

Groupwise Simultaneous Manifold Alignment for High-Resolution Dynamic MR Imaging of Respiratory Motion

Christian F. Baumgartner¹, Christoph Kolbitsch¹, Jamie R. McClelland²,
Daniel Rueckert³, and Andrew P. King¹

¹ Division of Imaging Sciences & Biomedical Engineering, King's College London, UK

² Centre for Medical Image Computing, University College London, UK

³ Biomedical Image Analysis Group, Department of Computing, Imperial College
London, UK

`christian.baumgartner@kcl.ac.uk`

Abstract. Respiratory motion is a complicating factor for many applications in medical imaging and there is significant interest in dynamic imaging that can be used to estimate such motion. Magnetic resonance imaging (MRI) is an attractive modality for motion estimation but current techniques cannot achieve good image contrast inside the lungs. Manifold learning is a powerful tool to discover the underlying structure of high-dimensional data. Aligning the manifolds of multiple datasets can be useful to establish relationships between different types of data. However, the current state-of-the-art in manifold alignment is not robust to the wide variations in manifold structure that may occur in clinical datasets. In this work we propose a novel, fully automatic technique for the simultaneous alignment of large numbers of manifolds with varying manifold structure. We apply the technique to reconstruct high-resolution and high-contrast dynamic 3D MRI images from multiple 2D datasets for the purpose of respiratory motion estimation. The proposed method is validated on synthetic data with known ground truth and real data. We demonstrate that our approach can be applied to reconstruct significantly more accurate and consistent dynamic images of the lungs compared to the current state-of-the-art in manifold alignment.

Keywords: Manifold learning, manifold alignment, MRI of the lungs, respiratory motion.

1 Introduction

Respiration is a complicating factor for many imaging techniques and image-guided interventions [10]. Motion caused by breathing is very complex with significant variations between respiratory cycles (inter-cycle variation), and also between inspiration and expiration (intra-cycle variation) [15,11].

Magnetic resonance imaging (MRI) offers an attractive means to image this complex motion because of its non-ionizing nature and high soft tissue contrast.

Many attempts have been made to estimate and model respiratory motion from MRI [9,6], but the state-of-the-art is limited by current MRI technology. Dynamic 3D scans suffer from poor contrast in the lungs, low image resolution, and relatively long acquisition times, which can lead to motion blurring. Dynamic 2D scans, on the other hand, can be acquired in a shorter time frame, have excellent in-plane resolution and high contrast in the lungs due to the in-flow of unpolarised blood, but lack the coverage of 3D scans. This is illustrated in Fig. 1. In this paper we propose a novel technique, based on manifold alignment, for combining the excellent contrast of dynamic 2D scans with the full thorax coverage of dynamic 3D scans.




Fig. 1. Comparison of (a) a coronal slice through a dynamic 3D MRI volume; (b) a coronal dynamic 2D MRI slice. The 2D slice has much improved contrast inside the lungs due to the in-flow of unpolarised blood.

Manifold learning is a powerful tool for non-linear dimensionality reduction of complex high-dimensional data and various manifold learning algorithms such as locally linear embedding (LLE) [12] or Laplacian eigenmaps (LEM) [1] have been proposed. In recent years manifold learning was shown to be useful in the analysis of motion in medical images. Recent applications include the region-wise separation of cardiac and respiratory motion [2], retrospective reconstruction of respiratory-gated lung computed tomography volumes [4] and extraction of respiratory gating navigators from MRI and ultrasound images [16].

Manifold alignment can be used to establish correspondences between multiple related datasets, which are not directly comparable in high-dimensional space, but have a similar low-dimensional manifold structure. There are two general approaches to manifold alignment: 1) The datasets are embedded in a common low-dimensional space in a single *simultaneous* embedding, either with prior knowledge of corresponding points [5,18], or without such knowledge, e.g. Joint Manifold Representation (JMR) [13]; 2) The datasets are embedded *separately*, and are then transformed to the same coordinate system in a subsequent alignment step either with known correspondences using a shape matching technique like Procrustes analysis [17], or without known correspondences using simple normalisation [4].

Although manifold alignment has proved effective at solving synthetic problems in computer vision, such as aligning the manifolds of video images of rotating 3D objects [5,13], there have been very few applications of these techniques to

real-world problems in medical imaging. To the best of our knowledge, the only such example is that by Bhatia et al. [2] who simultaneously aligned manifolds arising from different image patches in a hierarchical framework.

