



King's Research Portal

DOI:

[10.1109/ISBI.2018.8363583](https://doi.org/10.1109/ISBI.2018.8363583)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Clough, J. R., Balfour, D. R., Marsden, P. K., Prieto, C., Reader, A. J., & King, A. P. (2018). MRI slice stacking using manifold alignment and wave kernel signatures. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)* (Vol. 2018-April, pp. 319-323) Advance online publication. <https://doi.org/10.1109/ISBI.2018.8363583>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

MRI SLICE STACKING USING MANIFOLD ALIGNMENT AND WAVE KERNEL SIGNATURES

James R. Clough, Daniel R. Balfour, Paul K. Marsden, Claudia Prieto, Andrew J. Reader, Andrew P. King

School of Biomedical Engineering & Imaging Sciences
King's College London

ABSTRACT

MRI slice stacking involves retrospective combination of 2D MRI images to form pseudo 3D volumes. It is useful because physical constraints limit the temporal/spatial resolutions with which dynamic 3D MRI volumes can be acquired and so stacking fast high-resolution 2D images can yield pseudo 3D volumes with high in-plane spatial and temporal resolution. However, it is important that the stacked 2D images were acquired at consistent motion states. Assessing motion state consistency between slices representing different anatomy is challenging as the image contents are not easily comparable. Manifold alignment (MA) is a technique which provides a solution to this problem by embedding the 2D images for all slices into one globally consistent low-dimensional space. One successful approach to MA involves forming graphs from each slice dataset and using graph descriptors to find correspondences between datasets. Here we propose a new graph descriptor for the slice stacking problem, inspired by work in the computer vision literature, and evaluate it with two experiments. First, using a highly realistic synthetic MRI dataset in which reconstructed volumes can be compared to a ground truth, we find our method significantly outperforms the state of the art. Second, we use *in vivo* MRI data and show that the volumes reconstructed by our method have a higher degree of self-consistency.

Index Terms— Magnetic resonance imaging, Manifold alignment, Respiratory motion, Graphs

1. INTRODUCTION

High in-plane temporal and spatial resolution volumes can be retrospectively formed from dynamic 2D magnetic resonance imaging (MRI) data using slice stacking techniques. These volumes can be used for motion correction of PET [1] and the study of respiratory motion [2]. The slice stacking aims to group into volumes slices that were acquired at similar motion states. Techniques have been proposed that perform the grouping using an extra gating signal, such as a 1-D [3] or 2-D [2] MRI navigator. However, these approaches require extra scanning time. Alternatively, the grouping can be based on the acquired MRI slices themselves [4]. Our focus here is on the latter case. The underlying challenge of this self-gating approach is to find motion state correspondences between images which represent different anatomy.

The task of relating images from different datasets is common in medical imaging. Examples include images from different views

[5, 6] and different acquisition protocols [7]. Manifold alignment (MA) approaches have been proposed to address this problem [1].

MA is an extension of manifold learning (ML) in which multiple high-dimensional datasets are embedded into the same low dimensional space. Some MA techniques compute this embedding based on direct inter-dataset comparisons [1, 8]. However, sometimes the disparate nature of the different datasets makes this difficult. Alternative approaches have been proposed based on graph theory. Since images from the same dataset can be directly compared, a graph in which images are represented by nodes is formed for each dataset based on such intra-dataset comparisons. Inter-dataset comparisons are then made by comparing graph-based descriptors rather than the images themselves [6].

Our contribution in this paper is to apply and develop recently proposed graph based descriptors to MA to achieve state-of-the-art performance in MRI slice stacking. In particular, we use an adapted version of the Wave Kernel Signature (WKS) [9] which has previously been applied only to problems in computer vision. Furthermore, whereas slice-stacking methods have previously only used image intensities as the high-dimensional data, we also investigate the use of registration-derived motion field data. We find that in some cases, especially when the size of the high-dimensional dataset is limited, the use of motion fields can successfully reconstruct consistent 4D volumes where the image intensity approach cannot.

This paper is structured as follows. We firstly describe our method which is based on MA and graph theory. We then describe the synthetic and *in vivo* datasets we use, the experiments we performed, and our results, ending with discussion and conclusions.

2. METHODS

We describe our methodology in two stages. First, we outline how MA can be used to perform the global alignment of the data, and second we describe how graph theory can be used to define the inter-dataset comparison term used in MA.

