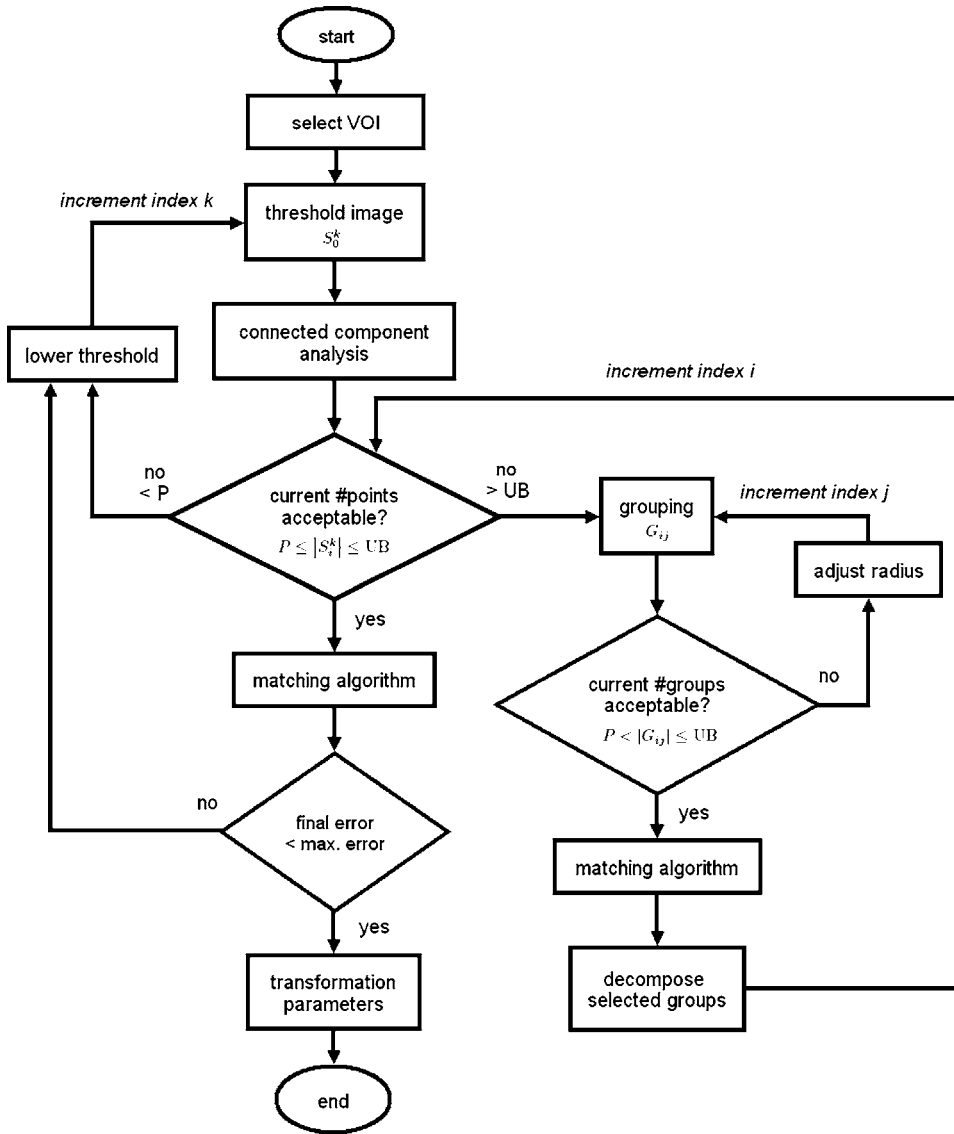


FIG. 2. Algorithm flowchart. Notation: P : number of markers, UB : upper bound used in the matching process.

(e) A matching algorithm is applied that attempts to transform the reference fiducial positions to the positions of the centroids of subsets g_i of the set G_{ij} . The algorithm produces a quality-of-fit metric which is used to select the best match. Let $g_p, p=1, 2, \dots, P$ be the set of groups that scores the best match. The current set of groups is then updated as

$$S_{i+1}^k = \bigcup_{p=1}^P g_p.$$

(f) Steps (d) and (e) are repeated until $|S_i^k| \leq UB$, where UB is an upper bound of numbers of groups to be used in the matching process. This parameter is chosen by the operator and its value determines the nature of the compromise between processing time and detection accuracy.