In this paper we show that the current state-of-the-art in manifold alignment is not robust to the wide variations in manifold structure which occur in real clinical images. We propose a novel method based on simultaneous group-wise embedding of datasets for the robust alignment of underlying manifolds without prior knowledge of correspondences. The method is demonstrated by retrospectively reconstructing dynamic high-resolution 3D MRI volumes from slice-by-slice 2D acquisitions. Such volumes can be used to estimate motion in the lungs, and have potential application in any scheme for which retrospective MRI-based respiratory motion estimates are required. One such application is the use of MRI data to motion-correct simultaneously acquired positron emission tomography (PET) data [14].

2 Theory

In the following we develop notation and theory for aligning the manifolds of large numbers of datasets by simultaneously embedding them in a groupwise framework.

2.1 Simultaneous Embedding of Two Datasets

In this section notation is established based on the example of simultaneously embedding two datasets. The derivations are conceptually similar to [5].

Given two high-dimensional datasets $\mathbf{X}_1, \mathbf{X}_2$, for each element $\mathbf{X}_i^{(1)}, \mathbf{X}_j^{(2)} \in \mathbb{R}^D$, we want to obtain the aligned embeddings $\mathbf{Y}_1, \mathbf{Y}_2$ with elements $\mathbf{Y}_i^{(1)}, \mathbf{Y}_j^{(2)} \in \mathbb{R}^d : d \ll D$. The total embedding error can be expressed as:

$$\phi_{tot}(\mathbf{Y}_1, \mathbf{Y}_2) = \phi_1(\mathbf{Y}_1) + \phi_2(\mathbf{Y}_2) + \mu \cdot \phi_{12}(\mathbf{Y}_1, \mathbf{Y}_2), \quad (1)$$

where ϕ_1, ϕ_2 , are the *intra-dataset* embedding errors and ϕ_{12} is the *inter-dataset* embedding error. The weighting parameter μ regulates the influence of the inter-dataset error term. Increasing μ forces the embeddings to be closer together, but setting it too high may alter the natural manifold structure of the data.

In contrast to most related manifold alignment works, which extend LEM [2,13], here we chose a LLE-based intra-dataset error term. As will be discussed in Section 4, a LEM type cost function was also investigated, but extending LLE proved advantageous.

LLE tries to preserve locally linear relations of the high-dimensional data in the low-dimensional embeddings. It is assumed that each high-dimensional point can be reasonably well reconstructed as a linear combination of its k nearest neighbours. The optimal reconstruction weights W_{ij} for each point i from its respective neighbours j can be calculated in closed form as described in [12].

The intra-dataset embedding errors, ϕ_1 , and ϕ_2 , which preserve the local relations of the high-dimensional data can then be expressed as

$$\phi_m(\mathbf{Y}_m) = \sum_i \left(\mathbf{Y}_i^{(m)} - \sum_{j \in \eta(i)} W_{ij}^{(m)} \mathbf{Y}_j^{(m)} \right)^2, \quad m \in \{1, 2\}, \quad (2)$$

where, for each dataset respectively, $\eta(i)$ is the neighbourhood of data point i .

For the inter-dataset error a different cost function is used. The embedding error of \mathbf{Y}_1 , and \mathbf{Y}_2 is defined as

$$\phi_{12}(\mathbf{Y}_1, \mathbf{Y}_2) = \sum_{i,j} \left(\mathbf{Y}_i^{(1)} - \mathbf{Y}_j^{(2)} \right)^2 U_{ij}, \quad (3)$$

where $U_{ij} = K(\mathbf{X}_i^{(1)}, \mathbf{X}_j^{(2)})$ is a (non-symmetric) similarity kernel. For high similarity values U_{ij} the error can only be minimised if $\mathbf{Y}_i^{(1)}$, and $\mathbf{Y}_j^{(2)}$ are close in the simultaneous embedding. The similarity kernel is application specific and may consist of a priori known labels [5,18], or image similarities [13].

