Dynamic Volume Reconstruction from Multi-slice Abdominal MRI Using Manifold Alignment

Xin Chen(✉), Muhammad Usman, Daniel R. Balfour, Paul K. Marsden, Andrew J. Reader, Claudia Prieto, and Andrew P. King

Division of Imaging Sciences and Biomedical Engineering,
King's College London, London, UK
xin.chen@kcl.ac.uk

Abstract. We present a novel framework for retrospective dynamic 3D volume reconstruction from a multi-slice MRI acquisition using manifold alignment. K-space data are continuously acquired under free breathing using a radial golden-angle trajectory in a slice-by-slice manner. Non-overlapping consecutive profiles that were acquired within a short time window are grouped together. All grouped profiles from all slices are then simultaneously embedded using manifold alignment into a common manifold space (MS), in which profiles that were acquired at similar respiratory states are close together. Subsequently, a 3D volume can be reconstructed at each of the grouped profile MS positions by combining profiles that are close in the MS. This enables the original multi-slice dataset to be used to reconstruct a dynamic 3D sequence based on the respiratory state correspondences established in the MS. Our method was evaluated on both synthetic and *in vivo* datasets. For the synthetic datasets, the reconstructed dynamic sequence achieved a normalised cross correlation of 0.98 and peak signal to noise ratio of 26.64 dB compared with the ground truth. For the *in vivo* datasets, based on sharpness measurements and visual comparison, our method performed better than reconstruction using an adapted central k-space gating method.

Keywords: Manifold alignment · Respiratory motion estimation · MRI · Dynamic 3D volume reconstruction

1 Introduction

Dynamic magnetic resonance imaging (MRI) involves imaging a region of interest ideally with high temporal resolution, and is useful in many applications in which knowledge of motion (e.g. respiratory motion) is of interest. However, the acquisition speed of MRI is not sufficiently fast to permit enough data to be acquired quickly enough to reconstruct fully sampled images with high temporal resolution, especially for 3D imaging. This problem can be tackled by using under-sampled reconstruction schemes such as compressed sensing (CS) [1], but such techniques typically involve a complex optimisation which can be time-consuming for large amounts of dynamic MRI data. Another approach is to use a gating technique, which involves the retrospective combination of imaging data that were acquired at different times but similar motion states.

S. Ourselin et al. (Eds.): MICCAI 2016, Part III, LNCS 9902, pp. 493–501, 2016.
DOI: 10.1007/978-3-319-46726-9_57

The term 'self-gating' refers to performing this gating using the acquired data itself. A common self-gating approach for radial acquisitions is to use the magnitude of the centre of k-space which is acquired in every radial profile as the self-gating signal [2]. A key weakness of this approach is the simplicity of the gating signal – gating using a scalar value necessarily causes complex intra- and inter-cycle variations in respiratory motion to be averaged out.

Alternative self-gating methods that are based on manifold learning (ML) have been reported [3–7]. The intuition behind such methods is that respiratory motion is pseudo-repetitive in nature, and can therefore be represented by a small number of motion variables. Usman et al. [7] applied Laplacian Eigenmaps to estimate respiratory motion from the central intersection region of a number of consecutive profiles using a radial golden-angle acquisition (RGA) [8]. Similarly, Bhatia et al. [5] used Laplacian Eigenmaps to estimate cardiac motion from repetitively sampled central k-space lines using a Cartesian acquisition. However, both methods estimated only 1D signals (i.e. scalar values over time) for self-gating and also only worked for 2D images. It would be nontrivial to extend these approaches to reconstruction of 3D volume. In [3, 4], Baumgartner et al. proposed a method for MRI self-gating using manifold alignment (MA) of 2D reconstructed image slices acquired at different anatomical positions. However, [3, 4] are image-based self-gating methods that assumed there was no motion within each of the fully sampled 2D image acquisitions (~ 300 ms per image), which may not be true for fast motion. The use of images for self-gating also limits the achievable temporal resolution to the time taken to acquire enough data to reconstruct the entire image.

