# Manifold Learning for Image-Based Breathing Gating in Ultrasound and MRI

Christian Wachinger[a,*], Mehmet Yigitsoy[a], Erik-Jan Rijkhorst[b], Nassir Navab[a]

[a] *Computer Aided Medical Procedures (CAMP), Technische Universität München, München, Germany*
[b] *Centre for Medical Image Computing (CMIC), University College London, London, UK*

---

## Abstract

Respiratory motion is a challenging factor for image acquisition and image-guided procedures in the abdominal and thoracic region. In order to address the issues arising from respiratory motion, it is often necessary to detect the respiratory signal. In this article, we propose a novel, purely image-based retrospective respiratory gating method for ultrasound and MRI. Further, we apply this technique to acquire breathing-affected 4D ultrasound with a wobbler probe and, similarly, to create 4D MR with a slice stacking approach. We achieve the gating with Laplacian eigenmaps, a manifold learning technique, to determine the low-dimensional manifold embedded in the high-dimensional image space. Since Laplacian eigenmaps assign to each image frame a coordinate in low-dimensional space by respecting the neighborhood relationship, they are well suited for analyzing the breathing cycle. We perform the image-based gating on several 2D and 3D ultrasound datasets over time, and quantify its very good performance by comparing it to measurements from an external gating system. For MRI, we perform the manifold learning on several datasets for various orientations and positions. We achieve very high correlations by a comparison to an alternative gating with diaphragm tracking.

*Keywords:* Image-based breathing gating, Manifold Learning, 4D, ultrasound, MRI

---

## 1. Introduction

Respiration is a cyclic, irregular motion that leads to deformations in the abdominal and thoracic region. The respiratory signal monitors the current breathing phase of the patient. For numerous applications, it is necessary to assign each image the corresponding respiratory phase in which it was acquired. One example is image mosaicing, where only the fusion of images from the same breathing state leads to consistent panorama images (Wachinger et al., 2008; Wachinger and Navab, 2009).An alternative to breathing gating are breath-hold acquisitions, but they further complicate the procedure and are dependent on the patients ability for breath-hold. A further application is radiation therapy, where organ motion due to respiration can lead to inaccuracies during target localization (Flampouri et al., 2006). These inaccuracies require the treatment margins to be much larger than the tumor size, resulting in exposure of the healthy tissue in the surroundings to a high radiation dose (Colgan et al., 2008). It is therefore necessary to study the respiratory motion in the planning process to decrease treatment margins. For an accurate analysis, 4D imaging techniques are necessary in order to handle the motion of the organs over time (Remmert et al., 2007; Li et al., 2008). Prospective and retrospective gating methods exist, with advantages of retrospective approaches being discussed in (Rohlfing et al., 2001). We describe the creation of 4D ultrasound (US) and 4D magnetic resonance (MR) data with a retrospective approach in further details throughout the article.

The measurement of the breathing phase is generally achieved with external gating systems, attached to the patient. The disadvantage of the usage of such systems are long setup times, the prolongation of the overall acquisition, and its high costs. Additionally, the image data stream and the respiratory signal have to be synchronized, which is not trivial. For certain imaging devices, such as CT and MR scanners, solutions for the synchronization exist, however, we are not aware of such a possibility for ultrasound; leaving the synchronization to the user. The consequence is that gating systems are rarely used in practice. In this article, we

---

[*] Corresponding Author

*Email addresses:* `wachinge@in.tum.de` (Christian Wachinger), `yigitsoy@in.tum.de` (Mehmet Yigitsoy), `e.rijkhorst@cs.ucl.ac.uk` (Erik-Jan Rijkhorst), `navab@in.tum.de` (Nassir Navab)

*Preprint submitted to Medical Image Analysis* *December 2, 2011*

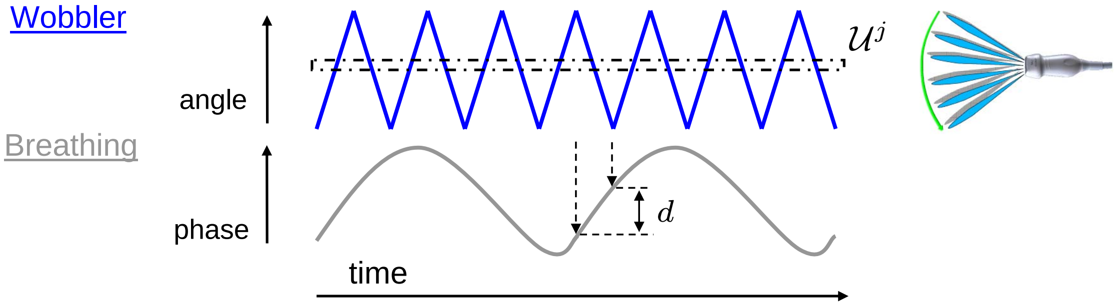

Figure 1: Wobbler angle (blue) and respiratory phase (gray) over time. Dashed lines indicate respiratory change $d$ within one sweep. Dash dotted line indicates frames from same angle over several breathing cycles.

propose a purely image-based retrospective respiratory gating system using manifold learning. The proposed method is fully automatic and does not need any prior information about the anatomy, training data, or user interaction. The basic performance of our algorithm is to assign to a stream of images acquired from the same position over time the corresponding respiratory signal. It builds upon the assumption that the breathing cycle forms a continuous manifold in image space with points at similar positions on the manifold related by the state of the breathing cycle. In this work, we deal with 2D images, however, we are not limited to it. In the following, we explain how we apply this basic technique for the creation of 4D ultrasound and 4D MR data. A preliminary version of this work was previously presented at conferences (Wachinger et al., 2010b; Yigitsoy et al., 2011).

### 1.1. *4D Ultrasound with Wobbler*

One application that we investigate in more details, and for which we have not yet found a solution proposed in the literature, is the acquisition of breathing-affected 4D ultrasound with a mechanically steered transducer, also referred to as wobbler. The problems for using a wobbler in such a scenario is that images in one sweep do not contain consistent information, but represent the anatomy in different breathing states. We illustrate this in Figure 1, where we schematically plot the deviation angle of the wobbler together with the respiratory signal over time. The phase difference $d$ indicates the range of breathing phases accumulated in one sweep. We propose to select all frames acquired from the same angle (dash dotted line) and apply the image-based gating on each of these sets of images. Having the respiratory signal estimated for each angle, we align these local curves and apply a robust spline curve fitting to create a globally consistent respiratory signal. This, consequently, allows us to reconstruct volumes for specific breathing stages.

An alternative to the application of a wobbler to obtain 4D ultrasound would be a native 3D transducer with elements arranged on a 2D array. Such systems, however, are still expensive and the access to data streaming and radio frequency data is very restricted. Our proposed method is, nevertheless, also interesting in conjunction with 2D array transducers, because we can perform the breathing-gating on the acquired volumetric data.

### 1.2. *4D MRI with Navigator Slices*

Several techniques based on MRI have been proposed in the literature to handle the respiratory motion. Among these are breath-hold and slice stacking techniques. The problem with breath-hold is that patients may not be able to hold their breath during the acquisition. von Siebenthal et al. (2007) use the slice stacking method where 2D slices from different locations having the same breathing state are stacked together to reconstruct a 3D image for that state. They acquire dedicated high quality slices, called *navigator slices*, at a fixed location to determine a similarity criterion to be applied during the sorting of data slices.

The acquisition process for the slice stacking approach is illustrated in Figure 2. It is differentiated between *data slices* $D_i^p$ with $p$ indicating the position and $i$ indicating time and *navigator slices* $N_i$. As already mentioned, the navigator slices are acquired at a fixed location and are used to estimate the breathing phase. An alternating acquisition scheme,

$$\{\ldots, N_i, D_{i+1}^p, N_{i+2}, D_{i+3}^{p+1}, N_{i+4}, D_{i+5}^{p+2}, \ldots\}, \qquad (1)$$

is applied to interleave the data slices and navigator slices, as shown in Figure 2. Further details about the acquisition process are presented in (von Siebenthal, 2008).

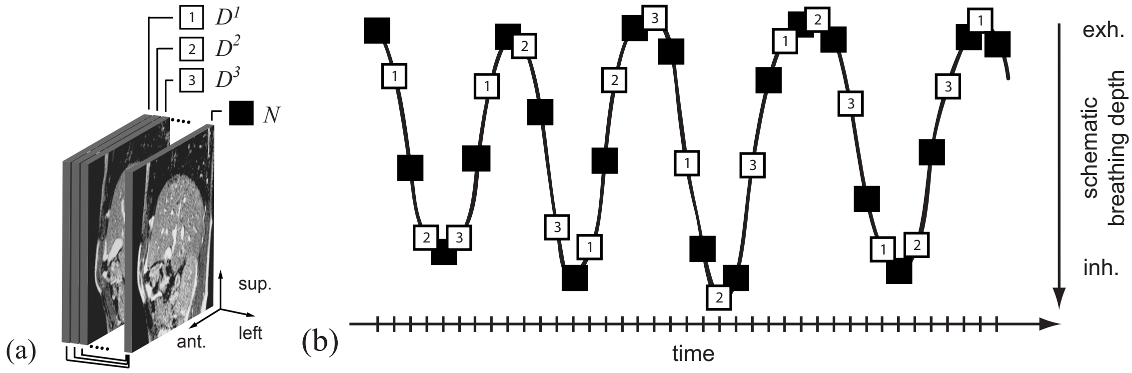

Figure 2: (a) Sagittal slices from the volume of interest. Data slices $D^1, D^2, D^3$ and the dedicated navigator slice, $N$. (b) The interleaved acquisition of data and navigator slices. Solid squares indicate navigator slices while the others indicate data slices with the position number displayed in the boxes. In this case only 3 positions are illustrated. Courtesy of von Siebenthal et al. (2007).

