Fully Automated Data-Driven Respiratory Signal Extraction from SPECT Images Using Laplacian Eigenmaps

James C. Sanders*, *Student Member, IEEE*, Philipp Ritt, Torsten Kuwert, A. Hans Vija, *Member, IEEE*, Andreas K. Maier, *Member, IEEE*

***Abstract*—We propose a data-driven method for extracting a respiratory surrogate signal from SPECT list-mode data. The approach is based on dimensionality reduction with Laplacian Eigenmaps. By setting a scale parameter adaptively and adding a series of post-processing steps to correct polarity and normalization between projections, we enable fully-automatic operation and deliver a respiratory surrogate signal for the entire SPECT acquisition. We validated the method using 67 patient scans from three acquisition types (myocardial perfusion, liver shunt diagnostic, lung inhalation/perfusion) and an Anzai pressure belt as a gold standard. The proposed method achieved a mean correlation against the Anzai of 0.81±0.17 (median 0.89). In a subsequent analysis, we characterize the performance of the method with respect to count rates and describe a predictor for identifying scans with insufficient statistics. To the best of our knowledge, this is the first large validation of a data-driven respiratory signal extraction method published thus far for SPECT, and our results compare well with those reported in the literature for such techniques applied to other modalities such as MR and PET.**

***Index Terms*—Data-driven, Laplacian eigenmaps, respiratory gating, single photon emission computed tomography (SPECT), surrogate signal.**

I. Introduction

SINGLE photon emission-computed tomographic (SPECT) imaging is vulnerable to blur and artifacts caused by respiratory motion occurring during respiratory cycles shorter than typical projection dwell times. One acquisition scenario where such issues arises is myocardial perfusion imaging, where physicians search for areas of the heart muscle receiving below average blood supply, which may be indicative of coronary artery disease or other cardiological pathologies. Motion of the diaphragm is known to induce corresponding movement on the heart. Previous work has determined that the predominant motion is along the inferior-superior axis, and mean amplitudes of this translation have been variously reported to be 16 mm, 18.1 mm, 9 mm, and 12.4 mm [1-4]. In simulation [5] and [6], as well as patient studies [2], it has been shown that respiratory motion of the heart can lead to hypoperfusion artifacts, particularly when the emission and transmission data used for attenuation correction are mismatched spatially.

The same motion of the diaphragm that is problematic for cardiac imaging also creates complications for hepatic imaging, where motion of the liver can approach 25 mm [7]. This is especially critical for therapy planning prior to Y-90 radioembolization, where an intraarterial injection of Tc-99m-MAA (macroaggregated albumin) into the hepatic artery is used to estimate dose to the lungs and healthy liver tissue incurred during therapy [8].

Another SPECT acquisition type where respiratory motion becomes problematic is lung ventilation/perfusion imaging. In this protocol, patients receive a first SPECT scan after inhaling a Tc-99m-labelled gas to assess airflow in the lungs. A second scan follows after injection of Tc-99m-MAA to examine blood perfusion through the lungs. The authors of [9] reported an increase in sensitivity for detection of ventilation and perfusion defects when the effect of respiratory motion was mitigated.

In order to overcome artifacts due to respiratory motion, a number of methods have been proposed that seek to subdivide the acquisition into time bins, or gates, during which motion is small. Individual gates may then be reconstructed and evaluated separately [10], or used to produce a single motion-corrected reconstruction [11-15]. Critical to each approach is a surrogate signal describing the respiratory state over time that can be used to drive the gating process.

One family of methods for deriving a surrogate signal is based on sensors placed directly on the patient. The most established of these is comprised of an elastic belt with a pressure sensor that measures the force exerted by the body

Manuscript received March 16, 2016; revised May 25th, 2016; accepted May 30th, 2016. Date of current version June 2nd, 2016. This work was supported in part by Siemens Molecular Imaging. *Asterisk indicates corresponding author.*

*J. C. Sanders and A. K. Maier are with the University of Erlangen-Nuremberg, Pattern Recognition Lab, Martensstr. 3, 91058 Erlangen, Germany (e-mail: james.sanders@uk-erlangen.de).

P. Ritt and T. Kuwert are with the University Hospital Erlangen, Clinic of Nuclear Medicine, Ulmenweg 18, 91054, Erlangen, Germany.

A. H. Vija is with Siemens Molecular Imaging, 2501 Barrington Rd., Hoffman Estates, IL 60169, USA.

surface due to respiration [16]. Other techniques utilizing spirometers [17] and piezoelectric belts [18] are available, as well as camera-based methods using marker tracking [19, 20] and range imaging [21] that offer more detailed, multi-dimensional information about respiratory motion.

Despite the utility of these approaches, they all necessitate that extra hardware be installed, calibrated, and synchronized with the SPECT system. Furthermore, the sensor-based methods, as well as the marker tracking technique, require that extra hardware be affixed to the patient, creating discomfort and an opportunity for user error if the requisite sensors are not used properly. A further issue in SPECT is that most systems in clinical use have two large detectors that rotate very close to the patient in order to maximize resolution. This proximity precludes the use of large external sensors and creates occlusion problems for techniques involving cameras. For these reasons, *data-driven* approaches are desirable that estimate a surrogate signal from the acquired data itself.