In this paper, we propose to retrospectively reconstruct 3D dynamic volume sequences from k-space data acquired using a multiple 2D slice (M2D) acquisition. The proposed framework provides a number of advantages over other methods in the literature: (1) In contrast to the central k-space gating (CKG) method, instead of using a 1D gating signal, the proposed method allows the use of a larger k-space region to estimate a multi-dimensional gating signal to better capture intra-cycle and inter-cycle motion variations. (2) Different from [3, 4], we perform MA directly on k-space profiles rather than reconstructed images. This enables our technique to achieve higher temporal resolution. (3) Compared to the 2D ML methods in [5, 7], our method works for reconstruction of 3D dynamic volumes. To the authors' knowledge, no work has been reported in the literature for reconstructing high temporal resolution dynamic 3D MRI sequences based on k-space data from a M2D acquisition.

2 Methodology

As illustrated in Fig. 1, the proposed framework consists of k-space data acquisition, manifold alignment and image reconstruction, which are described in the following subsections.

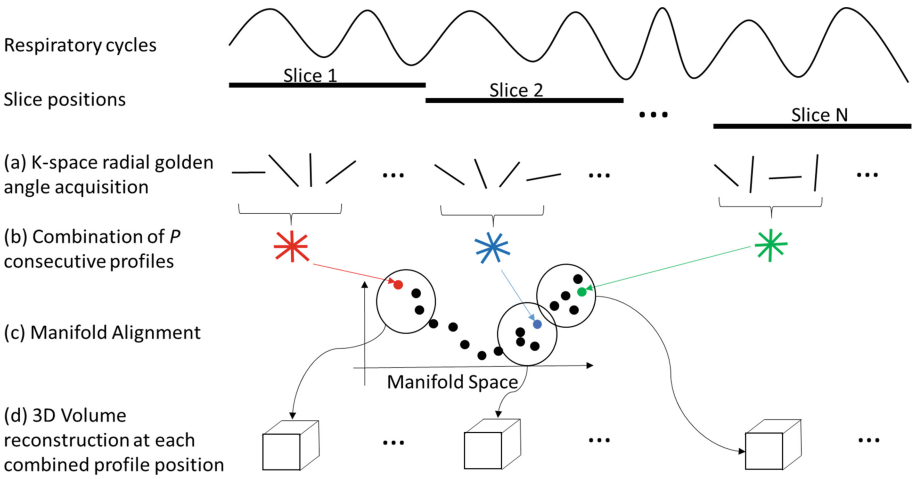

Fig. 1. Overview of the proposed framework.

2.1 K-Space Data Acquisition and Pre-processing

Data acquisition is performed under free breathing using a RGA trajectory (Fig. 1(a)), in which the angle between each two consecutive profiles is 111.25°. This enables a uniform coverage of k-space with high temporal incoherence for any arbitrary number of consecutive profiles.

M profiles are acquired at each slice position, followed by *M* profiles at the next slice position, and so on (see Fig. 1(a)). We denote the k-space data for each slice by $X_n = [x_{n,1}, \ldots, x_{n,M}]$ ($n \in \{1, 2, \ldots, N\}$ is the slice number), where the columns $x_{n,m}$ are k-space profiles. As reported in [7], using a larger central k-space region rather than the single sample of the k-space centre can produce more reliable estimation of respiratory motion. Therefore, since it takes only about 3 ms to acquire a single profile, we combine *P* consecutive profiles (Fig. 1(b)), and assume that no significant motion within this short time period occurred. Combination of profiles in this way enables a larger k-space central region to be used for respiratory self-gating. Since the radial profiles are sampled at non-uniform k-space positions, the convolutional gridding method [9] is applied to re-sample the non-uniform data to a Cartesian grid. Finally, the central $S \times S$ k-space region is used for the subsequent MA (see Sect. 2.2). After the profile combination and central region extraction, the original k-space data X_n are re-organised as $A_n = [a_{n,1}, \ldots, a_{n,L}]$. L is the quotient of *M/P*, and the remaining profiles after the $L \times P^{th}$ profile are not used. The columns $a_{n,l}$ are the vector representations of the $S \times S$ central k-space region. To reduce sensitivity to noise, we applied a 1D Gaussian smoothness filter ($\sigma = 15$ ms) along the temporal direction for each element in $a_{n,l}$.