In order to find data slices that are acquired in the same breathing state, its surrounding navigator slices are compared. The rationale is to assume that if the enclosing navigator slices are similar, then the data slices are in the same breathing state. For comparing navigator slices, non-rigid registrations are performed (von Siebenthal, 2008). Although this produces the desired results, the processing time, ranging from 10 to 100 hours, is the major drawback of this method for usage in clinical practice. Our proposed method is able to recover the breathing signal from the navigator slices within seconds. It can either be applied on its own or as a pre-processing step for the non-rigid registration by defining a search window. This significantly reduces the processing time for the non-rigid registration because of the pre-selection of navigator slices to be registered.

Once 4D MRI data is available, it can be used to build a continuous motion model by interpolating between the discrete time points of the 4D data set. To this end, a respiratory phase needs to be assigned to each volume in the set, in order to solve for the motion model coefficients (Rijkhorst et al., 2010). In our experiments, we show that the proposed method is suitable for identifying the respiratory phase on such data.

### 1.3. *Comparison of Both Applications*

The described applications for 4D imaging with ultrasound and MR are very similar. Both build upon the retrospective gating of 2D slices. With the gating information, it is possible to select frames that were acquired in the same breathing state and to consecutively construct a volume for each state. The major difference between both applications is the acquisition of navigator slices for MR. The availability of the navigator slices facilitates the application because the extracted respiratory signal serves as global gating signal and the sorting of the data slices is done relatively to it. As we will describe later on in more details, the proposed technique for 4D ultrasound imaging aligns local breathing curves to construct the global signal.

## 2. Related work

There are many articles on image-based gating in ultrasound for detecting the cardiac motion (Treece et al., 2002; Zhu et al., 2003; de Winter et al., 2003; Karadayi et al., 2006). These approaches apply techniques that are either (i) specific to detecting the cardiac signal *e.g.* centroid algorithm (Karadayi et al., 2006), (ii) based on user interaction (Treece et al., 2002), or (iii) designed for intravascular ultrasound (Zhu et al., 2003; de Winter et al., 2003). In (Sundar et al., 2009), a general technique for breathing gating is proposed and applied to ultrasound data. It bases on the phase correlation technique to estimate the motion between successive frames. The breathing phase is estimated from the energy change between consecutive frames. The inherent limitation of the phase correlation algorithm is that it finds the *global translation* in the image plane. Considering that ultrasound images and MR slices are 2D cross sections of the body, the organ motion is not necessarily in-plane, and consequently, there is no uniform global translation. This issue is further discussed in Section 4. In von Siebenthal (2008), the non-rigid registration of navigator slices is applied for image-based breathing gating in MRI.

Several manifold learning techniques were proposed in the literature with common techniques being Isomap (Tenenbaum et al., 2000), locally linear embedding (Roweis and Saul, 2000), and Laplacian eigenmaps (Belkin and Niyogi, 2003). Since its introduction, manifold learning has been applied for a multitude of applications, including segmentation (Zhang et al., 2006), registration (Rohde et al., 2008; Hamm et al., 2009; Wachinger and Navab, 2010), tracking (Lee and Elgammal, 2007), recognition (Arandjelovic and Cipolla, 2007; Wachinger et al., 2010a), detection (Jamieson et al., 2010; Suzuki et al., 2010), computational anatomy (Gerber et al., 2009), and 4D CT reconstruction of the lung (Georg et al., 2008). For the reconstruction approach, manifold learning is performed on slabs, where a manual inspection is necessary to crop the slabs to contain only lung tissue. Further, Isomap is used to create the low-dimensional embedding. In our work, we focus on Laplacian eigenmaps, since it provided us with better results in comparison to Isomap. Further, we deal with the specific challenges of the integration of 4D ultrasound wobbler data and 4D MRI slice stacking. Since each 4D imaging technique has its own acquisition protocol, the common manifold learning framework has to be adapted accordingly, as discussed in the next section.

## 3. Manifold Learning

The general idea of manifold learning is to project a manifold in high-dimensional space $\mathbb{R}^{\mathcal{N}}$ to a low-dimensional space $\mathbb{R}^n$, while preserving the local neighborhood. We propose the application of Laplacian eigenmaps (Belkin and Niyogi, 2003) for the respiratory phase estimation because the technique is well founded on mathematical concepts (Laplace Beltrami operator) and computationally efficient. Laplacian eigenmaps build upon the construction of a neighborhood graph that approximates the manifold, on which the data points lie. Subsequently, the graph Laplacian is applied to calculate a low-dimensional representation of the data that preserves locality.

Considering $k$ points $\mathbf{a}_1, \ldots, \mathbf{a}_k$ in $\mathbb{R}^{\mathcal{N}}$ lying on a manifold $\mathcal{M}$, we want to find a set of corresponding points $\mathbf{b}_1, \ldots, \mathbf{b}_k$ in the low-dimensional space $\mathbb{R}^n$ ($n \ll \mathcal{N}$). We define weights $\mathbf{W}$ between all pairs of input points, which reflect the locality of points. One possibility for the weights proposed in (Belkin and Niyogi, 2003) is derived from the heat kernel

$$\mathbf{W}_{i,j} = \exp(-\|\mathbf{a}_i - \mathbf{a}_j\|_2^2 / t). \tag{2}$$

The parameter $t$ weights the influence of neighboring points. The optimization in Laplacian eigenmaps that tries to preserve the locality of points in low-dimensional space is

$$\underset{\{\mathbf{b}_1,\ldots,\mathbf{b}_k\}}{\arg\min} \sum_{i,j} \mathbf{W}_{i,j} \cdot \|\mathbf{b}_j - \mathbf{b}_i\|_2^2. \tag{3}$$

We see that points that are close in high-dimensional space should be arranged close in low-dimensional space to avoid a high cost caused by the high weight. Calculating the diagonal matrix $\mathbf{D}_{i,i} = \sum_j \mathbf{W}_{i,j}$, we can construct the Laplacian matrix $\mathbf{L}$ of the graph

$$\mathbf{L} = \mathbf{D} - \mathbf{W}. \tag{4}$$

With the relationship $\sum_{i,j} \mathbf{W}_{i,j} \cdot \|\mathbf{b}_j - \mathbf{b}_i\|_2^2 = 2 \cdot \mathrm{trace}(\mathbf{B}^\top \mathbf{D}\mathbf{B} - \mathbf{B}^\top \mathbf{W}\mathbf{B})$, the optimization can be formulated as trace minimization

$$\min_{\substack{\mathbf{B} \in \mathbb{R}^{k \times n} \\ \mathbf{B}^\top \mathbf{D}\mathbf{B} = \mathbf{I} \\ \mathbf{b}^{(l)\top} \mathbf{D}\mathbf{1} = \mathbf{0}}} \quad \mathrm{trace}(\mathbf{B}^\top \mathbf{L}\mathbf{B}) \tag{5}$$

with $\mathbf{B} = [\mathbf{b}_1, \ldots, \mathbf{b}_k]^\top$ and further considering column vectors $\mathbf{b}^{(l)}$, enabling to write $\mathbf{B} = [\mathbf{b}^{(1)}, \ldots, \mathbf{b}^{(n)}]$. The constraints guarantee, among others, the orthogonality of the different dimensions, preventing a collapse onto a subspace of dimension less than $n$. Many dimensionality reduction techniques can be formulated as trace optimization problems (Kokiopoulou et al., 2011). The optimal solution of the optimization problem is given by the first $n$ eigenvectors corresponding to the lowest, non-zero eigenvalues of the generalized eigenvalue problem $(\mathbf{D} - \mathbf{W})\mathbf{v} = \lambda \mathbf{D}\mathbf{v}$. The solution is unique up to unitary transformations (Kokiopoulou et al., 2011).

### *3.1. Manifold Learning for Gating*

Considering $k$ images $\mathcal{U} = \{\mathbf{u}_1, \ldots, \mathbf{u}_k\}$ that are acquired over several breathing cycles, we want to find the corresponding breathing phase $\phi_i$ of each image $\mathbf{u}_i$. In terms of manifold learning, the images correspond to the high-dimensional points, $\mathbf{a}_i = \mathbf{u}_i$, and the phase to the low-dimensional embedding, $\mathbf{b}_i = \phi_i$. We consider one dimension of the ambient space for each image pixel, so $\mathcal{N}$ is corresponding to the resolution of the images. For the low-dimensional space, we investigated $n = \{1, 3\}$ as possible dimensions. We also experimented with the two dimensional space, however, the results were similar to the one dimensional case. This behavior may be due to the repeated eigendirections problem, which is associated to the specific shape of the manifold, as discussed in (Gerber et al., 2007).

