

Real-time 2D/3D Deformable Registration Using Metric Learning

Chen-Rui Chou¹ and Stephen Pizer^{1,2}

¹ Department of Computer Science, University of North Carolina at Chapel Hill,
Chapel Hill, NC 27599, USA,
cchou@cs.unc.edu

² Department of Radiation Oncology, University of North Carolina at Chapel Hill,
Chapel Hill, NC 27599, USA

Abstract. We present a novel 2D/3D deformable registration method, called Registration Efficiency and Accuracy through Learning Metric on Shape (*REALMS*), that can support real-time Image-Guided Radiation Therapy (*IGRT*). The method consists of two stages: planning-time learning and registration. In the planning-time learning, it firstly models the patient’s 3D deformation space from the patient’s time-varying 3D planning images using a low-dimensional parametrization. Secondly, it samples deformation parameters within the deformation space and generates corresponding simulated projection images from the deformed 3D image. Finally, it learns a Riemannian metric in the projection space for each deformation parameter. The learned distance metric forms a Gaussian kernel of a kernel regression that minimizes the leave-one-out regression residual of the corresponding deformation parameter. In the registration, *REALMS* interpolates the patient’s 3D deformation parameters using the kernel regression with the learned distance metrics. Our test results showed that *REALMS* can localize the tumor in 10.89 ms (91.82 fps) with 2.56 ± 1.11 mm errors using a single projection image. These promising results show *REALMS*’s high potential to support real-time, accurate, and low-dose *IGRT*.

1 Introduction

Tumor localization in 3D is the main goal of Image-guided Radiation Therapy (*IGRT*). It is usually accomplished by computing the patient’s treatment-time 3D deformations based on an on-board imaging system, usually x-ray. The treatment-time 3D deformations can be computed by doing image registration between the treatment-time reconstructed 3D image and the treatment-planning 3D image (3D/3D registration) or between the treatment-time on-board projection images and the treatment-planning 3D image (2D/3D registration). Recent advances of the *IGRT* registration methods emphasize real-time computation and low-dose image acquisition. Russakoff et al. [1,2], Khamene et al. [3], Munbodh et al. [4], Li et al. [5,6] rejected the time-consuming 3D/3D registration and performed 2D/3D registration by optimizing similarity functions defined in

the projection domain. Other than the optimization-based methods, Chou et al. [7,8] recently introduced a faster and low-dose 2D/3D image registration by using a linear operator that approximates the deformation parameters. However, all of the above registration methods involve computationally demanding production of Digitally-Reconstructed Radiographs (*DRRs*) in each registration iteration (e.g., 15ms on a modern GPU to produce a 256×256 DRR from a $256 \times 256 \times 256$ volume [9]), which makes them difficult to be extended to support real-time (> 30 fps) image registration.

We present a novel real-time 2D/3D registration method, called Registration Efficiency and Accuracy through Learning Metric on Shape (*REALMS*), that does not require DRR production in the registration. It calculates the patient’s treatment-time 3D deformations by kernel regression. Specifically, each of the patient’s deformation parameters is interpolated using a weighting Gaussian kernel on that parameter’s training case values. In each training case, its parameter value is associated with a corresponding training projection image. The Gaussian kernel is formed from distances between training projection images. This distance for the parameter in question involves a Riemannian metric on projection image differences. At planning time, *REALMS* learns the parameter-specific metrics from the set of training projection images using a Leave-One-Out (*LOO*) training.

To the best of our knowledge, *REALMS* is the first 2D/3D deformable registration method that achieves real-time (> 30 fps) performance. *REALMS* uses the metric learning idea firstly introduced in Weinberger and Tesauro [10] to tackle the 2D/3D image registration problem. Particularly, in order to make the metric learning work for the high dimensional ($D \gg 10^3$) projection space, *REALMS* uses a specially-designed initialization approximated by linear regression. The results have led to substantial error reduction when the special initialization is applied.