2.2 Manifold Alignment

The k-space dataset $A_1, \ldots, A_N$, now consists of N groups, each of which represents a different slice, where each group contains L vectors with S^2 elements each. The MA method is used to simultaneously reduce the dimensionality ($\mathbb{R}^{S\times S}$) of each group and align all groups of profiles into a common MS (with dimensionality $d < S \times S$), in which profiles acquired at similar respiratory states are close together (Fig. 1(c)). The dimensionality reduction and alignment of the groups is performed using the MA scheme proposed in [3]. Therefore, we briefly review this technique here. The MA scheme estimates the low dimensional embeddings $Y_n = \left[y_1^{(n)}, \ldots, y_L^{(n)}\right]$ of the input high dimensional data by minimising a joint cost function:

$$\emptyset_{total}(Y_1 \ldots Y_N) = \sum_{n=1}^{N} \varphi_n(Y_n) + \frac{\mu}{2} \sum_{\substack{n=1, m=1, \\ m \neq n}}^{N} \sum_{l,j}^{L} U_{lj}^{(nm)} ||y_l^{(n)} - y_j^{(m)}||^2 \tag{1}$$

The first term φ_n is the locally linear embedding (LLE) cost function [10], which represents the intra-group embedding errors. LLE forms the low-dimensional MS by preserving locally linear relations derived from the original high-dimensional data A_n. These relations at the l^{th} high dimensional data vector are represented by a weighted ($W_{lj}^{(n)}$) linear combination of its K^{LLE} nearest neighbours with indices j ($j \in \eta(l)$), i.e.

$$\varphi_n(Y_n) = \sum_l ||y_l - \sum_{j\in\eta(l)} W_{lj}^{(n)} y_j||^2 = Tr(YMY^T) \tag{2}$$

Tr(•) is the trace operator, and the centred reconstruction weight matrix M is calculated as $M = (I - W)^T(I - W)$. Note that $W_{lj}^{(n)}$ are computed from the high dimensional data A_n.

The second term in Eq. (1) represents the inter-group cost function. $y_l^{(n)}$ represents the MS coordinates of vector $a_{n,l}$. μ is a weighting parameter that balances the intra-group and inter-group terms. $U^{(nm)}$ is an inter-group similarity kernel derived from intra-group comparisons of the high-dimensional data $A_1, \ldots, A_N$. This is achieved by deriving a feature descriptor for each vector $a_{n,l}$ using the concept of random walks from graph theory. This feature descriptor allows the relationships between vectors from the same group to be preserved and uniquely represented, which permits the inter-group similarity to be robustly measured without inter-group comparisons of the original high-dimensional data. For us, the groups represent k-space data acquired from different slices, which are not directly comparable since they encode the frequency properties of different anatomical regions. For more details of the MA process please refer to [3].

2.3 3D Dynamic Sequence Reconstruction

As illustrated in Fig. 1(c), after the MA process, all of the combined profiles (CPs) from different slice positions are embedded into a common MS. (The CPs refer to the

combined consecutive profiles as described in Sect. 2.1.) The CPs that were acquired at similar respiratory states should be close together in the MS, even though they belong to different slice positions. Image reconstruction is subsequently performed by grouping the G CPs for each slice that have the closest Euclidean distances (in MS) to the current CPs (Fig. 1(d)). G is determined by the total number of profiles (T) which is required for reconstruction of each slice, divided by P (the number of combined profiles). Since a RGA acquisition is used, this ensures the selected profiles are approximately evenly distributed in k-space. Using a multi-slice reconstruction, a 3D volume can be reconstructed at each of the L respiratory states for each of the N slice positions. This enables the original multi-slice dataset to be used to reconstruct a dynamic 3D sequence. For the reconstruction, the non-uniform Fast Fourier Transform [9] is applied to reconstruct the final image from the selected radial k-space profiles.

3 Experiments and Results

The proposed method was evaluated on both synthetic and *in vivo* datasets. The synthetic datasets were used to establish a ground truth for quantitative evaluation. We also demonstrated the practical feasibility of our technique using 3D *in vivo* datasets, and the results were compared with an adapted CKG method.