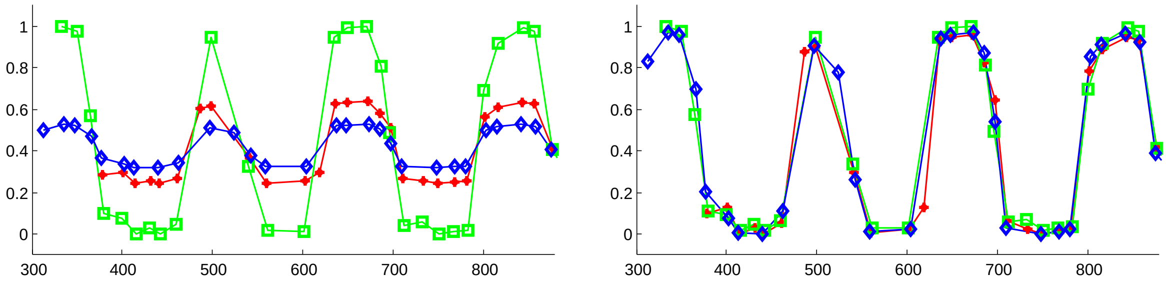

(a) Not Aligned

(b) Aligned

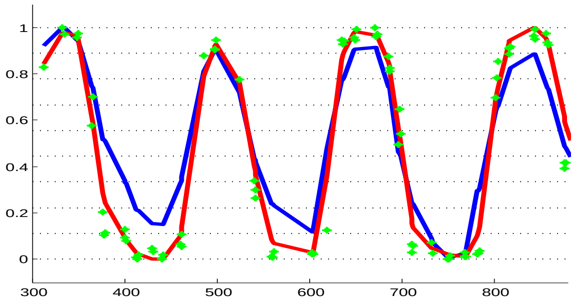

(c) Reconstructed

Figure 3: (a) Local breathing signals from manifold learning before alignment. Illustrated is the case for 3 angles (3 colors). X-axis indicates ultrasound frame number. (b) Local breathing signals after alignment. (c) Reconstructed global breathing signal (red) is calculated by robustly fitting a spline curve through the aligned local signals (green crosses). The ground truth signal is shown in blue. Dotted lines indicate the separation of the breathing cycle into several stages. For each stage a volume is compounded.

The manifold learning $m$ then assigns each image a coordinate in the low-dimensional space

$$m : \mathbb{R}^{\mathcal{N}} \to \mathbb{R}^{n} \tag{6}$$

$$\mathbf{u}_i \mapsto \phi_i, \tag{7}$$

with $1 \leq i \leq k$. The suggestion that images during free breathing lie on a low-dimensional manifold in the ambient space is justified because variations between neighboring slices are smooth, and further, slices from the same respiratory phase but different acquisition times are similar. Each image in the respiratory cycle corresponds to a point in high-dimensional space. During breathing, we are moving back and forth along this manifold or trajectory in high-dimensional space. The underlying optimization problem of manifold learning tries to optimally preserve local information, cf. equation (3), meaning that similar images are mapped to similar positions in the low-dimensional space. With manifold learning, we are therefore able to project the manifold in high-dimensional space, with the images lying on it, to low dimensions. Since the dominant change in the images in the abdominal region is due to respiration, the respiratory information is obtained with the first or first few lower dimensions. Consequently, it is reasonable to use $\phi_i$ as an estimate of the respiratory phase. Important to notice is that we do not use a parameterization of the expected motion, as it is the case in (Sundar et al., 2009) with global translation. The proposed method can therefore deal with complex motion patterns, *e.g.* local translation, rotation, deformation, out-of-plane motion, because the low-dimensional embedding is only based on the similarity of images.

The similarity measure is important for neighborhood selection and weighting, where the calculation of the Euclidean distance between the points is equivalent to calculating the sum of squared differences (SSD) between the images. A vast number of similarity measures is proposed in the context of medical image registration (Hajnal et al., 2001). Since we deal with monomodal data for our application, we investigate the performance of SSD and the correlation coefficient. The calculation of correlation coefficient is up to additive and multiplicative constants equivalent to the calculation of SSD on normalized input images (Viola, 1995). We therefore only have to normalize the input images to achieve the performance of the correlation coefficient. The normalization is done by subtracting the mean and by dividing by the standard deviation. Once the neigh-

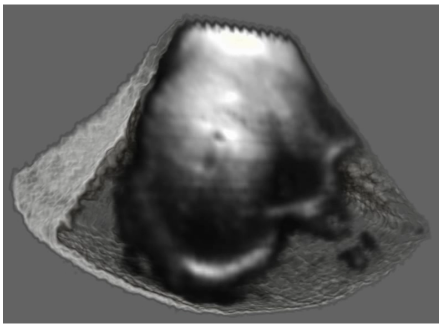

Figure 4: 3D rendering of compounded volume for one of the breathing stages.

borhood graph is constructed, the eigenvectors of the graph Laplacian provide the embedding map.

### *3.2. Global Consistency in 4D US*

The image acquisition processes for US and MR bear certain similarities, however, the availability of navigator slices for MR facilitates the application. For MR, we only perform the manifold learning on navigator slices, which are acquired at the same position, leading directly to a global respiratory signal. For ultrasound, we do not have those specific slices and have to perform the manifold learning on the data slices. The idea is to estimate a local breathing signal for each deflection angle and, subsequently, calculate the global respiratory signal from the local ones.

Given $\mathcal{U}$ the set of all acquired images, we partition the set in disjunct subsets $\mathcal{U}^1, \ldots, \mathcal{U}^\alpha$, corresponding to the number of different deflection angles $\alpha$ of the wobbler (dash dotted region in Figure 1). We perform the manifold learning for each of the subsets separately $m^j(\mathbf{u}_i) = \phi_i^j$, with $1 \leq j \leq \alpha$. So depending on the acquisition angle of the ultrasound image $\mathbf{u}_i$, the corresponding manifold learning $m^j$ is performed. Considering all the phases estimated from one angle, we have the local respiratory signals $\Phi^j = \{\phi_1^j, \ldots, \phi_v^j\}$, with $v$ the number of frames per angle. Each local signal contains a consistent estimation of the breathing signal. It is, however, not possible to directly compare local signals, because the 1D projection of the manifold learning can be in an arbitrary range. This is illustrated in Figure 3(a) with exemplary three local signals corresponding to three angular positions. A simple normalization of each of the local signals $\Phi^j$ is not sufficient because the extreme positions of the breathing cycle may not be reached within them. Consequently, we affinely register local signals in order to retrieve the best scaling $s^j$ and

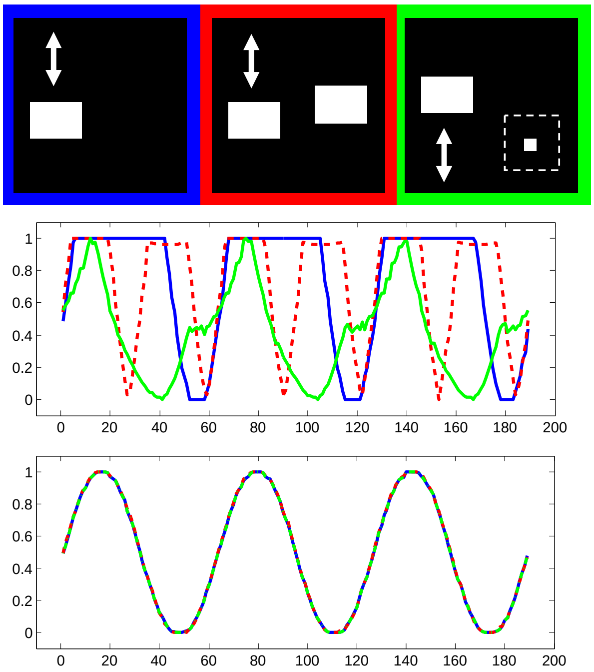

Figure 5: Analysis of the gating techniques for synthetic images. Three different motion scenarios are illustrated in the top row. The corresponding gating curves for the phase correlation are shown in the middle row. The last row shows results for manifold learning. Note that in this case the gating results for all three scenarios are identical and corresponding to the surrogate signal.

translation $t^j$

$$\Phi^j \mapsto s^j \cdot \Phi^j + t^j. \tag{8}$$

Note that this is a 1D affine registration and that scaling and translation are only performed in breathing phase direction (y-axis). We do not have to register in temporal direction, because the acquisition time of the images is provided by the ultrasound system. This is, in fact, a groupwise registration scenario, where we choose to align each pair of neighboring curves with a pairwise registration, starting from the middle one. The sum of the Euclidean distances between the temporally closest points serves as cost function. The result of the alignment is shown in Figure 3(b).

The values of the partial signals $\Phi^j$ are now comparable, but may still contain outliers. Consequently, we apply a robust curve fitting to all the sample points to retrieve the global breathing signal. We experimented with various curve models, including Fourier, sum of sine waves, and splines. We achieved best results with fitting a spline curve because it allows for the most flexibility, which is important due to irregularity of free breathing. The value of the fitted curve then represents the breathing phase of the ultrasound frames, see Figure 3(c).

In a final step, the breathing cycle is classified into several breathing stages. For each of the breathing stages, the ultrasound frames along the various angles are gathered, and compounded into a final volume, see Figure 3(c) and 4.

## 4. Experiments for 4D Ultrasound

For our experiments we use an ultrasound system from Ultrasonix (Richmond, Canada) and an optical tracking system from A.R.T. (Weilheim, Germany). Both systems are connected to a workstation PC. For the synchronization, we time stamp the data on the tracking system and use a network time server to calculate the time offset. For the ultrasound data, we use the direct streaming of B-mode images over the network. We perform tests on eight patient datasets acquired from different positions, focusing on the liver and kidney. We acquire 2D ultrasound images over time with a curved linear transducer (C5-2/60) and for the 3D+t experiments we use a wobbler probe (4DC7-3/40). The acquisition frequency was in the range of 3-5 MHz and the depth setting varied between 14 and 20 cm. All acquisitions were performed with the patient lying in the supine position.