(g) At this point, the set S_i^k contains the points \mathbf{x}_p and $P \leq |S_i^k| \leq UB$. The transformation parameters are computed to match the $\mathbf{x}_p \in S_i^k$ to the \mathbf{x}_p^R . The error of this transformation is denoted ϵ_i .

(h) If $\epsilon_i > \epsilon_{max}$, we repeat steps (b)–(g) using a lower value of intensity threshold γ . The parameter ϵ_{max} represents the maximum allowed transformation error. This is application-dependent and is selected by the operator.

Figure 2 describes the adaptive thresholding and matching process in more detail, as does the exposition below. Figures 3–7 provide an illustrative example of the application of the thresholding, grouping and matching processes to clinical images.

II.C. Grouping of points

After elimination of the contribution of large features using connected component analysis, the number of points $|S_i^k|$ may still be undesirably large for direct application of the matching algorithm. To speed up the matching process, a




FIG. 3. The upper panel shows transaxial slices from CT (upper) and MCVB CT (lower) images of the same patient. The presence of two fiducial markers within the prostate is evident in this 3 mm-thick CT slice. Only the lower left marker is visible in the 1 mm-thick MCVB CT slice owing to the smaller slice thickness. The enlarged volume-of-interest is shown to the right of this image. The MCVB CT image was obtained using a dose of 5 cGy.

dynamic grouping process is implemented whereby several single points within close mutual proximity are grouped into the groups g_l . This process is defined by the following pseudocode:




FIG. 4. The upper panel shows the MCVB CT slice that appears in Fig. 3 after the initial thresholding operation. The black spots represent voxels selected to be included as potential features. We see that along with the true fiducial marker, some of the bone volume also exceeds the threshold, thus adding spurious elements to the feature set. In the lower panel, which shows a slice where no seeds are present, a small amount of bone constitutes the entire feature set.




FIG. 5. After the application of grouping and matching steps to the features identified in Fig. 4 (and simultaneously to similar features in other slices within the volume that we have not shown), we find that the spurious voxels representing bone volume have been removed from the feature set. Only the true fiducial seed in the slice shown in the upper panel remains in this set, as is clear from the presence of only a single constellation of black voxels in the upper panel.

```

 $C := S_i^k; l := 1;$ 
for each voxel index in  $C$  or until  $|C|=0$ 
  find all voxel indices  $\in S_i^k$  within radius  $r$ 
  assign voxels indices to  $g_l$ 
   $C := C \setminus g_l$ 
   $l := l + 1$ 
end

```

The representative coordinate of a group g_l is computed as the centroid of its points. The maximum distance between two voxels belonging to the same group is smaller than or equal to $2r$. Depending on the actual number of groups $|G_{ij}|$, a regrouping may be performed where the proximity radius parameter r is dynamically adjusted to keep the number of groups within the range $P < |G_{ij}| \leq UB$. When such an appropriate group size is attained, the matching algorithm is applied to select the P most relevant groups. These groups are then decomposed into their constituent points. If more than




FIG. 6. After the procedures illustrated in Figs. 3–5 are complete, the algorithm outputs the detected seed locations as well as the geometrical transformation needed to register the patient to the planning CT. The three detected seeds lie at the centers of the white circles shown in the right column (MVCB CT). The left column shows the slices of the planning CT of the same patient for reference.

UB points exist after decomposition, the grouping process ensues, otherwise the points are fed directly to the matching algorithm.

II.D. Matching algorithm

The optimal match is computed by a pattern matching algorithm similar to RANSAC (random sample consensus).⁹ Let D denote the set of points or groups that are fed to the matching algorithm. For the matching of groups and points, we have $D = G_{ij}$ and $D = S_i^k$, respectively.

At each iteration, P elements are randomly chosen out of the set D of candidate matches. Since the order of the points is relevant, the maximum number of iterations necessary to evaluate all matches is $P! \times \binom{|D|}{P}$. The notation $|D|$ denotes the number of elements in set D (cardinality).

The cost function that is minimized by the matching algorithm is the root mean square of the distance between the transformed source points and the closest points (or group centroids) in the target dataset.