The rest of the paper is organized as follows: In section 2, we describe *REALMS*’s novel registration scheme that uses kernel regression. In section 3, we describe its deformation space modeling approach for generating training samples in the deformation space. In section 4, we describe the metric learning scheme and the specialized initialization in *REALMS*. We show our synthetic and real results in section 5. Finally, we discuss the results and conclude in section 6.

2 2D/3D Registration Framework

In this section, we describe *REALMS*’s 2D/3D registration framework. *REALMS* uses kernel regression (eq. 1) to interpolate the patient’s n 3D deformation parameters $\mathbf{c} = (c^1, c^2, \dots, c^n)$ separately from the on-board projection image $\Psi(\theta)$ where θ is the projection angle. It uses a Gaussian kernel $K_{\mathbf{M}^i, \sigma^i}$ with the width σ^i and a metric tensor \mathbf{M}^i on projection intensity differences to interpolate the patient’s i^{th} deformation parameter c^i from a set of N training projection images $\{\mathbf{P}(I \circ T(\mathbf{c}_\kappa); \theta) \mid \kappa = 1, 2, \dots, N\}$ simulated at planning time. Specifically, the

training projection image, $\mathbf{P}(I \circ T(\mathbf{c}_\kappa); \theta)$, is the DRR of a 3D image deformed from the patient's planning-time 3D mean image I with sampled deformation parameters $\mathbf{c}_\kappa = (c_\kappa^1, c_\kappa^2, \dots, c_\kappa^n)$. T and \mathbf{P} are the warping and the DRR operators, respectively. \mathbf{P} simulates the DRRs according to the treatment-time imaging geometry, e.g., the projection angle θ .

In the treatment-time registration, each deformation parameter c^i in \mathbf{c} can be estimated with the following kernel regression:

$$c^i = \frac{\sum_{\kappa=1}^N c_\kappa^i \cdot K_{\mathbf{M}^i, \sigma^i}(\Psi(\theta), \mathbf{P}(I \circ T(\mathbf{c}_\kappa); \theta))}{\sum_{\kappa=1}^N K_{\mathbf{M}^i, \sigma^i}(\Psi(\theta), \mathbf{P}(I \circ T(\mathbf{c}_\kappa); \theta))}, \quad (1)$$

$$K_{\mathbf{M}^i, \sigma^i}(\Psi(\theta), \mathbf{P}(I \circ T(\mathbf{c}_\kappa); \theta)) = \frac{1}{\sqrt{2\pi}\sigma^i} e^{-\frac{d_{\mathbf{M}^i}^2(\Psi(\theta), \mathbf{P}(I \circ T(\mathbf{c}_\kappa); \theta))}{2(\sigma^i)^2}}, \quad (2)$$

$$d_{\mathbf{M}^i}^2(\Psi(\theta), \mathbf{P}(I \circ T(\mathbf{c}_\kappa); \theta)) = (\Psi(\theta) - \mathbf{P}(I \circ T(\mathbf{c}_\kappa); \theta))^\top \mathbf{M}^i (\Psi(\theta) - \mathbf{P}(I \circ T(\mathbf{c}_\kappa); \theta)), \quad (3)$$

where $K_{\mathbf{M}^i, \sigma^i}$ is a Gaussian kernel (kernel width = σ^i) that uses a Riemannian metric \mathbf{M}^i in the squared distance $d_{\mathbf{M}^i}^2$ and gives the weights for the parameter interpolation in the regression. The minus signs in eq. 3 denote pixel-by-pixel intensity subtraction.

We describe in section 3 how REALMS, at planning time, parameterizes the deformation space and describe in section 4 how it learns the metric tensor \mathbf{M}^i and decides the kernel width σ^i .