In order to validate our results, we compare them to the measurements of an external gating system. Martinez-Möller et al. (2007) compared four different gating systems and obtained the best results with an elastic belt and an optical tracking system. We place a tracking target, consisting of 7 retro-reflective marker balls, on the chest of the patient. Our setup is more accurate than the one in (Martinez-Möller et al., 2007) for PET tracking, because our field of view is not hindered by the tube, enabling the tracking with four cameras from different views. Since we are only interested in the relative motion of the tracking target, and not in the absolute pose, constant target offsets and calibration errors do not influence the result; leading to a precise monitoring of the respiratory motion. The tracking system provides the pose of the tracking target in 3D space, which consists of 3 translational and 3 rotational components. We apply a principal component analysis of the 6D tracking data to find the principal component along which direction we measure the breathing motion. This is, for instance, done analogously for motion modeling in lung radiotherapy (McClelland et al., 2006). Further, we low-pass filter the signal to remove cardiac motion and extract the respiratory signal. We refer to the tracked signal as ground truth, which is not completely correct because it contains tracking errors. However, it is the best that can currently be achieved (Martinez-Möller et al., 2007) and is sufficient to validate the performance of our image-based approach.

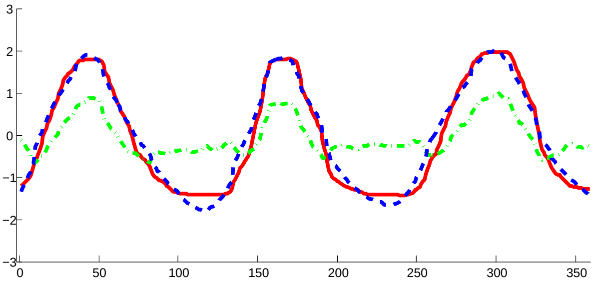

Figure 7: Breathing gating results for 2D (red solid: estimated signal, blue dashed: ground truth, green dash-dotted: phase correlation).

| 2D Corr Coeff | |
|---|---|
| liver1 | 95.4 % |
| liver2 | 94.4 % |
| liver3 | 93.6 % |
| kidney | 97.3 % |

Table 1: 2D+t results

| 3D Corr Coeff | |
|---|---|
| liver 30° | 94.3 % |
| liver 45° | 95.8 % |
| liver 60° | 96.8 % |
| kidney 45° | 94.4 % |

Table 2: 3D+t results

We compare our approach to the phase correlation technique for 2D+t images applied in (Sundar et al., 2009) [1]. In order to have a fair comparison, we use 2D images sequences over time from different parts of the liver. Unfortunately, with the phase correlation technique, we do not achieve meaningful results for our datasets. We think that this is due to the limitation of the approach to approximate the 3D motion with a global translation in 2D. In order to illustrate this limitation, we produced synthetic images that show periodic motion. The first scenario consists of a rectangle moving up and down, see Figure 5. For the second, we add a fixed rectangle, and for the third we add a rectangle that grows and shrinks (see additional material for videos). We plot the corresponding energy curves of the phase correlation technique. We further show the gating results for our manifold learning approach in Figure 5, which are identical for all three scenarios and corresponding to the ground truth signal. The result with the phase correlation technique for the first scenario (blue) is approximating the true signal. The addition of a fixed object (red)

---

[1] We want to thank the authors of (Sundar et al., 2009) for sharing source code.

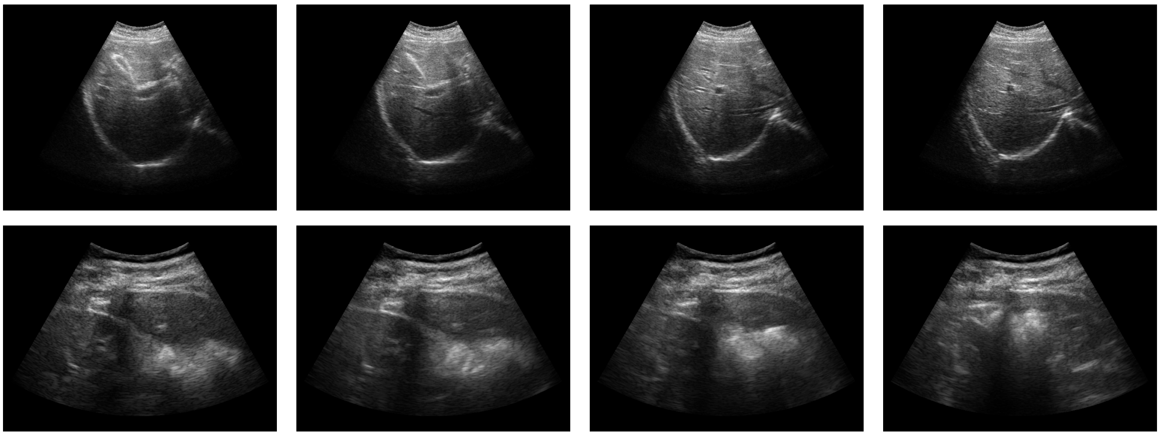

Figure 6: 2D Ultrasound images over time from liver (abdomen, right upper quadrant, oblique section) and kidney (left lateral decubitus position, right intercostal flank section).

leads to a signal with double frequency and the addition of the shrinking/growing object (green), leads to a further distortion of the motion signal. Since the results are already not optimal for this easy synthetic case, it is comprehensible that this approach is not best suited for breathing estimation in a noisy ultrasound or MR environment with 3D anatomy moving in and out of plane. Additional results of the phase correlation technique on real data are discussed in following paragraphs.

The resolution of our ultrasound images is $640 \times 480$ pixels. We downsample the images in each direction by a factor of 2, leading to $\mathcal{N} = \frac{1}{4} \cdot 640 \cdot 480$. This enables a faster processing and leads to no noticeable degradation of the manifold learning. We show excerpts of two data sets in Figure 6. We perform all our experiments with a graph neighborhood of $l = 14$. The number of images for manifold learning varies between 100 and 300, where we did not notice a dependency of the performance on the number of input samples. For the low-dimensional space, we perform experiments with embeddings to 1D and 3D. In the case of the ultrasound images, however, the results are very similar, so that we concentrate on the 1D case.

In Figure 7, we show the result of the respiratory gating for one of the 2D datasets together with the ground truth signal. For comparison, we also plot the result of the phase correlation technique, which confirms the findings in the synthetic experiments. We also calculate the correlation coefficient (CC) between the detected and ground truth signal for multiple 2D data sets, shown in table 1. It is remarkable that the ground truth signal is almost perfectly detected. All peaks in the ground truth signal also appear in the detection. Further, the calculation of the correlation, which is in the range of 95%, confirms the visual similarity of the graphs. We also experimented with normalizing the images before passing them to the manifold learning, noticed however no significant improvement.

For the 4D experiments, we show the result of a fitted curve in Figure 3. We also calculate the correlation coefficient between the fitted curves and ground truth for four datasets, see table 2. We experimented with three different angular ranges, 30°, 45°, and 60° (maximum of probe), for which the probe steers to 15, 21, and 29 different angular positions. We split the breathing signal into 9 different breathing stages, and compound a 3D volume for each of the stages. A volume rendering of one of the volumes is shown in Figure 4. The additional material contains a video showing the 4D volume rendering.

We perform one-sided significance tests for unknown mean and unknown standard deviation to assess the statistical properties of our results. Under the null-hypothesis $H_0$ that the correlation between tracking and gating is lower than 93.6%, we obtain a p-value lower than $p < 0.005$. The null-hypothesis can therefore be rejected and the result is considered to be statistically significant.

All image-based approaches rely on ultrasound acquisitions from the same position, because otherwise it is not possible to differentiate between probe motion and breathing motion. To investigate this assumption, we attached a tracking target to the transducer and analyzed its trajectory. This analysis showed only a negligible deviation. The still position therefore does not limit the applicability of our method, which is also confirmed

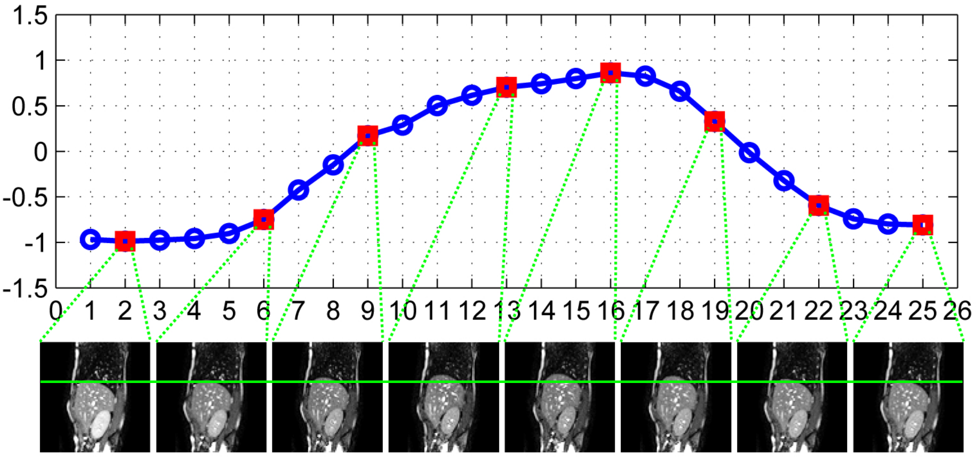

Figure 8: Sample slices selected from one breathing cycle. Red-square markers indicate the slice locations in the signal. The auxiliary line assists in observing the liver movement.

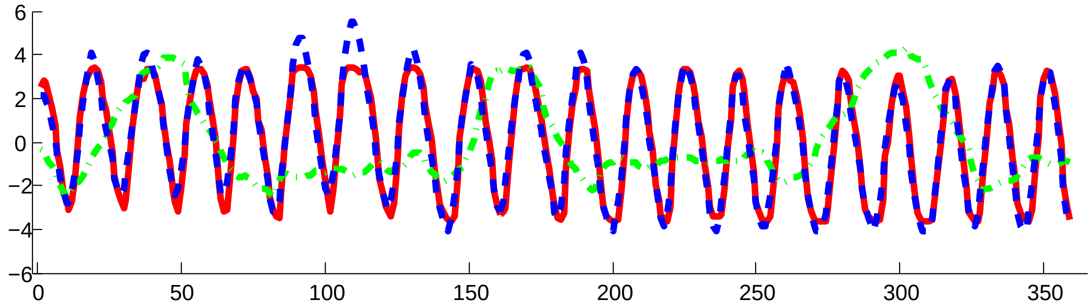

Figure 9: Breathing gating results for MRI navigator slices from ETH dataset (red solid: estimated signal, blue dashed: ground truth, green dash-dotted: phase correlation). The correlation between the manifold learning and ground truth is 98%, between phase correlation and ground truth is 2%.

by our good gating results.