3.1 Materials

Synthetic Dataset Generation: To mimic the M2D acquisition, we generated high temporal resolution 3D synthetic datasets, based on image registration of a respiratory gated high resolution (RGHR) 3D MRI volume to a dynamic 3D low-resolution (DLR) MRI sequence. The RGHR volume had 120 sagittal slices and was acquired at the end-exhale position, with TR = 4.4 ms, TE = 2.2 ms, flip angle = 90°, and acquired voxel size $2.19 \times 2.19 \times 2.74$ mm^3 with matrix size 160×120. The DLR sequences comprised 35 dynamics acquired under free-breathing using cardiac gating at late diastole. 20 slices were acquired for each volume using 3D TFEPI with TR = 10 ms, TE = 4.9 ms, flip angle 20°, acquired voxel size $2.7 \times 3.6 \times 8.0$ mm^3, acquired matrix size 128×77, TFE factor 26, EPI factor 13, TFE acquisition time 267.9 ms. The region-of-interest focused on the right lung and liver region. In order to generate a large amount of realistic synthetic high spatial/temporal resolution volumes, the following steps were followed.

(a) The RGHR volume was deformed to align with the corresponding end-exhale DLR volume using B-spline deformable image registration [11].
(b) The end-exhale DLR volume was registered with all other DLR volumes.
(c) The DLR volumes were grouped into three different respiratory groups (inhale, exhale and mid-inhale), according to their head-foot diaphragm positions in the central slice. We then randomly selected one volume from each of the inhale and exhale groups, and two volumes from the mid-inhale group. Extra volumes were interpolated between these four volumes using B-spline registrations over a

complete breathing cycle. In our experiments, 60 such breathing cycles were generated. Each breathing cycle lasted approximately 5 s and the interpolated volumes had a temporal resolution of ~ 3.2 ms.

(d) The RGHR volume was warped to each of the DLR positions based on the corresponding registration results, resulting in about 90000 high spatial/temporal resolution dynamics containing realistic intra-cycle and inter-cycle variations.

(e) One k-space profile was extracted from each of the synthetic volumes at the corresponding slice position. Approximately 9000 profiles were extracted for each slice before moving to the next slice position. We generated 6 (1 for parameter optimisation and 5 for validation) such highly realistic synthetic sequences to validate our method against known gold-standard volumes.

***In vivo* Dataset Acquisition:** A M2D acquisition with RGA trajectory was employed for data acquisition in the liver and lung region of five healthy volunteers. For each volunteer, $M = 9000$ profiles were acquired using the RGA trajectory in sagittal view at each slice position before moving to the next slice. Data with 8 slices were acquired on a Philips 1.5T scanner using a 28 channel-coil. The settings were TR = 3.2 ms, TE = 1.58 ms, flip angle = 70°, and acquired voxel size $2.0 \times 2.0 \times 8.0$ mm^3 with acquired matrix size of 160×160 for 8 slices. For each volunteer, a total of 72000 radial profiles were acquired under free-breathing in approximately 4 min.

3.2 Results

All free parameters were tuned using a parameter sweep based on a single synthetic dataset (which was not used for validation) and consistently applied to the remaining synthetic and *in vivo* datasets. The parameters were set as $S = 9$, $P = 20$, $\mu = 10^{-5}$, $d = 3$ and $T = 200$ for the following experiments.

Synthetic Dataset: The manifold embedding for one of the synthetic datasets is shown in Fig. 2 (left). The colours represent the normalised head-foot diaphragm position of the ground truth. Similar coloured points grouping together indicates a good manifold alignment. The larger distances between the red points in the MS are due to larger respiratory variations in the end inhale positions. The reconstructed image quality was quantitatively measured for all 5 datasets based on normalised cross correlation (NCC) and peak signal-to-noise ratio (PSNR), against the ground truth. Example ground truth and

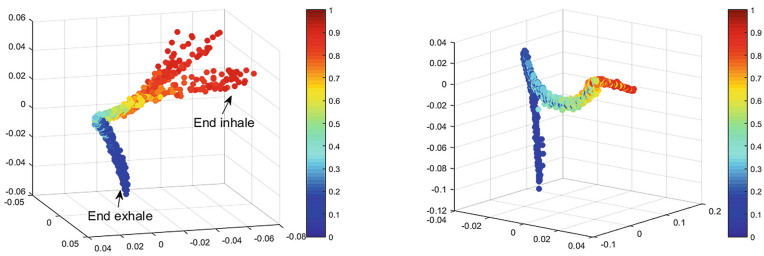

Fig. 2. Manifold embeddings for a synthetic dataset (left) and an *in vivo* dataset (right).

corresponding reconstructed images of 3 different sagittal slices, and coronal views at 4 different respiratory positions are shown in Fig. 3. The mean and standard deviation of NCC and PSNR values for all corresponding volumes of 5 synthetic datasets were 0.9801 ± 0.0057 and 26.64 ± 0.88 dB respectively.