With these choices of embedding error functions the minimisation of the whole cost function ϕ_{tot} can be rewritten in a single matrix expression,

$$\operatorname{argmin}_{\mathbf{Y}_1, \mathbf{Y}_2} Tr \left(\begin{bmatrix} \mathbf{Y}_1 \\ \mathbf{Y}_2 \end{bmatrix}^T \begin{bmatrix} \mathbf{M}_1 + \mu \mathbf{D}_1 & -\mu \mathbf{U} \\ -\mu \mathbf{U}^T & \mathbf{M}_2 + \mu \mathbf{D}_2 \end{bmatrix} \begin{bmatrix} \mathbf{Y}_1 \\ \mathbf{Y}_2 \end{bmatrix} \right), \quad (4)$$

where \mathbf{D}_1 and \mathbf{D}_2 are the row and column sums of \mathbf{U} as diagonal matrices, i.e. $D_{ii}^{(1)} = \sum_j U_{ij}$ and $D_{jj}^{(2)} = \sum_i U_{ij}$, and \mathbf{M}_1 and \mathbf{M}_2 , are the recentred reconstruction weight matrices $\mathbf{M}_m = (\mathbf{I} - \mathbf{W}_m)^T (\mathbf{I} - \mathbf{W}_m)$. This problem now has the same form as the standard LLE embedding [12], that is,

$$\operatorname{argmin}_{\mathbf{V}} Tr(\mathbf{V}^T \mathbf{L} \mathbf{V}), \quad (5)$$

where \mathbf{L} is the augmented matrix from Eq. (4) and \mathbf{V} are the augmented embeddings. Under the constraint that $\mathbf{V}^T \mathbf{V} = \mathbf{I}$, the simultaneous aligned embedding is given by the second smallest to the $(d + 1)$ -th smallest eigenvectors of \mathbf{L} .

2.2 Groupwise Simultaneous Embedding of Multiple Datasets

Eq. (4) can be easily extended to three or more manifolds by further augmenting \mathbf{V} and \mathbf{L} , as was described e.g. in [13]. However, this approach is not optimal for multiple reasons, in particular: 1) The problem becomes increasingly unstable when increasing the number of datasets \mathbf{X}_i ; 2) The manifold structure may not be similar enough across all datasets to justify a simultaneous embedding of all of them.

To overcome these limitations, we propose a novel scheme to embed the datasets simultaneously in overlapping groups of two, producing a much more

stable problem. For N high-dimensional inputs $\mathbf{X}_1, \dots, \mathbf{X}_N$, the datasets are embedded in $N - 1$ groups $\mathcal{G}^{(p)}$. Each group contains the simultaneous embeddings of two datasets $\mathbf{X}_{\mathbf{p}}$, and $\mathbf{X}_{\mathbf{p}+1}$, i.e. $\mathcal{G}_1^{(p)} = \mathbf{Y}_{\mathbf{p}}$, and $\mathcal{G}_2^{(p)} = \mathbf{Y}_{\mathbf{p}+1}$. The groups are interleaved such that $\mathcal{G}_2^{(p)} = \mathcal{G}_1^{(p+1)}$. Fig. 2 shows an example of a groupwise embedding in $d = 2$ dimensions and the relations between the groups.




Fig. 2. Schematic illustration of groupwise manifold alignment. The curved lines illustrate the manifold connections through the group overlap (solid), or through aligned embedding (dotted).

By embedding the datasets in overlapping groups the manifolds are all aligned. The two members of each group are aligned due to the simultaneous embedding, and the connections to the next group are deterministically known through the group overlap. For example, consider an arbitrary embedded point i in manifold p , i.e. $\mathbf{Y}_i^{(p)}$ (labelled with a square in Fig. 2), which is embedded in the first member manifold of group $\mathcal{G}^{(p)}$, i.e. $\mathcal{G}_1^{(p)} = \mathbf{Y}_{\mathbf{p}}$. Since within the group the manifolds are aligned, the closest neighbour on $\mathbf{Y}_{\mathbf{p}+1} = \mathcal{G}_2^{(p)}$ can be found directly (see dotted lines in Fig. 2). The two points ($\mathbf{Y}_{\mathbf{p}}$, and $\mathbf{Y}_{\mathbf{p}+1}$) can be looked up in the neighbouring groups since $\mathcal{G}_1^{(p)} = \mathcal{G}_2^{(p-1)}$ and $\mathcal{G}_2^{(p)} = \mathcal{G}_1^{(p+1)}$ (see solid lines in Fig. 2). In this manner the corresponding embeddings on all manifolds can be found iteratively.