We perform statical tests to analyze the significance of our results. Under the null-hypothesis $H_0$ that the correlation between tracking and gating is lower than 93.6%, we obtain a p-value lower than $p < 0.005$. The null-hypothesis can therefore be rejected and the result is considered to be statistically significant.

All image-based approaches rely on ultrasound acquisitions from the same position, because otherwise it is not possible to differentiate between probe motion and breathing motion. To investigate this assumption, we attached a tracking target to the transducer and analyzed its trajectory. This analysis showed only a negligible deviation. The still position therefore does not limit the applicability of our method, which is also confirmed by our good gating results.

## 5. Experiments for 4D MRI

We perform experiments on four different MRI datasets to evaluate the proposed method. Two of them originate from ETH and two from UCL. For each dataset, diaphragm tracking is performed to have a ground truth (GT) signal to compare to. More details on the utilized diaphragm tracking are presented in (Timinger et al., 2005; Nguyen et al., 2009; King et al., 2009). The general disadvantage of diaphragm tracking is its limitation to certain slice positions and orientations. For our experiments this is fine because we only need a reference signal and we can select the slice with the best visibility of the diaphragm for tracking. A further advantage of the proposed method is that it is applicable to almost any slice orientation and position, as is shown in following. Moreover, diaphragm

Table 3: Statistical analysis of results for gating on MRI. Correlation coefficients are calculated for all orientations and all positions for 3 datasets. Mean CC and standard deviations for each orientation are indicated. We further calculated the CC for the null hypothesis such that the p-value of a one-tailed t-test is $p < 0.005$.

| Datasets | Orientation | Number of Positions | Mean CC | STD CC | CC s.t. ($p < 0.005$) |
|---|---|---|---|---|---|
| ETH | Sagittal | 25 | 98.5 % | 0.90 | 98.0 % |
| | Coronal | 164 | 99.2 % | 0.31 | 99.1 % |
| | Axial | 256 | 99.1 % | 0.66 | 99.0 % |
| UCL | Sagittal | 78 | 97.1 % | 1.37 | 96.7 % |
| | Coronal | 143 | 95.4 % | 2.61 | 94.8 % |
| | Axial | 288 | 94.2 % | 4.46 | 93.5 % |
| UCL (filtered) | Sagittal | 78 | 97.7 % | 0.94 | 97.4 % |
| | Coronal | 143 | 95.7 % | 2.60 | 95.1 % |
| | Axial | 288 | 94.4 % | 3.19 | 93.9 % |
| UCL Motion Model | Sagittal | 159 | 98.9 % | 0.54 | 98.8 % |
| | Coronal | 168 | 98.7 % | 0.98 | 98.5 % |
| | Axial | 200 | 98.5 % | 1.13 | 98.3 % |

tracking necessitates the manual placement of a tracking window, while the proposed method is totally automatic.

### 5.1. ETH Datasets

The first dataset is the navigator slices acquired at ETH by von Siebenthal et al. (2007), as described in Section 1.2. These slices have a spatial resolution of 255 × 255, a temporal resolution of about 2.7 Hz, and are acquired with a Philips 1.5T Achieva. The sagittal plane was chosen as the imaging plane due to its ease of tracking vascular structure with minimal out-of-plane motion, since the dominant motion of the liver is in the superior-inferior direction (Rohlfing et al., 2004). A reconstructed breathing cycle with sample MR images is shown in Figure 8. We further show a comparison between the reconstructed signal of all navigator slices and the ground truth respiratory signal in Figure 9. The visual similarity of the curves is confirmed by a correlation of 98%. Also illustrated is the signal obtained from the phase correlation technique, leading to a correlation of 2%.

von Siebenthal (2008) proposes two methods for retrospective sorting. The first one involves the tracking of regions of interest. The disadvantage is that these regions have to be defined manually, taking the trackability and out-of-plane motion of vascular structures into account (von Siebenthal, 2008). More comparable to our method is the second approach, where a deformable registration is performed between all navigator slices, leading to $O(k^2)$ deformable registrations with $k$ the number of images. Although this still requires the segmentation of the liver on one of the slices, the manual interaction is reduced. In our case, most costly is the calculation of the neighborhood graph with the weights, leading to $O(k^2)$ calculations of SSD. Since SSD is evaluated in each iteration of the deformable registration, we clearly see the computational benefit of our approach, next to the advantage of being fully automatic.

The second dataset is the publicly available 4D data from ETH[2], which is created following the slice stacking principle. Each volume consists of 256 × 256 × 25 voxels with a spatial resolution of $1.37 \times 1.37 \times 4\text{mm}^3$. 14 volumes are available for one breathing cycle. We perform manifold learning on 2D slices at all possible positions and orientations, leading to 537 separately estimated respiratory signals with manifold learning in this experiment. This provides further insights whether certain regions or orientations are better suited for gating, which is further discussed in Section 6. We plot the correlation coefficient with respect to the GT signal for all orientations and positions in Figure 10. Also in this figure, we show cross-sectional views to have a better overview to which slice position and orientation a specific result corresponds. We further perform a statistical analysis of these results, which is summarized in table 3. We exclude the results of coronal slices showing only background from the statistics. Shown is the mean correlation coefficient, the standard deviation, and the correlation coefficient of the null-hypothesis of the t-test such that we obtain significant results ($p < 0.005$).

---

[2] http://www.vision.ee.ethz.ch/4dmri/

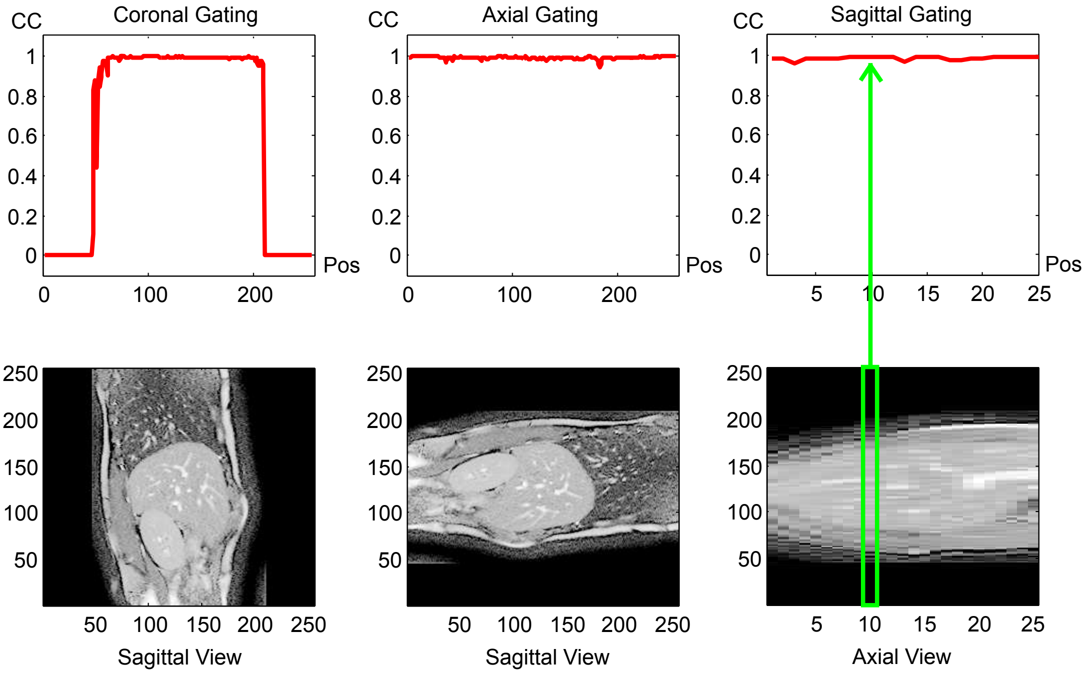

Figure 10: The three graphs show the correlation coefficient of the estimated signal with the ground truth for the different orientations of the ETH dataset. For each slice position and orientation, a separate gating is performed, leading to 537 separate gatings in this experiment. The images in the second row are cross-sections, with the x-axis of the graph and the image corresponding. To provide an overview of sagittal slices, for instance, we show an axial view.

Consider the coronal ETH slices, for instance, the hypothesis that the correlation between the ground truth and tracking signal is above 99.1% is statistically significant ($p < 0.005$). The CC in the test is lower for the sagittal direction because of the lower number of slices, which determines the degrees of freedom of the student's $t$ distribution function. But nevertheless, the CC is in the range of 98.0% to 99.1%.