The pseudocode for the matching algorithm is given as follows:

```

for each possible permutation of points or groups
  choose  $P$  points  $\mathbf{x}_p$  out of  $D$ 
  compute transformation between  $\mathbf{x}_p$  and  $\mathbf{x}_p^R$ 
  if the matching error  $\epsilon \leq \epsilon_+$ 
    break
end

```

The transformation is constrained to consist of rotations around three orthogonal axes and a 3D translation. The root mean square (RMS) distance ϵ between the transformed source points and the current target points is used as the

quality metric for the transformation. After evaluating all possible target point combinations, we choose the point set associated with the transformation that produces the smallest error. The iteration over all possible combinations is terminated before an exhaustive search is completed if the error of a transformation is below an operator-defined threshold ϵ_+ .

We denote the error obtained upon final application of the matching algorithm, where individual points are considered, as ϵ_{final} .

II.E. Determination of matching transform parameters

The transformation parameters between prospective point set matches are computed by optimization of the goodness-of-fit function using a nonlinear projected descent method. This approach iteratively minimizes the sum of squared distances between transformed source points $h(\mathbf{x}_p^R, \boldsymbol{\lambda})$ and the target points \mathbf{x}_p to yield the optimal transformation vector as

$$\boldsymbol{\lambda}^{\text{opt}} = \arg \min_{\boldsymbol{\lambda}} \sum_{p=1}^P \|\mathbf{x}_p - h(\mathbf{x}_p^R, \boldsymbol{\lambda})\|^2, \quad (3)$$

where the solution is subject to the bound constraints

$$\mathbf{b}^l \leq \boldsymbol{\lambda} \leq \mathbf{b}^h. \quad (4)$$

The function h applies the rigid transformation, which consists of a 3D rotation and translation. The 3D rotation is comprised of successive rotations around the x , y , and z axes. These axes are defined to intersect at the isocenter of the megavoltage cone beam (MVCB) CT image. The parameter vector $\boldsymbol{\lambda}$ determines the transform and is initialized using the operator-defined starting values $\boldsymbol{\lambda}^0$. The boundaries \mathbf{b}^l and \mathbf{b}^h represent the operator-defined maximum permitted rotation angles in each direction.

To perform the optimization in Eq. (3), we employ a stabilized Newton method. In this algorithm, the Newton descent direction is chosen when the Hessian is positive semidefinite. Otherwise, the steepest descent direction is used. The Armijo line search rule is employed to determine step size (Ref. 12, p. 58).

To implement the bound constraints, at iteration j the updated parameter vector

$$\boldsymbol{\lambda}^j = \boldsymbol{\lambda}^{j-1} + \Delta \boldsymbol{\lambda}^j \quad (5)$$

is projected onto the permissible solution set using the operator¹³

$$[\boldsymbol{\lambda}]_{\pm} = \begin{bmatrix} \max\{\min\{\lambda_1, b_1^u, b_1^l\}\} \\ \vdots \\ \max\{\min\{\lambda_n, b_n^u, b_n^l\}\} \end{bmatrix}. \quad (6)$$

III. EVALUATION

III.A. Datasets

A high quality treatment planning CT dataset, acquired using a Siemens multislice diagnostic CT system, was obtained for each patient prior to the pre-treatment MV image

TABLE III. Resolution of planning CT images.

Pixel resolution in nm	Slice thickness in mm
1.270×1.270	5.000
0.916×0.916	3.000

volumes. Typically, these images are obtained by delivering 200 mA s at 140 kVp to yield a body CT dose index of 2.2 cGy (center) and 3.6 cGy (surface).

The algorithm was tested on 164 target datasets obtained from ten patients. All target datasets were obtained using megavoltage linac beams of Siemens Primus and Oncor radiation therapy linear accelerators. Two different beam configurations were used in the studies. The majority of the images were produced using a conventional 6 MV treatment beam. This beam is generated by the action of ≈ 7 MeV electrons on a tungsten target. The resulting photons are filtered using a stainless steel flattening filter. The remaining images were obtained using a prototype imaging system in which ≈ 4.5 MeV electrons impinged on a carbon (diamond) target. No flattening filter was employed.¹⁴ We refer to this beam using the nominal designation “4.5 MV.”