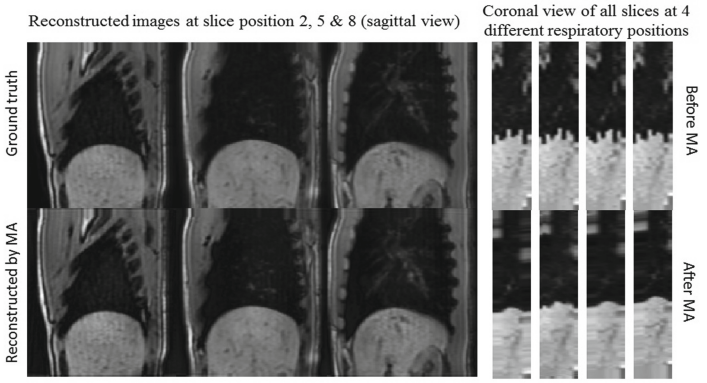

Fig. 3. Reconstructed volumes using MA and the corresponding ground truth of a synthetic dataset.

***In vivo* Dataset:** Due to the lack of ground truth and comparative techniques for the *in vivo* datasets, we adapted the CKG method (ACKG) for comparison. The central k-space magnitudes of each slice were normalised to the range of 0 to 1. Correspondences

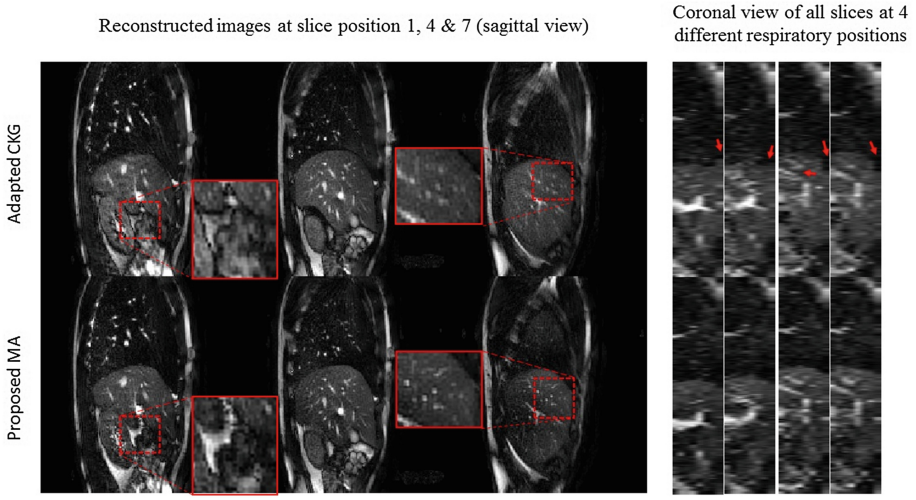

Fig. 4. Examples of reconstructed images using adapted CKG (top row) and the proposed MA method (bottom row) at different slice positions and different respiratory positions. (Color figure online)

across slices were then established by searching the same value of the normalised k-space magnitudes. To reconstruct a gated 3D volume for a specific profile, a gating window (0.05) was set around the current profile magnitude for each slice. Any profiles (T^{all}) with a magnitude within the gating window represent similar respiratory positions, and only the T temporally closest profiles were selected for image reconstruction. If $T^{all} < T$, the gating window was iteratively increased by 0.005 until $T^{all} \geq T$. Subsequently, the reconstructed slices were stacked into a volume.

The manifold embeddings for one of the *in vivo* datasets is shown in Fig. 2 (right). The colours represent the normalised k-space magnitude. A sharpness measurement was used as a quantitative measure of image quality. Sharpness was measured by the average of the image gradient magnitude in the liver and liver-lung boundary. Over all five *in vivo* datasets, the mean and standard deviation of the sharpness measure for our method was 0.0587 ± 0.0069. This was statistically significantly ($p < 0.01$) higher than the ACKG method (0.0559 ± 0.0071), based on two-sided Wilcoxon signed rank test of the paired datasets. Figure 4 shows examples of the reconstructed images using the ACKG method and our method, at different sagittal slice positions, and coronal views at 4 respiratory positions. Sharper organ details (red boxes) and more accurate slice alignment (red arrows) can be observed using our method.