### 5.2. *UCL Datasets*

The third dataset consists of free-breathing MR scans from UCL with a field-of-view covering the whole abdomen. A balanced-SSFP sequence (TR/TE=4.3/1.46 ms, 30° flip angle) was used to obtain high-resolution 4D dynamic scans during free breathing using a 1.5 T MR scanner (Philips Achieva, Best, The Netherlands) at Guy's Hospital, London. Parallel imaging with a 32-channel coil array using a SENSE acceleration factor of 4.6 resulted in scan times of approximately one second per dynamic volume. Since the largest liver motion normally occurs in the sagittal plane (von Siebenthal et al., 2007; Rijkhorst et al., 2010), the highest reconstruction resolution of $1.4 \times 1.4$ mm was chosen in this plane, resulting in a slice thickness of 4 mm. This leads $288 \times 288 \times 78$ voxels in the volume. 25 volumes were acquired over 4 breathing cycles. Once again, manifold learning is performed on 2D slices for all possible positions and orientations. We plot the correlation coefficient with respect to the GT signal for all orientations in Figure 11. The statistical analysis is summarized in table 3. In contrast to the ETH dataset, the UCL dataset contains more noise. This is comprehensible because the ETH dataset is the result of a sophisticated and time consuming slice stacking approach, while the volumes of the UCL dataset are acquired in real-time with a fast imaging protocol. This leads to slightly worse gating results. We are, however, able to improve the result by pre-processing the data with median filtering. The results of the gating on the original slices and noise reduced slices are shown in Figure 11. The results of the statistical analysis are summarized in table 3. The mean CC before filtering ranges between 94.2% and 97.1%, while after the filtering it is between 94.4% and 97.7%. The standard deviation is decreased by about 40% in sagittal and axial orientations after filtering.

The fourth 4D dataset is created by using a motion model similar to the one presented in (Rijkhorst et al., 2010). A set of dynamic 4D MR data was acquired

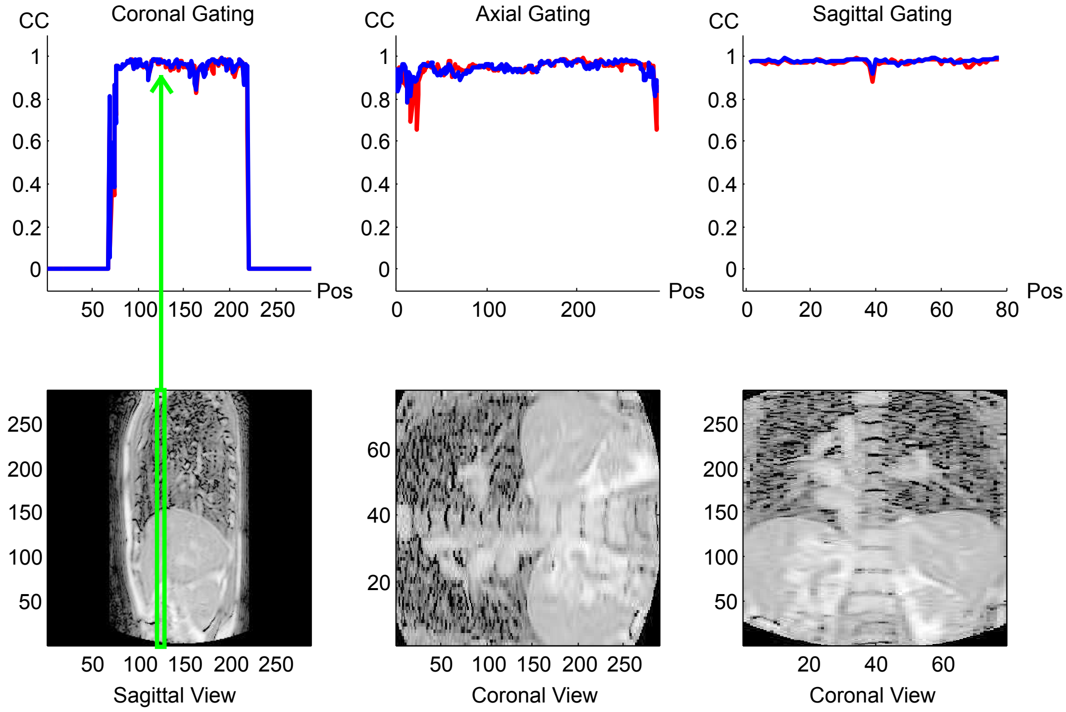

Figure 11: The three graphs show the correlation coefficient of the estimated signal with the ground truth for the different orientations of the UCL dataset. The red solid line shows the results of gating on original data, the blue dashed line shows the the results on filtered data. For each slice position and orientation, a separate gating with only those slices is performed, leading to 509 separate gatings in this experiment. The images in the second row are cross-sections, with the x-axis of the graph and the image corresponding. To provide an overview of coronal slices, for instance, we show a sagittal view.

as described above, and tissue displacements throughout the liver were computed by registering each volume within the set to a reference volume using a non-rigid fluid registration method (Crum et al., 2005). A ground truth surrogate respiratory signal is computed by positioning a sector inside each MR volume, and using a navigator window at the location of the diaphragm. By combining the signal with the registration results, a second order polynomial is fitted at each spatial location, allowing for the creation of motion fields at arbitrary time points. We apply the motion model to a separately obtained breath-hold 3D MRI to create 4D MRI. The resolution is $1.25 \times 1.25 \times 1.5$ mm and $190 \times 200 \times 160$ voxels. Once again, manifold learning is performed on 2D slices at all possible positions and orientations. We plot the correlation coefficient with respect to the GT signal for all orientations in Figure 12. The last position on sagittal slices is constant, due to boundary effects of the motion model, and does consequently not lead to usable gating results. The statistical analysis is summarized in table 3, with the mean CC ranging between 98.5% and 98.9%.

### *5.3. Embedding Dimensionality*

In this section, we want to further analyze the dimensionality of the embedding space and perform experiments for 3D and 1D. The projection of the ETH navigator slices to 1D space is shown in Figure 13, and the projection to 3D space in Figure 15. We arbitrarily select one of the slices as reference slice. We then look for the closest points in 1D and 3D space with respect to the reference image, calling them the best match in 1D and 3D, respectively. As distance measure we use the Euclidean distance. The image corresponding to the 1D best match is shown in the left column of Figure 14, and the 3D best match is shown in the right column. Considering the auxiliary lines, circles, and arrows, we observe that the breathing state of the reference image and 1D best match is not equivalent. In contrast, the slice corresponding to the 3D best match is almost completely identical to the reference image. Showing that the embedding to 3D has advantages in finding similar slices. In order to get further insights, we also show the best match in 3D on the 1D signal in Figure 13 and the best match in 1D on the 3D plot in Figure 15. We ob-

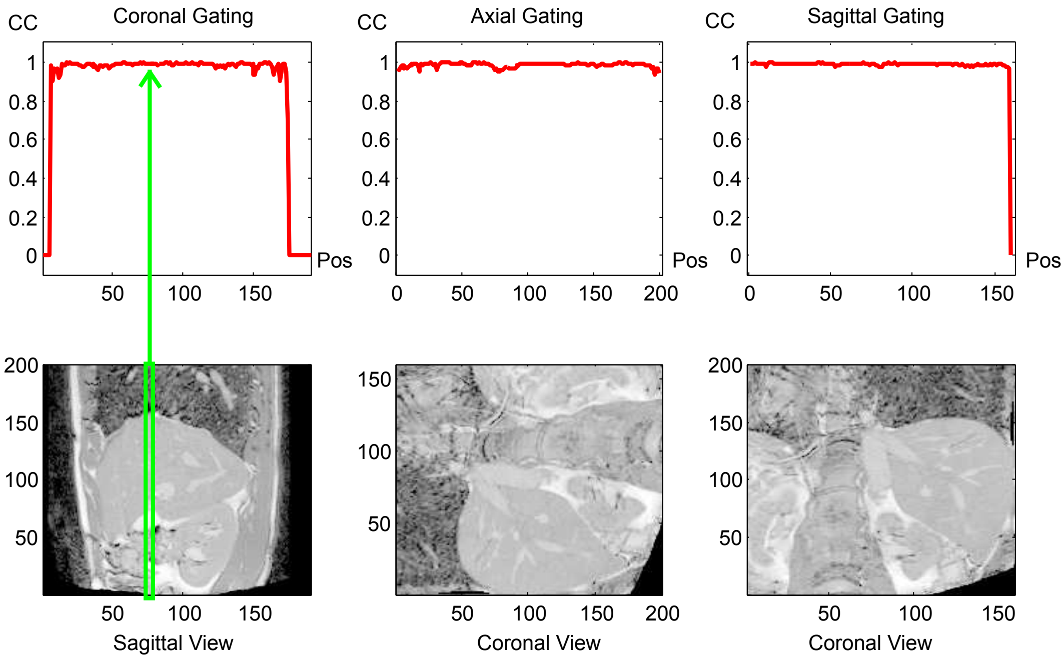

Figure 12: The three graphs show the correlation coefficient of the estimated signal with the ground truth for the different orientations of UCL dataset created with the motion model. For each slice position and orientation, a separate gating with only those slices is performed, leading to 550 gatings in this experiment. The images in the second row are cross-sections, with the x-axis of the graph and the image corresponding. To provide an overview of coronal slices, for instance, we show a sagittal view.

serve that the 1D best match is pretty far away from the reference image in the 3D plot.

For some applications, like finding the closest slice, the embedding into 3D is fine. For other purposes, such as visualization, instantiation of a motion model (Rijkhorst et al., 2010), or comparison to a 1D ground truth, a 1D signal is more appropriate. If we want to reduce the 3D signal to 1D, we have to perform another dimensionality reduction. Considering Figure 15, it seems possible to achieve an improved 1D embedding by performing a second dimensionality reduction on the 3D embedding. The successive application of multiple manifold learning causes no problems. In Figure 16, we compare the direct embedding to 1D and the combination of first an embedding to 3D and a successive reduction to 1D on navigator slices. We observe that the plots are very similar, which is confirmed by correlations of 98.2% and 98.8%, respectively. We perform the same experiment on the UCL motion model data. The results are very similar, with correlations of 99.8% and 99.9%, respectively. These results show that two successive dimensionality reductions do not significantly improve the result and that a direct embedding to 1D is reasonable, if a 1D signal is needed.