The 4.5 MV imaging studies were obtained over a wide dose range in order to determine the minimum possible useful dose for imaging. As a result, the quality of some of the images is extremely poor. In some cases it is not possible for a human observer to identify the fiducial markers.

All patient images were acquired at the University of California at San Francisco and the Savannah Oncology Center, Savannah, Georgia with the approval of the relevant

ethics boards. The algorithm was applied to the data retrospectively and did not influence patient treatment.

A treatment planning CT dataset, obtained using a multi-slice diagnostic CT system, was obtained for each patient prior to the before-treatment images. These datasets had different characteristics with respect to pixel resolution and slice thickness. The values employed are listed in Table III. MVCB images were all acquired with pixel resolution of $1070 \text{ mm} \times 1.070 \text{ mm}$ and slice thickness of 1.000 mm.

The applied dose for MVCB CT images is usually stated in monitor units (MU), where 1 MU produces a dose of 1 cGy at a depth of d_{max} in water (at the isocenter) for a $10 \text{ cm} \times 10 \text{ cm}$ radiation field, with a source-to-isocenter distance of 1 m. The quantity d_{max} is the depth along the dose-versus-depth curve at which the maximum dose is observed. Since we employ two different beams, we quote doses in cGy rather than MU and employ the relevant MU-to-cGy conversion factor for the particular beam.

All images were corrected for cupping artifact (due to beam hardening, lateral projection truncation, scatter, and nonuniform beam profile) using the method described in Ref. 15.

III.B. Specific implementation

In this section, we describe the selection of algorithm parameter values and provide other details relating to the algorithm evaluation procedure.

The single set of operator-chosen parameter values specified below was used to process all of the datasets.

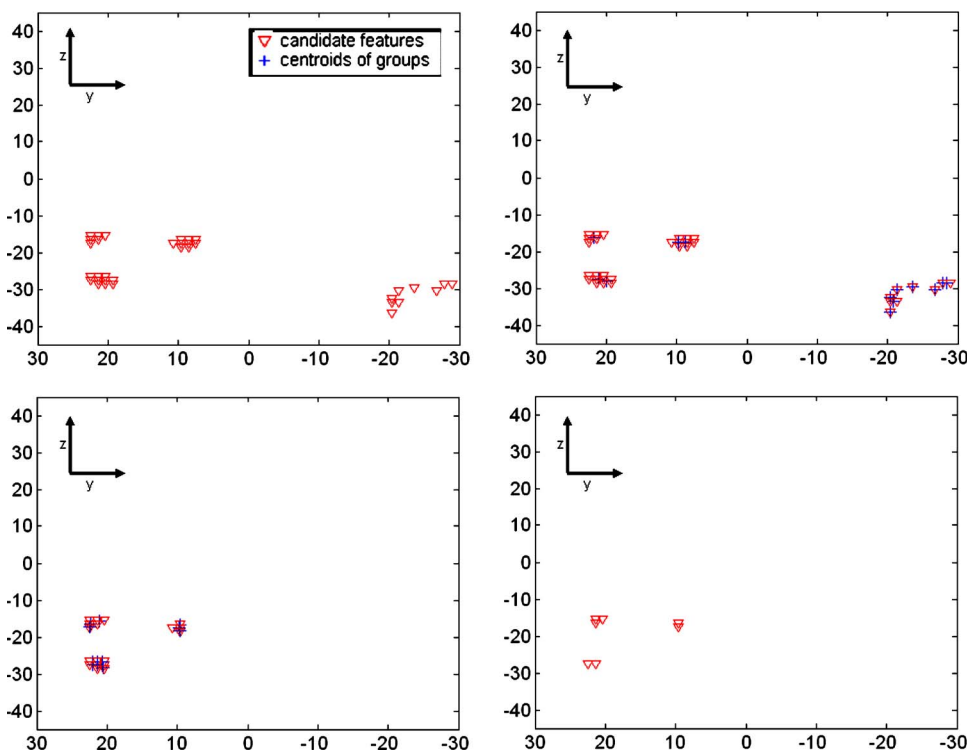


FIG. 7. These four figures provide an illustrative example of the point grouping and geometric matching process. Each plot represents a 2D projection of a 3D scenario that contains features such as those represented using black dots in Fig. 4. Each point in the candidate feature set is denoted by the symbol “ ∇ .” The upper left image shows a set of points obtained after thresholding. In this case, the three leftmost point constellations correspond to three gold seeds. The rightmost constellation represents spurious candidate points due to bone. The application of the first grouping operation yields 13 groups. The centroid of each group is denoted using the symbol “+.” The lower left panel shows the result of the application of the first matching step. Only points belonging to the groups that produce the lowest matching error ϵ_i remain in the feature set. The centroid positions after the subsequent grouping operation also appear in this plot. Finally, the lower right image illustrates the result of the second matching step. These points represent the final point candidates.