4 Conclusion and Discussion

We have presented a novel framework that is able to reconstruct high temporal resolution dynamic 3D MRI sequences using manifold alignment of k-space data from M2D acquisitions. Based on synthetic datasets, our method achieved high NCC and PSNR measurements compared with the ground truth. Our method also showed improved image sharpness and visual image quality over images reconstructed using an adapted CKG method on *in vivo* data. We believe the improved performance of our method is due to the fact that a multi-dimensional gating signal is extracted from a larger k-space region. However, similar to the CKG method, the MA method cannot achieve a good reconstruction if an insufficient number of profiles is available at similar respiratory positions. It is possible that such a limitation may be addressed by the subsequent use of CS techniques. High temporal resolution 3D motion information will be useful in any application for which retrospective motion estimations are required, such as motion correction of positron emission tomography (PET) data in a simultaneous PET-MR scenario, or motion-correction of the MR data.

Acknowledgements. This work was funded by the Engineering and Physical Sciences Research Council (Grant EP/M009319/1).

References

1. Lustig, M., Donoho, D., Pauly, J.: Sparse MRI: the application of compressed sensing for rapid MR imaging. Mag. Res. Med. **58**(6), 1182–1195 (2007)
2. Larson, A., White, R., Laub, G., McVeigh, E., Li, D., Simonetti, O.: Self-gated cardiac cine MRI. Mag. Res. Img. **51**(1), 93–102 (2004)
3. Baumgartner, C.F., Gomez, A., Koch, L.M., Housden, J.R., Kolbitsch, C., McClelland, J.R., Rueckert, D., King, A.P.: Self-aligning manifolds for matching disparate medical image datasets. In: Ourselin, S., Alexander, D.C., Westin, C.-F., Cardoso, M. (eds.) IPMI 2015. LNCS, vol. 9123, pp. 363–374. Springer, Heidelberg (2015)
4. Baumgartner, C., Kolbitsch, C., Balfour, D., Marsden, P., McClelland, J., Rueckert, D., King, A.: High-resolution dynamic MR imaging of the thorax for respiratory motion correction of PET using groupwise manifold alignment. Med. Img. Anal. **18**, 939–952 (2014)
5. Bhatia, K.K., Caballero, J., Price, A.N., Sun, Y., Hajnal, J.V., Rueckert, D.: Fast reconstruction of accelerated dynamic MRI using manifold kernel regression. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015, Part III. LNCS, vol. 9351, pp. 510–518. Springer, Heidelberg (2015). doi:10.1007/978-3-319-24574-4_61
6. Usman, M., Atkinson, D., Kolbitsch, C., Schaeffter, T., Prieto, C.: Manifold learning based ECG-Free free-breathing cardiac CINE MRI. Mag. Res. Img. **41**, 1521–1527 (2015)
7. Usman, M., Vaillant, G., Atkinson, D., Schaeffter, T., Prieto, C.: Compressive manifold learning: estimating one-dimensional respiratory motion directly from undersampled k-space data. Mag. Res. Med. **72**, 1130–1140 (2014)
8. Winkelmann, S., Schaeffter, T., Koehler, T., Eggers, H., Doessel, O.: An optimal radial profile order based on the golden ratio for time-resolved MRI. IEEE Trans. Med. Img. **26**(1), 68–76 (2007)
9. Greengard, L., Lee, J.-Y.: Accelerating the nonuniform fast fourier transform. SIAM Rev. **46**(3), 443–454 (2006)
10. Roweis, S., Saul, L.: Nonlinear dimensionality reduction by locally linear embedding. Science **290**, 2323–2326 (2000)
11. Rueckert, D., Sonoda, L., Hayes, C., Hill, D., Leach, M., Hawkes, D.: Nonrigid registration using free-form deformations: application to breast MR images. IEEE Trans. Med. Img. **18**(8), 712–721 (1999)