## 6. Discussion

For ultrasound, we achieved correlations between the proposed method and external gating of around 95%. For the ETH and UCL motion model data, the correlation is in the range of 98%. For the UCL free-breathing MR data it is in the range of 95%, with a slight improvement after filtering. Noticeable is the very low standard deviation. This shows that the proposed method is very versatile and leads in almost all scenarios to excellent gating results. For the two UCL datasets, we achieved the lowest performance on axial slices, and the best results on sagittal slices. This is in line with previous observations (von Siebenthal et al., 2007). The slightly lower correlation for the sagittal planes on the ETH data is a bit surprising and may be due to the large slice thickness.

Our experiments have further shown that the phase correlation technique was not able to extract the respiratory signal from the data. Diaphragm tracking is a valid alternative, does however require user interaction by placing the window and is further limited to views showing the diaphragm. Our results show that the proposed method can be applied to arbitrary orientations and that even gating on axial planes leads to good re-

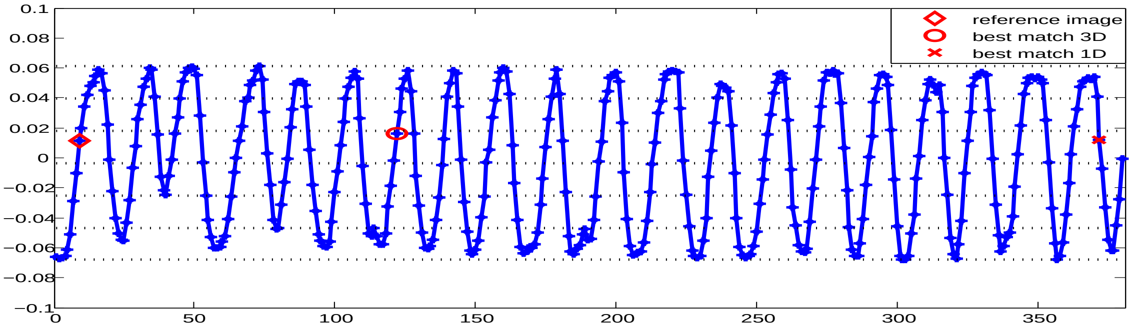

Figure 13: Calculated respiratory signal from 1D. X axis shows the slice numbers. Y axis shows corresponding 1D values from manifold learning. We highlight the reference slice and its best matches using 1D and 3D signals.

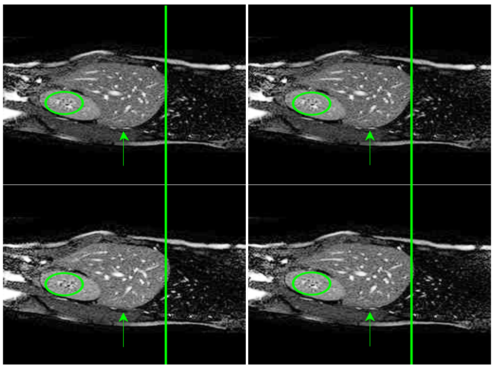

Figure 14: Results using 1D and 3D signals. Top row: reference slice, twice. Bottom row left: best match using 1D signal. Bottom row right: best match using 3D signal. Red line, circle and arrow assist comparison.

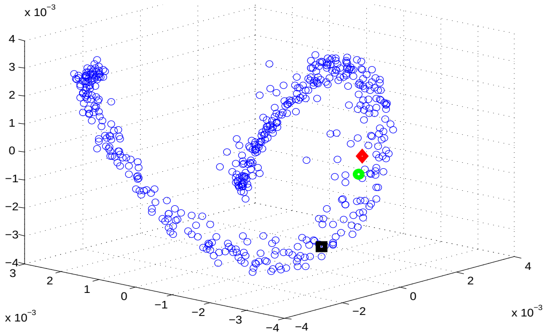

Figure 15: Scatter plot of manifold learning results using 3D. Red-diamond: reference slice, green-circle: 3D best match, black-square: 1D best match

sults. For ultrasound, diaphragm tracking is even more limited, because it is difficult to have a nice view of the diaphragm during the entire breathing cycle; showing a clear advantage of our method.

The respiratory signal is displaying the current respiratory state of the patient. It is, however, not entirely clear, how this state is defined. One possibility would be to measure the amount of air in the lungs, another possibility would be to measure the displacement of the diaphragm. While there is definitely a high correlation between such quantifications, they are not identical. Since we want to automatically extract the respiratory signal from the data, it is important for our evaluation to relate to a ground truth respiratory signal. Due to the lack of a global consensus, we compared our results to the tracking results from an external tracking system and the tracking of the diaphragm, which are both commonly used procedures. We also want to point out that these respiratory signals are prone to errors during acquisition, however, the error is generally low enough to make the comparison in our scenario meaningful.

The discussion about the definition of the respiratory signal is continued when it comes to the dimensionality. Our experiments show, that we achieve a better discrimination for MR with an embedding to 3D than 1D. The 3D signal over time can definitely not be as nicely visualized as the 1D signal, however, it may as well be considered as an adequate representation. The question whether we first have to extract a 1D signal out of the 3D signal to have a usable respiratory signal cannot be generally answered, but is dependent on the specific application. If the interest lies in finding the closest slices, as it is the case for 4D imaging, this search can directly be performed in 3D space. For an easy visualization and comparison to an alternative gating signal, we may however be interested in a 1D signal. We showed in our experiments that a 1D signal can be obtained from a 3D embedding through the subsequent application of

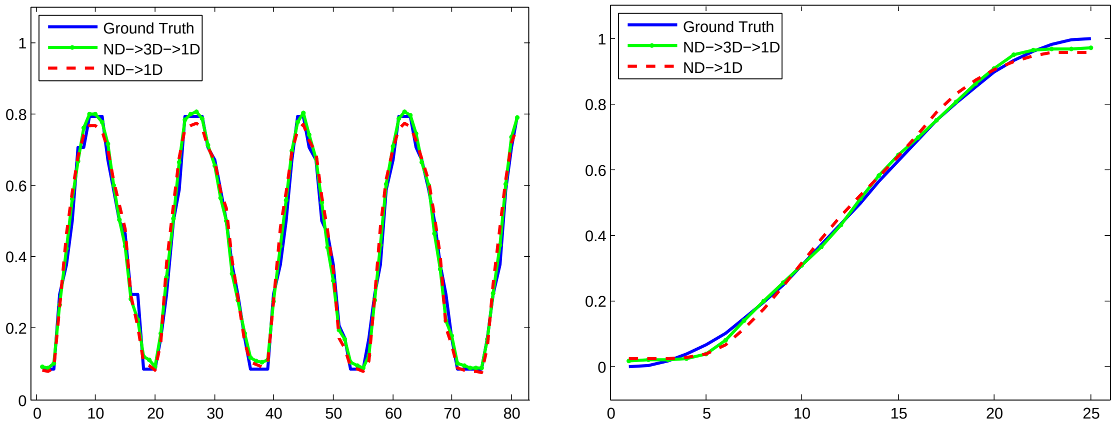

(a) ETH navigator slices

(b) UCL data

Figure 16: Gating experiment performed on the ETH navigator slices and UCL data to evaluate the difference between a direct 1D embedding and a 3D embedding with a successive 1D reduction. For the ETH data, the CC between the direct embedding to 1D and the GT is 98.2% and the CC between the 3D→1D embedding and GT is 98.8%. For the UCL data, the CC between the direct embedding to 1D and the GT is 99.8% and the CC between the 3D→1D embedding and GT is 99.9%.

second manifold learning, the performance is, however, similar to a direct reduction to 1D. Finding closest slices could also be performed in the original image space, it does, however, not respect the geometry of the manifold, which can lead to inaccuracies for a larger number of required neighbors. In contrast, manifold learning does respect the geometry and the number of nearest neighbors chosen for the graph construction can be much lower than the number of slices selected for reconstruction.

With the presented embedding to 1D and 3D we were not able to differentiate between images from inspiration and expiration. It is, however, possible to differentiate between them by looking at the extracted signal. The only requirement is to have several frames per cycle, which is the case for the presented applications. An analysis of the direct differentiation with manifold learning is subject to future research.

## 7. Conclusion

We presented an automatic, image-based respiratory gating method for ultrasound and MR using manifold learning. Moreover, we proposed a solution for acquiring 4D breathing data with a wobbler probe and also acquiring 4D MR with the slice stacking approach. Our method has the advantage that it is fully automatic and does not require a training phase or prior information about the underlying anatomy, nor the interaction of the user. To analyze the performance of our algorithm for ultrasound, we perform experiments on various datasets showing different organs and sections. The results of these experiments were very good, for both, 2D and 3D ultrasound data over time. For MRI, we worked on four different datasets and performed manifold learning on all positions and orientations. A comparison to a tracking-based gating approach is performed, leading to almost similar results and very high correlation. Finally, our approach presents an attractive alternative to external tracking and gating systems with their various setup issues and synchronization problems.

The proposed method is not limited to analyzing breathing motion, but could also be applied for studying cardiac motion. For the cardiac application, the motion curve was reported to be more regular (Brant and Helms, 2007) than respiratory motion which requires more flexibility during free breathing. Therefore, using a Fourier-based curve model may be more appropriate for studying cardiac motion (Ionasec et al., 2010).

## 8. Acknowledgment

This work was partly funded by the European Commission, the Graduate School of Information Science in Health (GSISH), and the TUM Graduate School. We want to thank Martin von Siebenthal and especially Dirk Boye for providing the MR data. We would like to thank Tobias Schäffter for his help in acquiring MR data. Further, we want to thank Athanasios Karamalis, Diana Mateus, and Oliver Kutter for helpful discussions, as well as, the unknown reviewers for insightful comments.