2.3 Inter-dataset Similarity Kernel Sparsification

Typically techniques that do not assume any a priori known correspondences evaluate the inter-dataset similarity kernel $U_{ij} = K(\mathbf{X}_i^{(r)}, \mathbf{X}_j^{(a)})$ (see Eq. 3) on all possible bipartite connections between two datasets. However, not all connections are desirable, and typically the graph is sparsified. The standard approach for sparsification is to keep only the k nearest neighbours for each data point. This may lead to unevenly distributed connectivities, which will distort the embedding. The JMR technique [13] instead performs an orthogonalisation step of the matrix representation of the full graph.

In this paper, we propose a sparsification technique based on a global bipartite maximum edge similarity (MES) matching. We calculate the matching in which every data point in $\mathbf{X}_{\mathbf{r}}$ is connected to exactly one data point in $\mathbf{X}_{\mathbf{q}}$, and the sum

of similarities over the corresponding edges U_{ij} is maximised. Fig. 3 illustrates this process. The bipartite matching that maximises the similarity is highlighted in red. This is equivalent to a combinatorial optimisation problem and can be solved using the Hungarian method [8].




Fig. 3. Graph sparsification of similarity kernel. The left figure shows the fully connected graph, and the right figure shows the optimal one-to-one mapping.

The resulting graph can be written as a sparse matrix \mathbf{U} , which in every row and every column has exactly one non-zero entry, $0 < U_{ij} \leq 1$. The values represent the likelihood that the connection between $\mathbf{X}_i^{(r)}$ and $\mathbf{X}_j^{(q)}$ is correct.

3 Materials and Methods

In this section we show how our method can be used to reconstruct high-resolution 3D MRI volumes from sequentially acquired 2D MRI slices.

3.1 Application of Groupwise Manifold Alignment

We image the volume of interest by sequentially acquiring coronal slices at shifting slice positions. In order to sufficiently sample all respiratory states the volume is covered several times. The specific acquisition details will be given in Sec. 4.2.

The 2D MRI slices from different body positions are not directly comparable in high-dimensional image space. However, our hypothesis is that, because the underlying mechanics of respiratory motion are the same, the respiratory states at different slice positions lie on similar low-dimensional manifolds.

The low-dimensional representations are not intrinsically aligned. To overcome this we apply our novel groupwise manifold alignment algorithm. We define the 2D image data acquired at a slice position p as the high-dimensional input \mathbf{X}_p , and the i -th slice acquired at this position as $\mathbf{X}_i^{(p)}$. The groups are chosen such that neighbouring slices belong to the same group. Applying the algorithm we arrive at the aligned d -dimensional embeddings \mathbf{Y}_p for each slice position, where the coordinates of the i -th acquired slice are given by $\mathbf{Y}_i^{(p)}$. Note that by virtue of the alignment the respiratory correspondences of $\mathbf{Y}_i^{(p)}$ to data from any other slice position are known. Since every $\mathbf{Y}_i^{(p)}$ is related to its original data slice $\mathbf{X}_i^{(p)}$ a volume can be reconstructed by looking up and combining the corresponding slices from other slice positions.

By reconstructing volumes from the corresponding embedded coordinates in the original acquisition order of the slices, a sequence of high-resolution 3D volumes, one for each time point of the acquisition, can be obtained.

3.2 Choice of Kernel

2D slices which represent different anatomical structures are not directly comparable in image space. However, neighbouring slices are similar to some degree. Since by design the similarity kernel K is only evaluated on neighbouring slices the \mathcal{L}_2 -distance between images will serve as a reasonable measure. To obtain a similarity measure we define a Gaussian Kernel [1]

$$K(\mathbf{X}_i^{(r)}, \mathbf{X}_j^{(q)}) = \exp\left(-\frac{\tilde{\mathcal{L}}_2(\mathbf{X}_i^{(r)}, \mathbf{X}_j^{(q)})^2}{2\sigma^2}\right), \quad (6)$$

where $\tilde{\mathcal{L}}_2$ denotes the normalised \mathcal{L}_2 -distance, such that the maximum distance is 1 and the minimum distance is 0. The parameter σ governs the kernel shape.

The \mathcal{L}_2 -distance might be misleading even for neighbouring slices because a shift in space also causes a shift in the diaphragm, which might suggest a misleading change in respiratory position. By using the proposed MES sparsification (Section 2.3) this problem can be avoided. The sum over all the edges of the graph can only be maximised if those misleading connections are avoided.