3 Deformation Modeling at Planning Time

REALMS limits the deformation to a shape space. It models deformations as a linear combination of a set of basis deformations calculated through PCA analysis. In our target problem – lung IGRT, a set of Respiratory-Correlated CTs (*RCCTs*, dimension: $512 \times 512 \times 120$) $\{J_\tau \mid \tau = 1, 2, \dots, 10\}$ are available at planning time. From these a mean image $I = \bar{J}$ and a set of deformations ϕ_τ between J_τ and \bar{J} can be computed. The basis deformations can then be chosen to be the primary eigenmodes of a PCA analysis on the ϕ_τ .

3.1 Deformation Shape Space and Mean Image Generation

REALMS computes a respiratory Fréchet mean image \bar{J} from the RCCT dataset via an *LDDMM* (Large Deformation Diffeomorphic Metric Mapping) framework described in Lorenzen et al. [11]. The Fréchet mean \bar{J} , as well as the diffeomorphic deformations ϕ from the mean \bar{J} to each image J_τ , are computed using a fluid-flow distance metric:

$$\bar{J} = \underset{J}{\operatorname{argmin}} \sum_{\tau=1}^{10} \int_0^1 \int_\Omega \|v_{\tau, \gamma}(x)\|^2 dx d\gamma + \frac{1}{s^2} \int_\Omega \|J(\phi_\tau^{-1}(x)) - J_\tau(x)\|^2 dx, \quad (4)$$

where $J_\tau(x)$ is the intensity of the pixel at position x in the image J_τ , $v_{\tau,\gamma}$ is the fluid-flow velocity field for the image J_τ in flow time γ , s is the weighting variable on the image dissimilarity, and $\phi_\tau(x)$ describes the deformation at the pixel location x : $\phi_\tau(x) = x + \int_0^1 v_{\tau,\gamma}(x) d\gamma$.

3.2 Statistical Analysis

With the diffeomorphic deformation set $\{\phi_\tau \mid \tau = 1, 2, \dots, 10\}$ calculated, our method finds a set of linear deformation basis vectors ϕ_{pc}^i by PCA analysis. The scores λ_τ^i on each ϕ_{pc}^i yield ϕ_τ in terms of these basis vectors.

$$\phi_\tau = \bar{\phi} + \sum_{i=1}^{10} \lambda_\tau^i \cdot \phi_{pc}^i \quad (5)$$

We choose a subset of n eigenmodes that captures more than 95% of the total variation. Then we let the n scores form the n -dimensional parametrization \mathbf{c} .

$$\mathbf{c} = (c^1, c^2, \dots, c^n) = (\lambda^1, \lambda^2, \dots, \lambda^n) \quad (6)$$

For most of our target problems, $n = 3$ satisfies the requirement.

4 Metric Learning at Planning Time

4.1 Metric Learning and Kernel Width Selection

REALMS learns a metric tensor \mathbf{M}^i with a corresponding kernel width σ^i for the patient's i^{th} deformation parameter c^i using a Leave-One-Out (*LOO*) training strategy. At planning time, it samples a set of N deformation parameter tuples $\{\mathbf{c}_\kappa = (c_\kappa^1, c_\kappa^2, \dots, c_\kappa^n) \mid \kappa = 1, 2, \dots, N\}$ to generate training projection images $\{\mathbf{P}(I \circ T(\mathbf{c}_\kappa); \theta) \mid \kappa = 1, 2, \dots, N\}$ where their associated deformation parameters are sampled uniformly within three standard deviations of the scores λ observed in the RCCTs. For each deformation parameter c^i in \mathbf{c} , REALMS finds the best pair of the metric tensor $\mathbf{M}^{i\dagger}$ and the kernel width $\sigma^{i\dagger}$ that minimizes the sum of squared LOO regression residuals \mathcal{L}_{c^i} among the set of N training projection images:

$$\mathbf{M}^{i\dagger}, \sigma^{i\dagger} = \arg \min_{\mathbf{M}^i, \sigma^i} \mathcal{L}_{c^i}(\mathbf{M}^i, \sigma^i), \quad (7)$$