TABLE IV. MV datasets obtained using a 6 MV beam.

ID	No. of datasets	1st MVCB— last MVCB (days)	CT to 1st MVCB (days)	No. of correct detections	Total correct (%)	Error RMS (mm)
c1	41	64	18	41	100	0.124–0.815
c2	25	52	12	25	100	0.021–1.233
c3	3	11	6	3	100	0.070–0.355

III.B.1. Volume-of-interest

The volume-of-interest VOI is defined by the user to enclose all possible seed locations. It is specified by a single parameter a_v , which corresponds to the sidelength of a square in the transaxial plane. The VOI thus defined is a rectangular prism of square cross section, centered at the isocenter, that spans the entire support of the image along the patient axis. We choose $a_v=70$ mm.

III.B.2. Thresholding

When gold seeds are imaged using MV beams, their intensity values typically exceed those due to bone by a significant margin. An intensity threshold γI_{\max} that is close to the maximum intensity within the volume of interest is thus appropriate for the thresholding of MV images. We select $\gamma=0.85$.

III.B.3. Connected component analysis

The actual voxel size of a seed depends on the resolution of the imaging modality, i.e., pixel spacing and slice thickness. The patients imaged in this study were implanted retroperitoneally using cylindrical gold seeds having a diameter of 1 mm and a length of 3 mm. The seed volume is thus 2.4 mm^3 . Due to the influences of finite imaging system focal spot size, projection image smoothing, image noise, and the partial-volume effect, the contribution of a seed is spread out over a much larger volume than this in image space. For our megavoltage conebeam pretreatment images we retained as potential seed features those connected groups having a volume of less than $V_{\max}=65 \text{ voxels}^3=69.57 \text{ mm}^3$.

III.B.4. Grouping

The maximum proximity r for group inclusion was initially set to 2.0 mm. Dependent on the actual number of

groups $|G_{ij}|$, the radius is either increased or decreased for the regrouping. We set the upper bound on the number of groups $UB=15$. The lower bound is set to $P=3$.

III.B.5. Transformation constraints

Rotation was constrained to the angular range between $-\pi/8$ and $\pi/8$. This range accommodates all observed prostate rotations. Translation was not constrained.

III.B.6. Matching

The threshold for early termination ϵ_+ in the grouping stage was set to 0.1 mm (RMS) and the threshold for the final error ϵ_{\max} to 2.0 mm. These values provide a positioning tolerance of similar magnitude to delivery system and treatment plan tolerances.

III.B.7. Algorithm implementation

The algorithm was coded in MATLAB (Version 2007b, Mathworks, Inc., Natick, MA) and was executed on a dual 3.2 GHz Intel Xeon (Family 15, Series 4) processor computer system.

IV. RESULTS

The results obtained through application of the algorithm to datasets acquired using the 6 and 4.5 MV beams appear in Tables IV and V, respectively. These tables also indicate the time range in days between the acquisition of a particular patient's first CT scan and first MVCB scan, as well as the number of days between the patient's first and last MVCB scans.

In Fig. 3, a transaxial slice of the planning CT that contains an implanted prostate seed, and a corresponding slice of a megavoltage cone beam pretreatment image are shown. Figures 4 and 5 illustrate intermediate results after the appli-

TABLE V. MV datasets obtained using a 4.5 MV beam.