4 Results

We evaluated our technique on two types of data: 2D MRI slices synthetically generated by extracting slices from warped high resolution 3D MRI volumes (7 volunteers), and real dynamic 2D MRI data (5 volunteers). The first made it possible to test our technique on realistic data with a known ground truth, whereas the second enabled us to test our technique on real data.

For both experiments we aligned the datasets originating from different slice positions using the simultaneous groupwise alignment method (SGA) proposed in this work, and also the JMR technique proposed in [13]. In addition, we performed the experiments on three variations of our technique, in order to evaluate each of the novelties of our proposed approach individually: SGA without groupwise embedding, i.e. with fully augmented matrices \mathbf{V} and \mathbf{L} (see Eq. 4) (SGA.FULL); SGA with a simple nearest neighbour sparsification instead of the proposed MES sparsification (SGA.NN); and SGA with a LEM-type cost function instead of the LLE cost function in Eq. 2 (SGA.LEM). For all examined techniques we empirically determined a good set of parameters before running the experiments. In particular, for all examined techniques we embedded the high-dimensional inputs into $d = 4$ dimensions, which proved to be sufficient to capture most intra- and inter-cycle breathing variations.

4.1 Experiments on Synthetic 2D Data

Fig. 4 illustrates the process of the synthetic data generation, from which we obtained realistic high-resolution coronal slices at different respiratory positions similar to those obtained by real slice-by-slice acquisitions from volunteers.

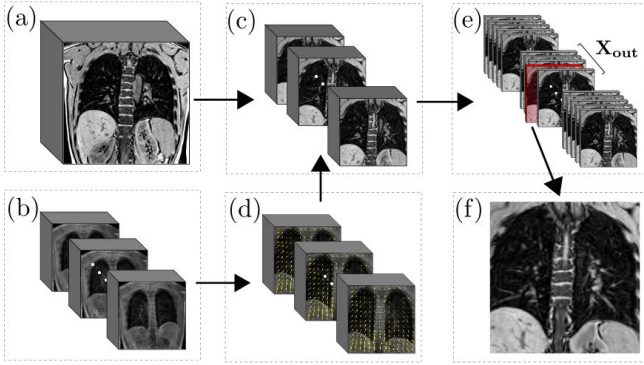


Fig. 4. Generation of synthetic data. A high-resolution breath-hold image such as that shown in (a) was warped to different respiratory states by generating motion fields (d) from 35 dynamic low-resolution volumes (b) [3]. A larger cross-section of a low-resolution volume is shown in Fig. 1a. Groups of slices from the synthetic high-resolution dynamics (c) obtained in this manner were combined to generate synthetic coronal slices with a realistic 2D dynamic slice thickness of 8 mm (e). A close-up of such a slice is shown in (f).

The high-resolution breath-hold volumes such as the one shown in Fig. 4a were acquired with a non-cardiac-triggered T1-weighted gradient echo sequence with an acquired image resolution of $1.5 \times 1.6 \times 1.5 \text{ mm}^3$. The low-resolution dynamic volumes such as the ones shown in Fig. 4b were acquired with a cardiac-triggered T1-weighted gradient echo sequence with an acquired image resolution of $1.5 \times 4.1 \times 5 \text{ mm}^3$, and typical acquisition time of 600 ms. Both scans were performed on a Philips Achieva 1.5T MRI scanner.

In order to estimate the performance of each of the manifold alignment approaches mentioned above we used a leave-out-one (LOO) cross validation framework. For each synthetic slice such as the one highlighted in red in Fig. 4e we left out the whole volume it belonged to (labelled by \mathbf{X}_{out} in the same Figure), but not the highlighted slice itself. For the remaining volumes plus the highlighted slice we calculated the aligned embedding using each of the approaches, and reconstructed a volume around the highlighted slice as described in Section 3.1. This resulted in an approximation $\hat{\mathbf{X}}_{\text{out}}$ of the left-out volume \mathbf{X}_{out} . The reconstruction error was estimated by calculating the \mathcal{L}_2 -distance, $\mathcal{L}_2(\hat{\mathbf{X}}_{\text{out}}, \mathbf{X}_{\text{out}})$, between the two volumes. This procedure was repeated for each of the slices originating from each of the warped volumes, and the mean and standard deviation of the \mathcal{L}_2 -distance were computed as overall measures of accuracy. As the magnitudes of these errors varied widely between volunteers, before combining them we normalised the maximum error to 1 for each subject.