$$\mathcal{L}_{c^i}(\mathbf{M}^i, \sigma^i) = \sum_{\kappa=1}^N \left(c_\kappa^i - \hat{c}_\kappa^i(\mathbf{M}^i, \sigma^i) \right)^2, \quad (8)$$

$$\hat{c}_\kappa^i(\mathbf{M}^i, \sigma^i) = \frac{\sum_{\chi \neq \kappa} c_\chi^i \cdot K_{\mathbf{M}^i, \sigma^i}(\mathbf{P}(I \circ T(\mathbf{c}_\kappa); \theta), \mathbf{P}(I \circ T(\mathbf{c}_\chi); \theta))}{\sum_{\chi \neq \kappa} K_{\mathbf{M}^i, \sigma^i}(\mathbf{P}(I \circ T(\mathbf{c}_\kappa); \theta), \mathbf{P}(I \circ T(\mathbf{c}_\chi); \theta))}, \quad (9)$$

where $\hat{c}_\kappa^i(\mathbf{M}^i, \sigma^i)$ is the estimated value for parameter c_κ^i interpolated by the metric tensor \mathbf{M}^i and the kernel width σ^i from the training projection images χ other than κ ; \mathbf{M}^i needs to be a positive semi-definite (*p.s.d*) matrix to fulfill the pseudo-metric constraint; and the kernel width σ^i needs to be a positive real number.

To avoid high-dimensional optimization over the constrained matrix \mathbf{M}^i , we structure the metric tensor \mathbf{M}^i as a rank-1 matrix formed by a basis vector \mathbf{a}^i : $\mathbf{M}^i = \mathbf{a}^i \mathbf{a}^{i\top}$. Therefore, we can transform eq. 7 into a optimization over the unit vector \mathbf{a}^i where $\|\mathbf{a}^i\|_2 = 1$:

$$\mathbf{a}^{i\dagger}, \sigma^{i\dagger} = \underset{\mathbf{a}^i, \sigma^i}{\operatorname{argmin}} \mathcal{L}_{c^i}(\mathbf{a}^i \mathbf{a}^{i\top}, \sigma^i) \quad (10)$$

Then we can rewrite the squared distance $d_{\mathbf{M}^i}^2 = d_{\mathbf{a}^i \mathbf{a}^{i\top}}^2$ used in the Gaussian kernel $K_{\mathbf{M}^i, \sigma^i}$ as follows:

$$d_{\mathbf{a}^i \mathbf{a}^{i\top}}^2(\mathbf{P}(I \circ T(\mathbf{c}_\kappa); \theta), \mathbf{P}(I \circ T(\mathbf{c}_\chi); \theta)) = (\mathbf{a}^{i\top} \cdot \mathbf{r}_{\kappa, \chi})^\top (\mathbf{a}^{i\top} \cdot \mathbf{r}_{\kappa, \chi}), \quad (11)$$

$$\mathbf{r}_{\kappa, \chi} = \mathbf{P}(I \circ T(\mathbf{c}_\kappa); \theta) - \mathbf{P}(I \circ T(\mathbf{c}_\chi); \theta), \quad (12)$$

where $\mathbf{r}_{\kappa, \chi}$ is a vector of intensity differences between projection images generated by parameters \mathbf{c}_κ and \mathbf{c}_χ ; and \mathbf{a}^i is a metric basis vector where the magnitude of the inner product of \mathbf{a}^i and the intensity difference vector $\mathbf{r}_{\kappa, \chi}$, $\mathbf{a}^{i\top} \cdot \mathbf{r}_{\kappa, \chi}$ gives the Riemannian distance for the parameter c^i (eq. 11).

The learned metric basis vector $\mathbf{a}^{i\dagger}$ and the selected kernel width $\sigma^{i\dagger}$ form a weighting kernel $K_{\mathbf{a}^{i\dagger} \mathbf{a}^{i\dagger\top}, \sigma^{i\dagger}}$ to interpolate the parameter c^i in the registration (see eq. 1).