## 9. Multimedia Material

We attach several videos to this submission for illustration. These include 2D+t breathing series, with the result of the manifold learning shown in the video, and a volume rendering of a constructed 4D ultrasound data set from the liver. Moreover, we show the gating results on a 2D+t MR sequence from ETH. Please open the file readme.html for an easy navigation.

## References

Arandjelovic, O., Cipolla, R., 2007. A manifold approach to face recognition from low quality video across illumination and pose using implicit super-resolution. ICCV .

Belkin, M., Niyogi, P., 2003. Laplacian eigenmaps for dimensionality reduction and data representation. Neural Comput. 15.

Brant, W., Helms, C., 2007. Fundamentals of diagnostic radiology. Lippincott Williams & Wilkins.

Colgan, R., McClelland, J., McQuaid, D., Evans, P., Hawkes, D., Brock, J., Landau, D., Webb, S., 2008. Planning lung radiotherapy using 4D CT data and a motion model. Physics in Medicine and Biology 53, 5815.

Crum, W., Tanner, C., Hawkes, D., 2005. Anisotropic multi-scale fluid registration: evaluation in magnetic resonance breast imaging. Physics in Medicine and Biology 50, 5153.

Flampouri, S., Jiang, S., Sharp, G., Wolfgang, J., Patel, A., Choi, N., 2006. Estimation of the delivered patient dose in lung IMRT treatment based on deformable registration of 4D-CT data and Monte Carlo simulations. Physics in Medicine and Biology 51, 2763.

Georg, M., Souvenir, R., Hope, A., Pless, R., 2008. Manifold learning for 4d ct reconstruction of the lung, in: Computer Vision and Pattern Recognition Workshops.

Gerber, S., Tasdizen, T., Joshi, S., Whitaker, R., 2009. On the manifold structure of the space of brain images, in: MICCAI.

Gerber, S., Tasdizen, T., Whitaker, R., 2007. Robust non-linear dimensionality reduction using successive 1-dimensional laplacian eigenmaps, in: Proceedings of the 24th international conference on Machine learning, ACM. pp. 281–288.

Hajnal, J., Hawkes, D., Hill, D. (Eds.), 2001. Medical Image Registration. CRC Press, Baton Rouge, Florida.

Hamm, J., Davatzikos, C., Verma, R., 2009. Efficient large deformation registration via geodesics on a learned manifold of images, in: MICCAI, pp. 680–687.

Ionasec, R., Wang, Y., Georgescu, B., Voigt, I., Navab, N., Comaniciu, D., 2010. Robust motion estimation using trajectory spectrum learning: Application to aortic and mitral valve modeling from 4D TEE, in: Computer Vision, 2009 IEEE 12th International Conference on, IEEE. pp. 1601–1608.

Jamieson, A., Giger, M., Drukker, K., Li, H., Yuan, Y., Bhooshan, N., 2010. Exploring nonlinear feature space dimension reduction and data representation in breast CADx with Laplacian eigenmaps and t-SNE. Medical physics 37, 339.

Karadayi, K., Hayashi, T., Kim, Y., 2006. Automatic image-based gating for 4d ultrasound, in: Engineering in Medicine and Biology Society.

King, A., Boubertakh, R., Rhode, K., Ma, Y., Chinchapatnam, P., Gao, G., Tangcharoen, T., Ginks, M., Cooklin, M., Gill, J., et al., 2009. A subject-specific technique for respiratory motion correction in image-guided cardiac catheterisation procedures. Medical Image Analysis 13, 419–431.

Kokiopoulou, E., Chen, J., Saad, Y., 2011. Trace optimization and eigenproblems in dimension reduction methods. Numerical Linear Algebra with Applications 18, 565–602.

Lee, C., Elgammal, A., 2007. Modeling view and posture manifolds for tracking, in: ICCV.

Li, G., Citrin, D., Camphausen, K., Mueller, B., Burman, C., Mychalczak, B., Miller, R., Song, Y., 2008. Advances in 4D medical imaging and 4D radiation therapy. Technology in cancer research & treatment 7, 67.

Martinez-Möller, A., Bundschuh, R., Riedel, M., Navab, N., Ziegler, S., Schwaiger, M., Nekolla, S., 2007. Comparison of respiratory sensors and its compliance for respiratory gating in emission tomography, in: Journal of Nuclear Medicine, p. 426.

McClelland, J., Blackall, J., Tarte, S., Chandler, A., Hughes, S., Ahmad, S., Landau, D., Hawkes, D., 2006. A continuous 4D motion model from multiple respiratory cycles for use in lung radiotherapy. Medical Physics 33, 3348.

Nguyen, T., Moseley, J., Dawson, L., Jaffray, D., Brock, K., 2009. Adapting liver motion models using a navigator channel technique. Medical physics 36, 1061.

Remmert, G., Biederer, J., Lohberger, F., Fabel, M., Hartmann, G., 2007. Four-dimensional magnetic resonance imaging for the determination of tumour movement and its evaluation using a dynamic porcine lung phantom. Physics in medicine and biology 52, N401.

Rijkhorst, E.J., Heanes, D., Odille, F., Hawkes, D., Barratt, D., 2010. Simulating dynamic ultrasound using mr-derived motion models to assess respiratory synchronisation for image-guided liver interventions, in: Navab, N., Jannin, P. (Eds.), Information Processing in Computer-Assisted Interventions. Springer Berlin / Heidelberg. volume 6135 of *Lecture Notes in Computer Science*, pp. 113–123.

Rohde, G.K., Wang, W., Peng, T., Murphy, R.F., 2008. Deformation-based nonlinear dimension reduction: Applications to nuclear morphometry, in: ISBI.

Rohlfing, T., Maurer, Jr., C.R., O'Dell, W.G., Zhong, J., 2001. Modeling liver motion and deformation during the respiratory cycle using intensity-based free-form registration of gated MR images, in: Mun, S.K. (Ed.), Medical Imaging: Visualization, Display, and Image-Guided Procedures, pp. 337–348.

Rohlfing, T., Maurer, Jr., C.R., O'Dell, W.G., Zhong, J., 2004. Modeling liver motion and deformation during the respiratory cycle using intensity-based free-form registration of gated MR images. Medical Physics 31, 427–432.

Roweis, S.T., Saul, L.K., 2000. Nonlinear Dimensionality Reduction by Locally Linear Embedding. Science 290, 2323–2326.

von Siebenthal, M., 2008. Analysis and Modelling of Respiratory Liver Motion using 4DMRI. Ph.D. thesis. Eidgenössische Technische Hochschule ETH Zürich.

von Siebenthal, M., Székely, G., Gamper, U., Boesiger, P., Lomax, A., Cattin, P., 2007. 4D MR imaging of respiratory organ motion and its variability. Physics in medicine and biology 52, 1547.

Sundar, H., Khamene, A., Yatziv, L., Xu, C., 2009. Automatic image-based cardiac and respiratory cycle synchronization and gating of image sequences, in: MICCAI, pp. 381–388.

Suzuki, K., Zhang, J., Xu, J., 2010. Massive-training artificial neural network coupled with laplacian-eigenfunction-based dimensionality reduction for computer-aided detection of polyps in ct colonography. Medical Imaging, IEEE Transactions on 29, 1907 –1917.

Tenenbaum, J., Silva, V., Langford, J., 2000. A global geometric framework for nonlinear dimensionality reduction. Science 290, 2319.

Timinger, H., Krueger, S., Dietmayer, K., Borgert, J., 2005. Motion compensated coronary interventional navigation by means of diaphragm tracking and elastic motion models. Physics in Medicine and Biology 50, 491.

Treece, G., Prager, R., Gee, A., Cash, C., Berman, L., 2002. Grey-scale gating for freehand 3D ultrasound, in: IEEE International Symposium on Biomedical Imaging, pp. 993–996.

Viola, P.A., 1995. Alignment by Maximization of Mutual Information. Ph.d. thesis. Massachusetts Institute of Technology.

Wachinger, C., Mateus, D., Keil, A., Navab, N., 2010a. Manifold Learning for Patient Position Detection in MRI, in: IEEE International Symposium on Biomedical Imaging, Rotterdam, The Netherlands.

Wachinger, C., Navab, N., 2009. Similarity Metrics and Efficient Optimization for Simultaneous Registration, in: IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR).

Wachinger, C., Navab, N., 2010. Manifold learning for multi-modal image registration, in: 11st British Machine Vision Conference (BMVC).

Wachinger, C., Wein, W., Navab, N., 2008. Registration strategies and similarity measures for three-dimensional ultrasound mosaicing. Academic Radiology 15, 1404–1415.

Wachinger, C., Yigitsoy, M., Navab, N., 2010b. Manifold Learning for Image-Based Breathing Gating with Application to 4D Ultrasound, in: International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI), Beijing, China.

de Winter, S., Hamers, R., Degertekin, M., Tanabe, K., Lemos, P., Serruys, P., Roelandt, J., Bruining, N., 2003. A novel retrospective gating method for intracoronary ultrasound images based on image properties, in: Computers in Cardiology.

Yigitsoy, M., Wachinger, C., Navab, N., 2011. Manifold learning for image-based breathing gating in mri, in: SPIE Medical Imaging, Orlando, Florida, USA.

Zhang, Q., Souvenir, R., Pless, R., 2006. On Manifold Structure of Cardiac MRI Data: Application to Segmentation. CVPR 1, 1092–1098.

Zhu, H., Oakeson, K., Friedman, M., 2003. Retrieval of cardiac phase from IVUS sequences, in: Proceedings of SPIE, p. 135.