Fig. 5a shows the reconstruction errors for each of the tested techniques over all 7 volunteers. We grouped the reconstruction errors by region of origin of the input slice, i.e. anterior, medial, or posterior. The distribution of the error was symmetric but not normal. Therefore, we used a 1-tailed Wilcoxon signed

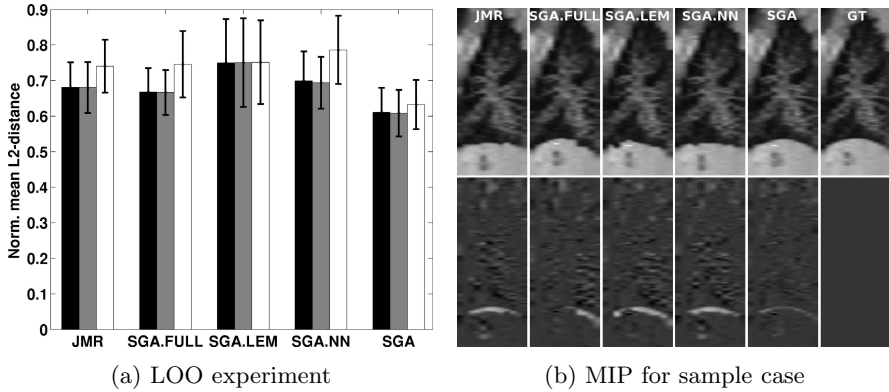


Fig. 5. (a) Mean \mathcal{L}_2 distances of synthetic data reconstruction for the tested approaches for anterior (black), medial (grey), posterior (white) input slice regions, and (b) MIP of reconstructed volumes in sagittal direction (top), and their difference to the same MIP of the left out volume (bottom).

rank test and found a statistically significant improvement for SGA versus all compared techniques ($p < 0.001$). In the top row of Fig. 5b maximum intensity projections (MIP) over the right lung of reconstructed volumes for a sample volunteer at an arbitrary respiratory position are shown. The rightmost column contains the MIP of the ground truth (GT). The bottom row contains the differences between the top row and the GT MIP.

4.2 Experiments on Real Data

We validated our method on real dynamic MRI data acquired using the slice-by-slice acquisition protocol described in Sec. 3.1. To maximise vessel contrast, and to minimise cardiac motion, only one slice was acquired per heartbeat at systole. The acquisition time for each slice was 160 ms. The slice position was shifted for each acquisition and to cover the whole thorax typically 15-19 slice positions were needed. Each position was acquired 50 times resulting in an overall acquisition time of 13-16 minutes. The scans were carried out on a Phillips Achieva 3T MRI scanner using a T1-weighted gradient echo sequence with an acquired in-plane image resolution of $1.4 \times 1.4 \text{ mm}^2$ and a slice thickness of 8 mm.

We used the manifold alignment techniques to embed the data from all slice positions and based on the respective embeddings reconstructed a volume for each slice. Thus, we obtained dynamic high-resolution volumes for each time point of the acquisition. We examined the same set of techniques as in the previous section, i.e. JMR, SGA.FULL, SGA.LEM, SGA.NN, and SGA.

Fig. 6a shows MIPs of reconstructed volumes through the right lung of a sample volunteer at a typical end-exhale position (top-row), and a typical end-inhale position (bottom-row).

In addition to reconstructing high-resolution dynamic volumes we evaluated the reconstruction consistency. For each acquired slice we reconstructed a

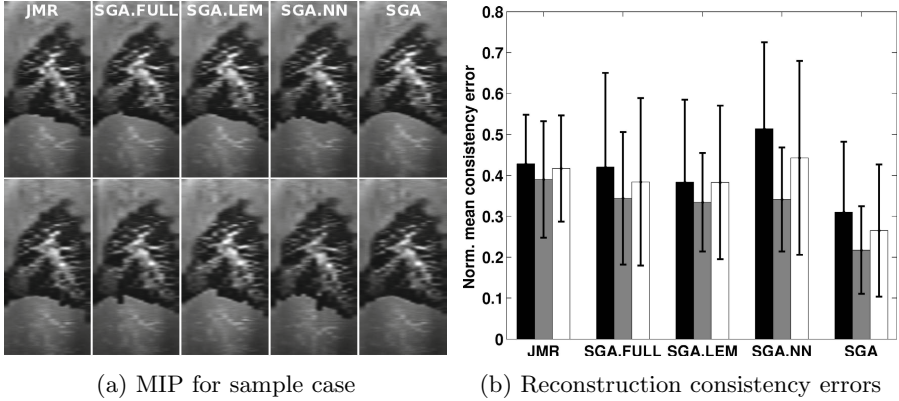


Fig. 6. (a) Sagittal MIP through reconstructed volumes from an input slice close to end-expiration (top row), end-inspiration (bottom row), and (b) consistency of reconstruction over different input slice positions for anterior (black), medial (grey), posterior (white) input slice regions for the five tested approaches.