4.2 Linear-Regression Implied Initial Metric

Since the residual functional \mathcal{L} (see eq. 7) that we want to minimize is non-convex, a good initial guess of the metric basis vector \mathbf{a} is essential. Therefore, REALMS uses a vector \mathbf{w}^i as an initial guess of the metric basis vector \mathbf{a}^i for the parameter c^i . Let $\mathbf{W} = (\mathbf{w}^1 \mathbf{w}^2 \dots \mathbf{w}^n)$ list these initial guesses. The matrix \mathbf{W} is approximated by a multivariate linear regression (eq. 13 and eq. 14) between the projection difference matrix $\mathbf{R} = (\mathbf{r}_1 \mathbf{r}_2 \dots \mathbf{r}_N)^\top$ and the parameter differences matrix $\Delta \mathbf{C}$. In particular, the projection difference vector $\mathbf{r}_\kappa = \mathbf{P}(I \circ T(\mathbf{c}_\kappa); \theta) - \mathbf{P}(I; \theta)$ is the intensity differences between the DRRs calculated from the deformed image $I \circ T(\mathbf{c}_\kappa)$ and the DRRs calculated from the mean image I (where $\mathbf{c} = \mathbf{0}$).

$$\Delta \mathbf{C} = \begin{pmatrix} c_1^1 & c_1^2 & \dots & c_1^n \\ c_2^1 & c_2^2 & \dots & c_2^n \\ \vdots & \vdots & \ddots & \vdots \\ c_N^1 & c_N^2 & \dots & c_N^n \end{pmatrix} - \mathbf{0} \approx \begin{pmatrix} \mathbf{r}_1^\top \\ \mathbf{r}_2^\top \\ \vdots \\ \mathbf{r}_N^\top \end{pmatrix} \cdot (\mathbf{w}^1 \mathbf{w}^2 \dots \mathbf{w}^n) \quad (13)$$

$$\mathbf{W} = (\mathbf{R}^\top \mathbf{R})^{-1} \mathbf{R}^\top \Delta \mathbf{C} \quad (14)$$

The inner product of the matrix \mathbf{W} , calculated by the pseudo-inverse in eq. 14, and the projection intensity difference matrix \mathbf{R} , $\mathbf{W}^\top \mathbf{R}$, gives the best linear approximation of the parameter differences $\Delta \mathbf{C}$. Therefore, we use \mathbf{w}^i as the initial guess of the metric basis vector \mathbf{a}^i for the parameter c^i .

4.3 Optimization Scheme

REALMS uses a two-step scheme to optimize the metric basis vector \mathbf{a}^i and the kernel width σ^i in eq. 10.

First, for each candidate kernel width σ^i , it optimizes the metric basis vector \mathbf{a}^i using the quasi-Newton method (specifically, the BFGS method) with the vector \mathbf{w}^i as the initialization. The gradient of the function \mathcal{L}_{c^i} with respect to \mathbf{a}^i can be stated as

$$\frac{\partial \mathcal{L}_{c^i}}{\partial \mathbf{a}^i} = \frac{2\sqrt{2}}{\sigma^i} \mathbf{a}^i \sum_{\kappa=1}^N (\hat{c}_\kappa^i - c_\kappa^i) \sum_{\chi=1}^N (\hat{c}_\chi^i - c_\chi^i) K_{\mathbf{a}^i \mathbf{a}^i \tau, \sigma^i}(\mathbf{P}(I \circ T(\mathbf{c}_\kappa); \theta), \mathbf{P}(I \circ T(\mathbf{c}_\chi); \theta)) \mathbf{r}_{\kappa, \chi} \mathbf{r}_{\kappa, \chi}^\top \quad (15)$$

Second, REALMS selects a kernel width $\sigma^{i\dagger}$ among the candidate kernel widths where its learned metric basis vector $\mathbf{a}^{i\dagger}$ yields minimum LOO regression residuals \mathcal{L}_{c^i} for parameter c^i .