ID	No. of datasets	1st MVCB— last MVCB (days)	CT to 1st MVCB (days)	No. of correct detections	No. of correct rejections	Total correct (%) validated	Error RMS (mm)	No. of false detections
i1	8	55	7	7	0	100	0.075–0.248	1*
i2	6	40	15	5	0	100	0.101–0.538	1*
i3	6	22	36	4	1*	80	0.538–1.618	1
i4	4	47	10	4	0	100	0.033–0.240	0
i5	3	3	0	3	0	100	0.087–0.107	0
i6	37	67	0	37	0	100	0.052–0.370	0
i7	31	66	0	29	2	100	0.075–0.890	0

cation of two steps of the algorithm. A typical result is presented in Fig. 6. The detected seed locations lie at the centers of the white circles.

The validated detection rate takes into account only those images that were acquired with doses of 1 cGy or above. The images that were excluded from the calculation of this statistic are marked with the superscript “*” in Table V. We justify this cut-off by observing that the seeds in the <1 cGy images are not visible to human observers.

The average processing time per dataset (mean \pm standard deviation) was 4.7 ± 2.8 s for the 6 MV studies and 7.7 ± 2.9 s for the 4.5 MV studies.

V. DISCUSSION

We formulated and demonstrated a practical algorithm for the automatic 3D coregistration of metallic fiducial markers and demonstrated its effectiveness through application to a large number of volumetric megavoltage conebeam datasets.

When applied to images obtained using a commercially available 6 MV imaging system, the algorithm was 100% successful in identifying the fiducial positions (69 studies on three patients).

The failure of registration in one study of series i1 was due to poor image quality. This dataset was produced with dose of 0.3 cGy and was too noisy to allow even manual seed identification.

In series i2, the matching error in one study was due to a very large amount of prostate motion relative to the reference image. One of the fiducial markers was displaced from the entire volume-of-interest in this case.

The result obtained for a very low dose (0.6 cGy) image in series i3 represents a true negative. The algorithm was not able to find a match with an acceptably low error and indicated this to the user.

Unexpectedly, the application of the algorithm to one higher dose (10 cGy) image in series i3 led to a false positive match. More than five seeds were present within the prostate of this patient. On the day this image was acquired, the patient presented with a highly deformed prostate likely due to a full bladder and rectum. The deformation permitted the algorithm to erroneously find an acceptable match to an incorrect marker. Selection of the correct match in this case is not straightforward, even for a trained human operator. Nonetheless, these cases re-enforce the clinical importance of final operator validation.

In series i7, two images were correctly identified as true negatives due to bad image quality. Both images showed artifacts due to missing tomographic projections, so even manual seed identification was not possible.

In the study presented here, the algorithm is employed as a tool for effecting pretreatment IGRT. While it is relatively rare for the prostate gland to move significantly during treatment, motion of over 2 cm is encountered in some cases. A recent study of intrafraction prostate motion determined that displacements of more than 3 mm were observed during one in eight treatment sessions during the first five minutes after initial patient alignment. 10 min after initial alignment, such

motion was observed in a quarter of all treatment episodes.¹⁷ The authors of this study conclude by emphasizing “the importance of initiating treatment shortly after initially positioning the patient.” This supports our focus on the application of our algorithm to pretreatment IGRT, since any time that elapses in between 3D imaging and determination of the correct patient position shift that must be applied increases the probability of prostate motion introducing targeting error. In addition, it appears that the dosimetric consequences of intrafraction prostate motion are rather limited, even when narrow 3 mm margins are employed to extend the clinical target volume (CTV). Langen *et al.* report, as a worst-case observation, a patient in which prostate motion would have caused the CTV to extended more than 4 mm from the planned target volume for 3.1% of the treatment time had the patient been planned with 3 mm margins.¹⁷ Such small dosimetric consequences are far less likely to influence treatment outcome than inaccurate initial patient setup.

The impact of any potential transient or systematic motion error on the method presented here is identical to its impact on contemporary standard workflows that involve obtaining orthogonal portal images for pretreatment seed detection. An ideal IGRT system would continue to track the fiducials during treatment. Future delivery systems are envisaged which make use of an integrated magnetic resonance imaging unit or multisource x-ray tomosynthesis system to provide real-time 3D imaging during treatment. While tomography systems provide rather limited 3D information in general owing to poor depth resolution, the latter is greatly enhanced for objects characterized by high spatial frequencies, such as fiducial markers. There is no reason to expect that this algorithm will not perform as well as demonstrated here when applied to magnetic resonance images or to tomography images acquired at diagnostic x-ray energies. Streaking of gold seeds due to beam-hardening may require the upward adjustment of the seed volume parameter. However, this is unlikely to be problematic since the higher spatial resolution will decrease the partial volume effect and better contrast between bone and gold will decrease the probability of identifying bony features as fiducials.