volume \mathbf{X} , and for each slice in \mathbf{X} we reconstructed a new volume $\hat{\mathbf{X}}_s$. Ideally, $\hat{\mathbf{X}}_s$ should be equal to \mathbf{X} . In practice, however, a different input slice position will give a different reconstruction. To estimate this reconstruction consistency we calculated the \mathcal{L}_2 -distance, $\mathcal{L}_2(\hat{\mathbf{X}}_s, \mathbf{X})$, for each slice of each reconstructed volume. The mean consistency errors for all the tested approaches are shown in Fig. 6b. In order to combine the mean errors of different subjects we normalised the errors for each subject as before. Again, we grouped the errors by body region of the input slice. Using a 1-tailed Wilcoxon signed rank test SGA’s improvements over all other techniques were found to be statistically significant ($p < 0.001$).

5 Discussion and Conclusions

We have presented a novel, fully automatic approach for the simultaneous alignment of large numbers of related datasets based on an extension of LLE to groupwise simultaneous embeddings and a global, MES based optimisation of the inter-dataset kernel.

Applying the technique to the reconstruction of high-resolution dynamic MRI volumes from a slice-by-slice acquisition protocol gave better and more consistent results than the current state-of-the-art for manifold alignment with no known correspondences [13]. On synthetic data with a known GT our method yielded the most accurate reconstructions from all of the tested approaches. On real data our technique gave the most plausible and most consistent reconstructions. For all experiments the proposed method gave good reconstructions from input slices from anterior, medial and posterior body regions, which is important in forming a coherent dynamic sequence, as input slice positions will vary over time.

The experiments showed that each of the elements of the proposed technique, i.e. the MES sparsification, groupwise embedding, and extension of LLE instead of LEM were necessary for robust reconstructions. Generally, we noticed that LLE was better able to capture the subtle intra-, and inter-cycle breathing motion variations that occur in some subjects, than the commonly used LEM technique.

In this paper we investigated simultaneous manifold alignment approaches. Separate alignment without prior knowledge can be performed by simple normalisation [4]. However, this is only reliable for 1-dimensional embeddings, which may not be sufficient to capture all variations in the data.

A different approach to retrospective reconstruction of high-resolution volumes from a slice-by-slice acquisition protocol was proposed by von Siebenthal et al. [15]. The authors used an interleaved slice acquisition protocol where a navigator slice at a constant slice position was acquired before and after each data slice. Manually selected image features derived from the navigator slices were then matched retrospectively to reconstruct a volume. However, in this method only half of the data is actually used for reconstruction, which doubles the already long acquisition times. In contrast our method is fully automatic, doesn't require additional navigator slices and makes use of all the acquired data.

Slice-by-slice 2D acquisitions allow excellent image contrast and make it possible to image vessel structures inside the lungs which cannot be visualised using a dynamic 3D MRI acquisition protocol. These additional structures could be used to extract very detailed and reliable motion fields of the whole thorax, which would have potential application for motion correction of simultaneously acquired PET-MRI data. This could be achieved with minimal scanning overheads, since one slice per cardiac cycle is sufficient to retrospectively obtain high contrast volumes for the entire duration of a PET imaging session. Currently, the volumes only have a high in-plane resolution. Reducing the slice thickness is a natural extension of our work, but will increase acquisition times. PET imaging of the thorax typically takes 15-30 minutes which is sufficient time to acquire enough data for reconstructions which also have a high through-plane resolution.

The proposed technique also has potential application in MRI-guided treatments such as MRI-guided high-intensity focused ultrasound (HIFU) [7]. The aligned embeddings contain information about different respiratory motion states and can thus be seen as a motion model [10]. A volume can be generated from previously unseen data by acquiring a new slice at a convenient slice position and embedding it into its appropriate manifold. In MR-guided HIFU this could be applied for online updating of guidance information.