4.4 Projection Normalization

To account for variations caused by x-ray scatter that produces inconsistent projection intensities, REALMS normalizes both the training projection images $\mathbf{P}(I \circ T(\mathbf{c}_\kappa); \theta)$ and the on-board projection image $\Psi(\theta)$. In particular, it uses the localized Gaussian normalization introduced in Chou et al. [8], which has shown promise in removing the undesired scattering artifacts.

5 Results

5.1 Synthetic Tests

We used coronal DRRs (dimension: 64×48) of the target CTs as synthetic on-board cone-beam projection images. The target CTs were deformed from the patient's Fréchet mean CT by normally distributed random samples of the first three deformation parameters.¹ We generated 600 synthetic test cases from 6 lung datasets and measured the registration quality by the average *mTRE* (mean Target Registration Error) over all cases and all voxels at tumor sites.

¹ In our lung datasets, the first three deformation parameters captured more than 95% lung variation observed in their RCCTs.

With REALMS's registrations, the average mTRE and its standard deviation are down from 6.89 ± 3.53 mm to 0.34 ± 0.24 mm using $N = 125$ training projection images. The computation time for each registration is 11.39 ± 0.73 ms (87.79 fps) on Intel Core2 Quad CPU Q6700. As shown in figure 1, REALMS reduces the minimum errors produced by kernel regressions that use the Euclidean metric ($\mathbf{M}^i = \mathbf{I}$).

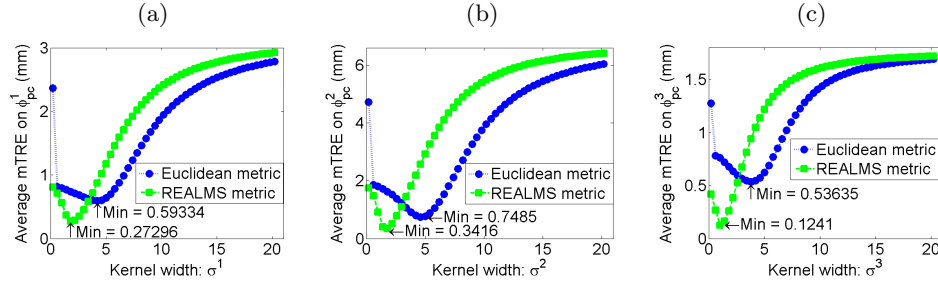


Fig. 1. Average mTREs over 600 test cases projected onto the (a) first, (b) second, and (c) third deformation basis vector versus the candidate kernel widths using $N = 125$ training projection images.

Figure 2 shows the computation time and registration accuracy tradeoff in REALMS.




Fig. 2. (a) Time and (b) accuracy v.s. the number of training projection images N .

5.2 Real Tests

We tested REALMS on 6 lung datasets with an on-board CBCT system where a *single* coronal on-board CB projection (dimension downsampled to 64×48 for efficient computation) at both *EE* (End-Expiration) and *EI* (End-Inspiration) phases were used for the testing. See the top image of figure 4(b) for illustration. For each dataset, we generated $N = 125$ training DRRs to learn the metrics and select optimal interpolation kernel widths. The learned metrics and the selected kernel widths were used to estimate deformation parameters for the testing EE and EI on-board projections. The estimated CTs were deformed from

the Fréchet mean CT with the estimated deformation parameters. The results were validated with reconstructed CBCTs at target phases.² Table 1 shows the 3D Tumor Centroid Differences ($TCDs$) between REALMS-estimated CTs and the reconstructed CBCTs at the same respiratory phases. Tumor centroids were computed via Snake active segmentations. As shown in table 1, REALMS reduces the TCD from 5.58 ± 3.14 mm to 2.56 ± 1.11 mm in 10.89 ± 0.26 ms (91.82 fps).