2.1. Manifold Alignment

ML is a tool for non-linear dimensionality reduction which aims to extract low dimensional manifolds from high-dimensional datasets. We denote the high-dimensional data by $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T]$, which consists of T points in \mathbb{R}^D . In previously published MRI slice stacking techniques the data \mathbf{X} have consisted of image intensities [4, 1, 6], but in this paper we investigate a novel approach of using image-derived motion fields as discussed in section 4. Assuming that \mathbf{X} lie on or close to a manifold \mathcal{M} of dimension d , ML constructs a map from $\mathbb{R}^{T \times D}$ to $\mathbb{R}^{T \times d}$ where $d \ll D$. The resulting points form a new low-dimensional dataset, $\mathbf{Y} = [\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_T]$ which describe each point's position on \mathcal{M} . Many methods have

This work was supported by the Engineering and Physical Sciences Research Council under Grant EP/M009319/1 and by the Wellcome EP-SRC Centre for Medical Engineering at Kings College London (WT 203148/Z/16/Z).

been proposed to achieve this embedding, such as Isomap [10], Locally Linear Embedding [11], and Laplacian Eigenmaps [12].

In some cases, two or more different datasets may be described by similar manifolds, even if the high-dimensional data do not appear to be similar or easily comparable. MA is an extension of ML, in which multiple high-dimensional datasets are mapped to one common low-dimensional space, whilst preserving relations between points in the same dataset, and also some specified inter-dataset relations.

ML techniques which work by optimising a cost function can be extended to perform MA by adding terms to the cost function which represent inter-dataset alignment [8]. Here we use Laplacian Eigenmaps (LE) [12], so we briefly review how LE can be extended to perform MA.

LE involves firstly forming a graph \mathcal{G} based on $k_{\mathcal{G}}$ nearest-neighbour relations, with the edge between i and j weighted by

$$W_{ij} = \exp \left[\frac{-\|\mathbf{x}_i - \mathbf{x}_j\|^2}{2\sigma_{\mathcal{G}}^2} \right], \quad (1)$$

where $\sigma_{\mathcal{G}}$ determines the strength of neighbourhood relations. The graph Laplacian, \mathbf{L} is given by

$$\mathbf{L} = \mathbf{D} - \mathbf{W} \quad (2)$$

where \mathbf{D} is the diagonal degree matrix $D_{ii} = \sum_j W_{ij}$. The eigenvectors of \mathbf{L} corresponding to the d smallest non-zero eigenvalues provide the low-dimensional coordinates \mathbf{Y} . These are the coordinates minimising the cost term $\Phi(\mathbf{Y})$,

$$\Phi(\mathbf{Y}) = \sum_{i,j} W_{ij} \|\mathbf{y}_i - \mathbf{y}_j\|^2 \quad (3)$$

subject to the constraint that $\mathbf{Y}^T \mathbf{D} \mathbf{Y} = \mathbf{I}$. MA is achieved by extending this formulation to the case of N high-dimensional datasets. The joint cost term becomes

$$\Phi_{total} = \sum_{\ell} \Phi^{(\ell)} + \frac{\mu}{2} \sum_{n \neq m} \sum_{i,j} U_{ij}^{(n,m)} \|\mathbf{y}_i^{(n)} - \mathbf{y}_j^{(m)}\|^2 \quad (4)$$

where $\Phi^{(\ell)}$ is the cost term for each individual dataset, $\mathbf{U}^{(n,m)}$ is some similarity kernel between datasets $\mathbf{X}^{(n)}$ and $\mathbf{X}^{(m)}$, and μ is the parameter that weights the intra-dataset matching versus the inter-dataset matching.

The aligned coordinates $\mathbf{Y}^{(\ell)}$ minimising this cost are given by the smallest non-zero eigenvectors of the block matrix \mathbf{M} where

$$\mathbf{M} = \begin{bmatrix} \mathbf{L}^{(1)} + \mu \mathbf{I} & -\mu \mathbf{U}^{(1,2)} & \dots & -\mu \mathbf{U}^{(1,N)} \\ -\mu \mathbf{U}^{(2,1)} & \mathbf{L}^{(2)} + \mu \mathbf{I} & \dots & -\mu \mathbf{U}^{(2,N)} \\ \vdots & \vdots & \ddots & \vdots \\ -\mu \mathbf{U}^{(N,1)} & -\mu \mathbf{U}^{(N,2)} & \dots & \mathbf{L}^{(N)} + \mu \mathbf{I} \end{bmatrix}. \quad (5)$$