Each MVCB image is acquired over a period of less than 45 s. The influence of motion during this period on algorithm performance is difficult to predict, since its effects on image quality are highly dependent on the amount of motion that occurs and the time at which it occurs during the imaging process. However, since no attempt was made to regulate prostate motion in the studies to which the algorithm is applied here, we expect that such motion will not increase the average error rate from that observed in this investigation.

A major determinant of the practical utility of automatic algorithms is their sensitivity to the values of algorithm parameters such as thresholds and parameter bounds. While we have demonstrated that the present algorithm is able to provide excellent performance over a large number of datasets using a single set of parameter values, a discussion of the sensitivity of the algorithm to different parameter values selections is in order.

The algorithm requires the user to set the volume-of-interest in order to bound the space inside which the algorithm will search for fiducial markers. Since the planning CT is available, it is relatively straightforward to select a reasonable value for this parameter based on the known maximum radius of the seeds from the isocenter and an allowance for the maximum expected patient setup error. While selection of an overly large value will increase processing time, it is unlikely to increase the detection error rate unless additional non-relevant markers are present in the expanded VOI.

The selection of the threshold intensity value for the initial identification of highly attenuating features within the search volume is a key algorithm parameter. Implementation of a hard threshold is unattractive, since the HU values of seeds may vary due to beam characteristics, incorrect HU scaling, cupping artifacts, noise levels, and partial volume effects. The present algorithm makes use of an adaptive thresholding technique that adjusts the threshold based on the size and nature of the feature set yielded by a particular threshold setting. When the initial threshold is set too high, too few feature sets are detected. When it is set too low, too many features are identified with the result that processing time and false positive rates are increased. The adaptive threshold adjustment strategy provides a practical mechanism for reducing computation time and increasing detection accuracy.

Fundamental to the successful automatic adjustment of the intensity threshold is valid feature identification. We employ CCA to remove features having a volume in excess of those expected for fiducial markers. CCA is applied only for removing objects such as prosthetics and bone from the set of potential features and does not play a role in the grouping of remaining feature points into seeds. The algorithm is not very sensitive to the precise selection of the parameter limiting the maximum volume of feasible seed features within the search space. This is because highly attenuating objects other than seeds are normally much larger than the fiducial markers.

A major feature of the algorithm is the sophisticated iterative grouping process. It is important to note that while such a complex process is not essential for the successful identification of seeds in most imaging studies, it does provide necessary robustness against noise and partial volume effects. An additional advantage of the scheme we employ is that initial proximity-based feature grouping greatly accelerates algorithm execution by reducing the number of possible point correspondences in the matching process.

A good choice of the parameter limiting the maximum proximity of voxels for group inclusion leads to fewer regrouping iterations and thus decreases processing time. However, since these regrouping steps constitute only a small fraction of the total computational burden, the influence of this factor on efficiency is minor. The UB of the matching process plays a far more significant role. When the value of UB is too small, detection accuracy is decreased, while values that are too high will increase execution time considerably, since the algorithm performs an exhaustive search among all possible combinations.

The parameters constraining the geometric transform must be set after consideration of typical motion and set-up errors encountered in the specific clinical application. It should be possible to employ a single set of limits for a particular IGRT protocol, such as pretreatment positioning for prostate radiotherapy. These limits might be obtained from published distributions of prostate motion and set-up errors.¹⁶ The transformation constraint limits the maximal movement between source and target images. The transformation error bound distinguishes between successful detection on the one hand and incorrect detection or unacceptably large deformation on the other.

ACKNOWLEDGMENTS

The authors would like to thank Dr. Jean Pouliot of the Department of Radiation Oncology at the University of California, San Francisco, and Dr. Morris Geffen at the Savannah Oncology Center of Savannah, Georgia, for providing the patient data for our algorithm evaluation.

^{a)}Electronic mail: martin.koch@informatik.stud.uni-erlangen.de

¹E. Vigneault, J. Pouliot, J. Laverdière, J. Roy, and M. Dorion, "Electronic portal imaging device detection of radioopaque markers for the evaluation of prostate position during megavoltage irradiation: A clinical study," *Int. J. Radiat. Oncol., Biol., Phys.* **37**, 205–212 (1997).