dataset#	TCD at EE phase (mm)	TCD at EI phase (mm)	Time (ms)
1	2.42 (9.70)	4.06 (7.45)	10.40
2	3.60 (4.85)	3.60 (4.89)	10.92
3	2.30 (8.71)	3.60 (4.03)	10.91
4	1.27 (2.69)	2.80 (2.29)	10.91
5	0.70 (9.89)	3.28 (8.71)	11.15
6	1.98 (2.03)	1.12 (1.72)	11.08

Table 1. Tumor Centroid Differences (TCD) after REALMS’s registration at EE and EI phases of 6 lung datasets. Numbers inside the parentheses are the initial TCDs.

Figure 3 illustrates an example REALMS registration on a lung dataset where the tumor, the diaphragm, and most of the soft tissues are correctly aligned.




Fig. 3. (a) Image overlay of the reconstructed CBCT at EE phase (red) and the Fréchet mean CT (green) (b) Image overlay of the reconstructed CBCT at EE phase (red) and the REALMS-estimated CT (green) calculated from an on-board cone-beam projection image at EE phase. The yellow areas are the overlapped region.

5.3 The Learned Metric Basis Vector

The learned metric basis vector $\mathbf{a}^{i\dagger}$ will emphasize projection pixels that are significant for the distance calculation of the deformation parameter c^i (e.g. give

² The CBCTs were reconstructed by the retrospectively-sorted CB projections at target breathing phases.

high positive or high negative values). As shown in figure 4(a), the learned metric basis vector $\mathbf{a}^{1\dagger}$ emphasized the diaphragm locations and the lung boundaries as its corresponding deformation basis vector ϕ_{pc}^1 covers the expansion and contraction motion of the lung. See the bottom image of figure 4(b) for illustration.

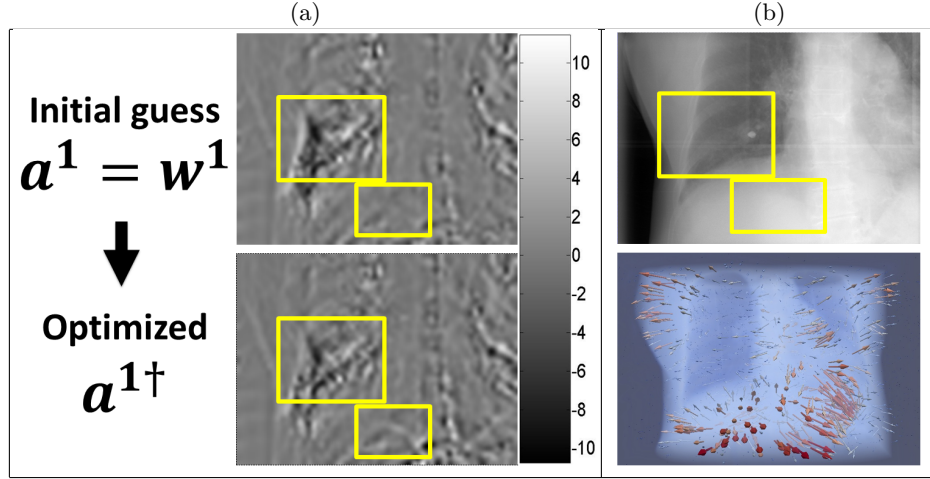


Fig. 4. (a) Initial guess of the metric basis vector $\mathbf{a}^1 = \mathbf{w}^1$ (top) and the optimized metric basis vector $\mathbf{a}^{1\dagger}$ (bottom) of a lung dataset. They are re-shaped into projection image domain for visualization. As shown in the figure, the diaphragm locations and the lung boundaries (yellow boxes) were emphasized after metric learning. (b) Top: a coronal on-board CB projection at EE phase of the lung dataset used in (a). The yellow boxes in (a) and (b) correspond to the same 2D locations. Bottom: the first deformation basis vector ϕ_{pc}^1 (the color arrows indicate heat maps of the deformation magnitudes) overlaid with the volume rendering of the Fréchet mean CT of the lung dataset used in (a). For this dataset, ϕ_{pc}^1 covers the expansion and contraction motion of the lung.