2.2. Graph-Based Descriptors

The key to implementing Eq (5) is the definition of the inter-dataset similarity kernel, $\mathbf{U}^{(n,m)}$. Since direct comparison of data from different high-dimensional datasets is often not possible, we must transform the original data in some way that facilitates comparison. One approach is to form graphs representing each high-dimensional dataset and compute descriptors based on the graph structures. This was the approach taken in [6], in which a descriptor based on the steady-state distribution of a lazy random walker on the graph was proposed. In the next section we outline a new descriptor that is based on work from the computer vision literature.

2.2.1. Wave Kernel Signature

The wave kernel signature (WKS) [9] is part of a family of graph matching methods which use the eigenvectors of the graph's Laplacian to compare vertices. The graph Laplacian, \mathbf{L} , can be interpreted as a discrete version of the Laplace-Beltrami operator and so can be used to describe diffusive processes on the graph [12]. We denote the Laplacian's eigenvalues as E_k and eigenvectors as ϕ_k . The WKS is a function $\omega_i(z)$ for each node i in the graph, defined as

$$\omega_i(z) = C(z) \sum_k \phi_{k,i}^2 \exp \left[\frac{-(z - \log(E_k))^2}{2\sigma_z^2} \right] \quad (6)$$

where $C(z)$ is a normalisation term given by

$$C(z) = \left(\sum_k \exp \left[\frac{-(z - \log(E_k))^2}{2\sigma_z^2} \right] \right)^{-1}. \quad (7)$$

This function is a stable and highly informative descriptor [13] which corresponds to the diffusion of a quantum mechanical particle of energy z on the graph [9]. The parameter σ_z is a measure of the 'smoothness' of this descriptor which is normally constant and manually chosen for the task at hand.

The similarity between two nodes, i and j in the two graphs n and m can be assessed by measuring a distance, $\Delta_{ij}^{(n,m)}$, between their wave kernel signatures, where

$$\Delta_{ij}^{(n,m)} = \int_{z_{\min}}^{z_{\max}} \frac{\omega_i^{(n)}(z) - \omega_j^{(m)}(z)}{\omega_i^{(n)}(z) + \omega_j^{(m)}(z)} dz \quad (8)$$

The similarity kernels are then given by

$$U_{ij}^{(n,m)} = \exp \left[\frac{-(\Delta_{ij}^{(n,m)})^2}{2\sigma_{\text{WKS}}^2} \right] \quad (9)$$

which ensures that vertices with similar wave kernels have a high similarity in U . We set $\sigma_{\text{WKS}} = 1$.

As noted above, the value of σ_z is normally manually chosen based on how similar the graphs are. Very similar graphs can be better matched with a low σ_z and vice versa. In MRI slice stacking the task is to align the manifolds of multiple 2D slice datasets, some of which may be similar (e.g. adjacent slices) and some of which may be dissimilar (e.g. far apart slices). Therefore, in this paper we propose a novel extension of the WKS in which different values of σ_z are used for each pair of datasets and these are automatically computed based on the similarities of the graph's eigenvalues:

$$\sigma_z^{(i,j)} = a + b \sum_k |\log(E_k^{(i)}) - \log(E_k^{(j)})| \quad (10)$$

where a and b are constants found by grid search. We refer to this method as *adaptive* WKS, or adWKS.

In Section 4, we also evaluate the widely used heat kernel signature (HKS) [14], which differs in that the exponential term in (6) and (7) is replaced with $\exp[-E_k z]$.

As in [6], once a feature descriptor has been used to generate a similarity kernel $\mathbf{U}^{(n,m)}$, the kernel is then sparsified by using the Hungarian algorithm to establish one-to-one correspondences with the maximal similarity. These sparsified kernels are used for MA.