Acknowledgements. This work was funded by EPSRC programme grant EP/H046410/1. This research was supported by the National Institute for Health Research (NIHR) Biomedical Research Centre at Guy's and St Thomas' NHS Foundation Trust and King's College London. The views expressed are those of the authors and not necessarily those of the NHS, the NIHR or the Department of Health. This work was partly funded by the EU FP7 SUBLIMA project.

References

1. Belkin, M., Niyogi, P.: Laplacian eigenmaps and spectral techniques for embedding and clustering. *Adv. Neur. In.* 14, 585–591 (2001)
2. Bhatia, K.K., Rao, A., Price, A.N., Wolz, R., Hajnal, J., Rueckert, D.: Hierarchical manifold learning. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (eds.) *MICCAI 2012, Part I. LNCS*, vol. 7510, pp. 512–519. Springer, Heidelberg (2012)
3. Buerger, C., Schaeffter, T., King, A.P.: Hierarchical adaptive local affine registration for fast and robust respiratory motion estimation. *Med. Image Anal.* 15(4), 551–564 (2011)
4. Georg, M., Souvenir, R., Hope, A., Pless, R.: Manifold learning for 4D CT reconstruction of the lung. In: *Proc. IEEE CVPRW*, pp. 1–8 (2008)
5. Ham, J., Lee, D., Saul, L.: Semisupervised alignment of manifolds. In: *AI and Statistics*, vol. 10, pp. 120–127 (2005)
6. King, A.P., Buerger, C., Tsoumpas, C., Marsden, P.K., Schaeffter, T.: Thoracic respiratory motion estimation from MRI using a statistical model and a 2-D image navigator. *Med. Im. Anal.* 16(1), 252–264 (2012)
7. Köhler, M.O., Denis de Senneville, B., Quesson, B., Moonen, C.T.W., Ries, M.: Spectrally selective pencil-beam navigator for motion compensation of MR-guided high-intensity focused ultrasound therapy of abdominal organs. *Magn. Reson. Med.* 66(1), 102–111 (2011)
8. Kuhn, H.W.: The Hungarian method for the assignment problem. *Nav. Res. Logist. Q* 2(1-2), 83–97 (1955)
9. Manke, D., Nehrke, K., Börnert, P.: Novel prospective respiratory motion correction approach for free-breathing coronary MR angiography using a patient-adapted affine motion model. *Magn. Reson. Med.* 50(1), 122–131 (2003)
10. McClelland, J.R., Hawkes, D.J., Schaeffter, T., King, A.P.: Respiratory motion models: A review. *Med. Image Anal.* 17(1), 19–42 (2013)
11. McClelland, J.R., Hughes, S., Modat, M., Qureshi, A., Ahmad, S., Landau, D.B., Ourselin, S., Hawkes, D.J.: Inter-fraction variations in respiratory motion models. *Phys. Med. Biol.* 56(1), 251–272 (2011)
12. Roweis, S.T., Saul, L.K.: Nonlinear dimensionality reduction by locally linear embedding. *Science* 290(5500), 2323–2326 (2000)
13. Torki, M., Elgammal, A., Lee, C.S.: Learning a joint manifold representation from multiple data sets. In: *Proc. IEEE ICPR*, pp. 1068–1071 (2010)
14. Tsoumpas, C., Mackewn, J.E., Halsted, P., King, A.P., Buerger, C., Totman, J.J., Schaeffter, T., Marsden, P.K.: Simultaneous PET–MR acquisition and MR-derived motion fields for correction of non-rigid motion in PET. *Ann. Nucl. Med.* 24(10), 745–750 (2010)
15. von Siebenthal, M., Székely, G., Gamper, U., Boesiger, P., Lomax, A., Cattin, P.: 4D MR imaging of respiratory organ motion and its variability. *Phys. Med. Biol.* 52(6), 1547–1564 (2007)
16. Wachinger, C., Yigitsoy, M., Rijkhorst, E.J., Navab, N.: Manifold learning for image-based breathing gating in ultrasound and MRI. *Med. Im. Anal.* 16(4), 806–818 (2011)
17. Wang, C., Mahadevan, S.: Manifold alignment using Procrustes analysis. In: *Proc. ICML* (2008)
18. Zhai, D., Li, B., Chang, H., Shan, S., Chen, X., Gao, W.: Manifold alignment via corresponding projections. In: *Proc. BMVC*, pp. 3–11 (2010)