6 Conclusion and Discussion

This paper presents an accurate and real-time 2D/3D registration method, REALMS, that estimates 3D deformation parameters from a single projection image using kernel regressions with learned rank-1 projection distance metrics. The learned distance metrics are optimized with an initialization approximated by linear regression that we found, is essential to the success of this high dimensional metric learning. Without this special initialization, the optimization would have easily converged to local minimum and thus produce wrong distance metrics. With this special initialization, the regression estimation on both synthetic and real test cases showed its good promise in supporting real-time and low-dose IGRT by using a single projection image. In this paper, we use highly down-sampled projection images for efficient learning at planning time. To support efficient learn-

ing for projection images of higher dimensions, the future work of REALMS will incorporate neighborhood approximation methods in the leave-one-out training such that the computation complexity will be reduced from $O(N^2)$ to $O(kN)$ if only k nearest training neighbors are considered for the regression estimation.

References

1. Russakoff, D., Rohlfing, T., Maurer, C.: Fast intensity-based 2D-3D image registration of clinical data using light fields. In: Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on. Volume 1. (2003) 416–422
2. Russakoff, D.B., Rohlfing, T., Mori, K., Rueckert, D., Ho, A., Adler, J.R., Maurer, C.R.: Fast generation of digitally reconstructed radiographs using attenuation fields with application to 2d-3d image registration. *IEEE Transactions on Medical Imaging* **24** (2005) 1441–1454
3. Khamene, A., Bloch, P., Wein, W., Svatos, M., Sauer, F.: Automatic registration of portal images and volumetric ct for patient positioning in radiation therapy. *Medical Image Analysis* **10** (2006) 96–112
4. Munbodh, R., Jaffray, D.A., Moseley, D.J., Chen, Z., Knisely, J.P.S., Cathier, P., Duncan, J.S.: Automated 2d-3d registration of a radiograph and a cone beam ct using line-segment enhancement. *Medical Physics* **33** (2006) 1398–1411
5. Li, R., Jia, X., Lewis, J.H., Gu, X., Folkerts, M., Men, C., Jiang, S.B.: Real-time volumetric image reconstruction and 3d tumor localization based on a single x-ray projection image for lung cancer radiotherapy. *Medical Physics* **37** (2010) 2822–2826
6. Li, R., Lewis, J.H., Jia, X., Gu, X., Folkerts, M., Men, C., Song, W.Y., Jiang, S.B.: 3d tumor localization through real-time volumetric x-ray imaging for lung cancer radiotherapy. *Medical Physics* **38** (2011) 2783–2794
7. Chou, C.R., Frederick, B., Chang, S., Pizer, S.: A Learning-Based patient repositioning method from Limited-Angle projections. In: Brain, Body and Machine. Volume 83 of *Advances in Soft Computing*. Springer (2010) 83–94
8. Chou, C.R., Frederick, B., Liu, X., Mageras, G., Chang, S., Pizer, S.: Claret: A fast deformable registration method applied to lung radiation therapy. In: Fourth International (MICCAI) Workshop on Pulmonary Image Analysis. (2011) 113–124
9. Miao, S., Liao, R., Zheng, Y.: A hybrid method for 2-d/3-d registration between 3-d volumes and 2-d angiography for trans-catheter aortic valve implantation (tavi). In: ISBI. (2011) 1215–1218
10. Weinberger, K., Tesauro, G.: Metric learning for kernel regression. In: Eleventh international conference on artificial intelligence and statistics. (2007) 608–615
11. Lorenzen, P., Prastawa, M., Davis, B., Gerig, G., Bullitt, E., Joshi, S.: Multi-modal image set registration and atlas formation. *Medical Image Analysis* **10**(3) (2006) 440–451