3. MATERIALS

We evaluate our MRI slice stacking method on synthetic and real MRI datasets. The first dataset consists of a sequence of synthetic but highly realistic dynamic 3D MRI volumes of the thorax as described in [15]. Briefly, this synthetic dataset was generated using image registration of a respiratory-gated high spatial resolution 3D MRI volume to a series of dynamic 3D low spatial resolution MRI volumes. The high resolution volume was warped using the registration results to create a series of realistic high spatial resolution volumes at different respiratory motion states. We use 250 volumes sampled with an interval of 0.64s each containing 40 sagittal slices. We direct the reader to [15] for further details on the generation of this dataset. To use this dataset for evaluating the slice stacking methods, at each time point, a single slice was extracted from the current volume (and the rest left out); the datasets used for slice stacking comprised this extracted slice and all slices from the other time points. This process was repeated for all time points to simulate a real slice-by-slice acquisition. We repeat this process for each slice, giving 40 reconstructed volumes for each time point and report the median result across these 40 volumes.

The real MRI slice datasets we use are those presented in [16]. The data was acquired from four healthy subjects and had a field of view covering the entire thorax, including the lungs and liver. The data consists of N sagittal slices of thickness 8mm, where N is typically around 30. The 2D images were acquired by taking one image from each slice position, iterating through the slices one by one, and then repeating this process until 40 images were obtained for each slice position. In this acquisition, one image was acquired per heart-beat (at systole) so as to isolate respiratory motion. The acquisitions were carried out on a Philips Achieva 3T MR scanner using a T1-weighted gradient echo sequence with an acquired in-plane image resolution of $1.4 \times 1.4 \text{ mm}^2$, a slice thickness of 8 mm, repetition and echo times (TR and TE) of 3.1 and 1.9 ms, a flip angle of 30 degrees, and a SENSE-factor of 2. The field of view covering the entire thorax was $400 \times 370 \text{ mm}^2$, and each slice took around 180 ms to acquire.

For both experiments we assess the use of image intensities and registration derived motion fields as the high-dimensional data used to build the manifold. As in [16], NiftyReg [17] was used to estimate the 2D motion fields. In the synthetic dataset, images are registered to an end-exhale respiratory position, and in the real dataset to end-exhale breath-hold images.

4. EXPERIMENTS AND RESULTS

We quantitatively evaluate the use of MA for MRI slice stacking based on our proposed adWKS descriptor and compare it to the use of three alternatives: the standard WKS descriptor (i.e. without our novel adaptive selection of σ_ω), the random walk based descriptor as described in [6], and the heat kernel signature [14] (see Section 2.2.1). For each descriptor, and at each time point i , to reconstruct a volume around a slice acquired at position n , we stack together images from other slice positions $m \neq n$ such that the whole volume, V_i has a consistent respiratory position. Specifically, if our current image is $\mathbf{x}_i^{(n)}$, we find the $j \neq i$ for each slice $m \neq n$ to minimise $\|\mathbf{y}_i^{(n)} - \mathbf{y}_j^{(m)}\|^2$ (i.e. the distance between the images in the low-dimensional manifold), and then use the image $\mathbf{x}_j^{(m)}$ for slice position m in the volume V_i . Stacking all of these slices together produces a respiratory-resolved 3D volume corresponding to the 2D image $\mathbf{x}_i^{(n)}$. Repeating this process for each time point results in a

sequence of dynamic 3D volumes created from the dynamic 2D data.

Figure 1 shows a sample embedding \mathbf{Y} of the synthetic data, produced using the adWKS descriptor. Note that each point in this plot represents a 2D slice (all slice positions are included), and the colour indicates the respiratory position of the ground truth volume from which it was extracted (estimated using a virtual navigator). The smooth variation in colour demonstrates that the manifold has accurately captured the variation in respiratory state across all slices.

4.1. 4D MRI Reconstruction of Synthetic Volumes




Fig. 1. Example of points in aligned manifold for all sagittal slices in synthetic MRI dataset. Each point here represents a 2D sagittal MRI image. The colours represent respiratory position.

For the first experiment we have ground truth volumes to compare the stacked slices to. To quantify the accuracy of the reconstructions we calculate the L_2 distance between these reconstructed volumes and the corresponding ground truth volumes. Because the error distribution was skewed, we use the median error across all slices for each time point as a summary statistic, resulting in a median reconstruction error for each time point for each evaluated method. Figure 2 summarises these errors. It can be seen that the WKS significantly outperforms the random walk descriptor and the heat kernel signature, and that the adaptive choice of σ_ω produces a further small but statistically significant reduction in error.