²A. Nederveen, J. Lagendijk, and P. Hofman, "Detection of fiducial gold markers for automatic on-line megavoltage position verification using a marker extraction kernel (MEK)," *Int. J. Radiat. Oncol., Biol., Phys.* **47**, 1435–1442 (2000).

³E. J. Harris, H. A. McNair, and P. M. Evans, "Feasibility of fully automated detection of fiducial markers implanted into the prostate using electronic portal imaging: a comparison of methods," *Int. J. Radiat. Oncol., Biol., Phys.* **66**, 1263–1270 (2006).

⁴H. Dehnad, A. J. Nederveen, U. A. van der Heide, R. J. A. van Moorseelaar, P. Hofman, and J. J. Lagendijk, "Clinical feasibility study for the use of implanted gold seeds in the prostate as reliable positioning markers during megavoltage irradiation," *Radiother. Oncol.* **67**, 295–302 (2003).

⁵P. Kupelian, T. Willoughby, A. Mahadevan, T. Djemil, G. Weinstein, S. Jani, C. Enke, T. Solberg, N. Flores, D. Liu, D. Beyer, and L. Levine, "Multi-institutional clinical experience with the Calypso system in localization and continuous, real-time monitoring of the prostate gland during external radiotherapy," *Int. J. Radiat. Oncol., Biol., Phys.* **67**, 1088–1098 (2007).

⁶S. Pouliot, A. Zaccarin, D. Laurendeau, and J. Pouliot, "Automatic detection of three radio-opaque markers for prostate targeting using EPID during external radiation therapy," *Proceedings of the 2001 International Conference on Image Processing*, Oct. 2001, 2, pp. 857–860 (unpublished).

⁷M. Miften, O. Gayou, B. Reitz, R. Fuhrer, B. Leicher, and D. S. Parda, "IMRT planning and delivery incorporating daily dose from mega-voltage cone-beam computed tomography imaging," *Med. Phys.* **34**, 3760–3767 (2007).

⁸H. Chui and A. Rangarajan, "A new algorithm for non-rigid point matching," *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2000, Vol. 2, pp. 44–51 (unpublished).

⁹M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Commun. ACM* **24**, 381–395 (1981).

¹⁰R. Sedgewick, *Algorithms in C* (Addison-Wesley, Reading, MA, 1998).

¹¹R. M. Haralick and L. G. Shapiro, *Computer and Robot Vision* (Addison-Wesley, Reading, MA, 1992).

¹²E. Polak, *Optimization: Algorithms and Consistent Approximation* (Springer, Berlin, 1997).

¹³A. Schwartz and E. Polak, "Family of projected descent methods for optimization problems with simple bounds," *J. Optim. Theory Appl.* **92**, 1–31 (1997).

¹⁴B. A. Faddegon, I. C. Hsu, F. Ghelkansarai, A. Bani-Hashemi, and J.

- Pouliot, "Significant improvement in megavoltage cone-beam CT in a side-by-side comparison of a cadaver imaged with x-rays from low and high-Z targets," *Int. J. Radiat. Oncol., Biol., Phys.* **66**, S147 (2006).
- ¹⁵T. Schiwietz, S. Bose, J. Maltz, and R. Westermann, "A fast and high-quality cone beam reconstruction pipeline using the GPU," *SPIE Medical Imaging 2007: Physics of Medical Imaging*, Vol. 6501, p. 65105H (2007).
- ¹⁶J. Wu, T. Haycocks, H. Alasti, G. Ottewell, N. Middlemiss, M. Abdoell, P. Warde, A. Toi, and C. Catton, "Positioning errors and prostate motion during conformal prostate radiotherapy using on-line isocentre set-up verification and implanted prostate markers," *Radiother. Oncol.* **61**, 127–133 (2001).
- ¹⁷K. M. Langen, T. R. Willoughby, S. L. Meeks, A. Santhanam, A. Cunningham, L. Levine, and P. A. Kupelian, "Observations on real-time prostate gland motion using electromagnetic tracking," *Int. J. Radiat. Oncol., Biol., Phys.* **71**, 1084–1090 (2008).