Statistical significance was determined with $p < 0.01$ according to a two-tailed Wilcoxon signed-rank test. In this experiment fully connected graphs were used with $\sigma_G = 1.5$ for the image-intensity data, and $\sigma_G = 80$ for the motion-field data. Method specific parameters were chosen by grid search and the values used here were $\sigma_{RW} = 0.01$, $\sigma_\omega = 0.8$, $a = 0.5$, $b = 0.5$, and $\mu = 0.05$. It should be noted though that the superior performance of the WKS and adWKS methods held for a wide range of parameters.

4.2. 4D MRI Reconstruction of Real Data

Our second experiment used real data acquired from four health subjects. In this experiment there is no ground truth volume to compare the stacked slices to. Therefore we quantify the quality of the reconstructed volumes by computing the consistency of the positions of the left and right hemidiaphragms. We automatically estimate the diaphragm position in a coronal slice through the reconstructed volumes by finding the pixels with the largest difference to their neigh-




Fig. 2. Sum of square differences between reconstructed volumes and ground truth volumes for different graph descriptors. Each box plot represents 250 values, one for each time point of the synthetic 4D dataset. The wave kernel methods outperform the HKS and RW approaches, regardless of the type of high-dimensional data used. The minimal error is for the adWKS method using image-intensity data which has a small but statistically significant advantage over the other wave kernel methods.

hours in the head-foot direction as shown in the example in Figure 3. We then measure the Pearson’s correlation coefficient between the head-foot positions of the points in the left lung with those in the right lung over all time points. If this correlation is high then we judge the reconstructed volumes to be consistent. We use the correlation rather than a distance measure between the two hemidiaphragms because one may be higher than the other even if they are moving together and our aim is to observe consistent motion which is captured by measuring the correlation. Results of this experiment are shown in table 1.

Volunteer	Data	Graph descriptor			
		adWKS	WKS	HKS	RW
A	Registration	0.958	0.957	0.803	0.954
	Image	0.539	0.336	0.401	0.918
B	Registration	0.914	0.904	0.674	0.870
	Image	0.518	0.524	0.814	0.480
C	Registration	0.903	0.864	0.533	0.339
	Image	0.746	0.754	0.569	0.890
D	Registration	0.810	0.801	0.685	0.722
	Image	0.617	0.458	0.629	0.624

Table 1. Pearson’s correlation coefficient between left and right hemidiaphragm positions of reconstructed volumes. A volume can be reconstructed from each sagittal slice, each giving its own correlation coefficient; here we report the median across these slices. The best result for each volunteer is shown in bold.




Fig. 3. Top: examples from volunteers A (left) and D (right) of a coronal slice through the original unaligned volumes. The sagittal slices are not in consistent motion states resulting in discontinuities in hemidiaphragm positions. Bottom: examples from a volume reconstructed by stacking sagittal slices aligned by motion state. Estimated diaphragm positions for left and right lungs are shown in red.

5. DISCUSSION AND CONCLUSIONS

MA often requires a way of comparing data between disparate datasets, and graph-based descriptors have been proposed for this purpose. In this paper we have proposed a novel descriptor, which is an extension of one previously proposed in the computer vision literature. Our results showed that the WKS-based descriptors outperformed the state-of-the-art descriptors. Furthermore, the adWKS descriptor allows for graph-based information to be exchanged between different MRI slice positions whilst accounting for how similar or different their graphs are. Intuitively, those slices which represent very different anatomy will have different graphs, and so they should have their images matched whilst taking this difference into account, whereas those which are more similar, such as neighbouring slices, can be matched more precisely. Our results demonstrated that this flexibility resulted in an improvement in image reconstruction over the use of a constant σ_ω value.

We also found that in the second experiment, in which the amount of data was more limited, the use of motion fields to generate the low-dimensional manifold was crucial to reconstructing consistent volumes. Although the use of motion fields requires more computation than using raw images, it may allow for successful slice-stacking even when the underlying manifold of motion states is more sparsely sampled.

Future work will consider whether alternative approaches in the MA step, tailored towards situations where data is limited, can also improve image reconstruction accuracy. We will also investigate the application of these techniques to other problems in medical imaging, such as 4D ultrasound compounding of images acquired from different acoustic windows [6].

6. REFERENCES

- [1] C.F. Baumgartner, C. Kolbitsch, D.R. Balfour, P.K. Marsden, J.R. McClelland, D. Rueckert, and A.P. King, “High-resolution dynamic MR imaging of the thorax for respiratory motion correction of PET using groupwise manifold alignment,” *Med. Image Anal.*, vol. 18, no. 7, pp. 939–952, 2014.
- [2] M. von Siebenthal, G. Szekely, U. Gamper, P. Boesiger, A. Lomax, and P. Cattin, “4D MR imaging of respiratory organ motion and its variability,” *Phys. Med. Biol.*, vol. 52, no. 6, pp. 1547, 2007.
- [3] C. Würslin, H. Schmidt, P. Martirosian, C. Brendle, A. Boss, N.F. Schwenzler, and L. Stegger, “Respiratory motion correction in oncologic PET using T1-weighted MR imaging on a simultaneous whole-body PET/MR system,” *J. Nucl. Med.*, vol. 54, no. 3, pp. 464–471, 2013.
- [4] N. Dikaios, D. Izquierdo-Garcia, M.J. Graves, V. Mani, Z.A. Fayad, and T.D. Fryer, “MRI-based motion correction of thoracic PET: initial comparison of acquisition protocols and correction strategies suitable for simultaneous PET/MRI systems,” *European radiology*, vol. 22, no. 2, pp. 439–446, 2012.
- [5] C. Wachinger, M. Yigitsoy, E.J. Rijkhorst, and N. Navab, “Manifold learning for image-based breathing gating in ultrasound and MRI,” *Med. Image Anal.*, vol. 16, no. 4, pp. 806–818, 2012.
- [6] C.F. Baumgartner, A. Gomez, L.M. Koch, J.R. Housden, C. Kolbitsch, J.R. McClelland, D. Rueckert, and A.P. King, “Self-aligning manifolds for matching disparate medical image datasets,” in *IPMI*. Springer, 2015, pp. 363–374.
- [7] R. Guerrero, C. Ledig, and D. Rueckert, “Manifold alignment and transfer learning for classification of Alzheimers disease,” in *International Workshop on Machine Learning in Medical Imaging*. Springer, 2014, pp. 77–84.
- [8] C.S. Lee, A. Elgammal, and M. Torki, “Learning representations from multiple manifolds,” *Pattern Recognit.*, vol. 50, pp. 74–87, 2016.
- [9] M. Aubry, U. Schlickewei, and D. Cremers, “The wave kernel signature: A quantum mechanical approach to shape analysis,” in *Proc. IEEE Int’l Conf Computer Vision Workshops*. IEEE, 2011, pp. 1626–1633.
- [10] J.B. Tenenbaum, V. de Silva, and J.C. Langford, “A global geometric framework for nonlinear dimensionality reduction,” *Science*, vol. 290, no. 5500, pp. 2319–2323, 2000.
- [11] S.T. Roweis and L.K. Saul, “Nonlinear dimensionality reduction by locally linear embedding,” *Science*, vol. 290, no. 5500, pp. 2323–2326, 2000.
- [12] M. Belkin and P. Niyogi, “Laplacian eigenmaps for dimensionality reduction and data representation,” *Neural Computation*, vol. 15, no. 6, pp. 1373–1396, 2003.
- [13] N. Hu, R.M. Rustamov, and L. Guibas, “Stable and informative spectral signatures for graph matching,” in *Proc. CVPR*, 2014, pp. 2305–2312.
- [14] J. Sun, M. Ovsjanikov, and L. Guibas, “A concise and provably informative multi-scale signature based on heat diffusion,” in *Proc. Symposium on Geometry Processing*, 2009, SGP ’09, pp. 1383–1392.
- [15] X. Chen, M. Usman, C.F. Baumgartner, D.R. Balfour, P.K. Marsden, A.J. Reader, C. Prieto, and A.P. King, “High-resolution self-gated dynamic abdominal MRI using manifold alignment,” *IEEE Trans. Med. Imaging*, vol. 36, no. 4, pp. 960–971, 2017.
- [16] C.F. Baumgartner, C. Kolbitsch, J.R. McClelland, D. Rueckert, and A.P. King, “Autoadaptive motion modelling for MR-based respiratory motion estimation,” *Med. Image Anal.*, vol. 35, pp. 83–100, 2017.
- [17] M. Modat, G.R. Ridgway, Z.A. Taylor, M. Lehmann, J. Barnes, D.J. Hawkes, N.C. Fox, and S. Ourselin, “Fast free-form deformation using graphics processing units,” *Computer methods and programs in biomedicine*, vol. 98, no. 3, pp. 278–284, 2010.