Data-driven approaches based on tracking count rate variations [22], image center of mass (CoM) [23, 24], and spectral analysis [25] have been previously investigated for positron emission tomography (PET). However, whereas PET employs a ring detector, SPECT acquisitions are composed of sequential discrete projections. Due to attenuation and background, focal regions contributing to count rate variations and CoM movement may not be visible from each angle. One CoM tracking approach was applied to cardiac SPECT with a three-headed system in [26], but the method was found to be unstable at clinical count levels. Another CoM implementation in [27] proved more successful, but for a specialized SPECT system capable of acquiring 19 projections simultaneously.

Similar to traditional SPECT, cone-beam CT (CBCT) relies on slowly acquired discrete projections and, hence, faces similar problems. For this modality, methods based on diaphragm tracking [28, 29] have shown promise. SPECT projections have a lower resolution and much higher noise than their CBCT counterparts, however, limiting the applicability of this approach.

Despite the fact that temporal variations due to respiration in image characteristics such as count rate and CoM location may not enable a surrogate signal estimate *individually*, it is nevertheless safe to assume that some type of variation exists in each projection. A successful data-driven technique in SPECT should therefore be sensitive to *all* variation in the image and be capable of "summarizing" it in a surrogate signal. This problem statement corresponds to the goal of dimensionality reduction (DR) methods, which seek to map high-dimensional data points into a low-dimensional space while preserving their structure and variations between them. DR methods are useful for exploratory data analysis [30] as well as medical image processing [31].

The most well-known DR method is principal component analysis (PCA), which computes a linear mapping from a high- to a low-dimensional space based on an analysis of the input data's covariance matrix. It has been successfully applied to PET in [32] for the purpose of respiratory surrogate signal estimation, as well as in a specialized form to CBCT in [33]. An alternative to PCA is Laplacian Eigenmaps (LE) [34], a nonlinear DR technique adapted to magnetic resonance imaging (MR) and ultrasound in [35, 36]. LE was compared to PCA and other methods for PET in [37] and found to deliver acceptable but slightly poorer results. However, in preliminary work by our group for SPECT [38] we compared PCA and LE in a phantom experiment and on cardiac patient data and found that LE outperformed PCA in both cases, particularly at high noise levels. Nevertheless, the goal of [38] was to compare the two methods, and a number of steps still remain before LE can be deemed fit for clinical practice.

The aim of this paper is to make an effort to resolve these remaining issues. Specifically, we detail a solution for the issue of normalization and polarity correction of LE-based surrogate signal estimates, which may be arbitrarily scaled and flipped for each projection view. Furthermore, we perform a more extensive validation and characterization of the method using patient data from three different tracer/scan protocols where surrogate signals obtained using an Anzai belt serve as a gold standard. We also propose a predictor of surrogate quality that can be used to identify cases where the method is likely to fail due to insufficient photon counts.

II. Methods

A. *Laplacian Eigenmaps*

The goal of a DR method is to map data natively represented in a high-dimensional space to a low-dimensional space while preserving its intrinsic structure. In our case, the input data consists of T vectorized projection images $\mathbf{x}_i \in \mathbb{R}^M$, where i denotes the i-th time bin and M is the number of pixels in the image. The desired output is the respiratory surrogate signal with one $\mathbf{y}_i \in \mathbb{R}^N$ for each time bin, where $N \ll M$.

Laplacian Eigenmaps accomplishes DR by finding the matrix $\mathbf{Y} = [\mathbf{y}_1, \mathbf{y}_2 \ldots \mathbf{y}_i \ldots \mathbf{y}_T]$ that minimizes the following function:

$$\sum_{i,j} \left\| \left(\mathbf{y}_{*,i} - \mathbf{y}_{*,j} \right) \right\|_2^2 \mathrm{w}_{i,j}, \tag{1}$$

where the subscripts $*,i$ and $*,j$ denote column vectors of $\mathbf{Y}$ corresponding to N-dimensional output points at time bins i and j, respectively. $\mathrm{w}_{i,j}$ is an element of the data's affinity matrix $\mathbf{W}$ and is a weight inversely proportional to the distance between points $\mathbf{x}_i$ and $\mathbf{x}_j$. Determining a series of $\mathbf{y}_{*,i}$'s that minimizes (1) ensures that points that are close in the high dimensional space will remain close in the low dimensional space as well.

The choice of distance metric originally proposed in [34] is motivated by thermodynamics and based on the heat kernel:

$$\mathrm{w}_{i,j} = \begin{cases} e^{-\frac{\left\| \mathbf{x}_i - \mathbf{x}_j \right\|_2^2}{\alpha}}, \text{for } \left\| \mathbf{x}_i - \mathbf{x}_j \right\|_2^2 < \epsilon \\ 0, \text{ otherwise,} \end{cases} \tag{2}$$

where ϵ is a threshold value and α a constant scale parameter. An alternative method also proposed in [34] utilizes the same exponential kernel, but employs a k-nearest-neighbor approach instead of thresholding, where only the k largest $\mathrm{w}_{i,j}$'s for

each $\mathbf{x}_i$ receive non-trivial values, with all others being set to zero. Thresholding and *k*-nearest-neighbor rules serve to pare elements of $\mathbf{W}$ and have the effect of imposing an aspect of locality on (1), allowing only points in a given $\mathbf{x}_i$'s local neighborhood to affect the corresponding $\mathbf{y}_i$. For reasons discussed below, we found that, in our application case, judicious selection of α allowed both of these rules to be discarded, eliminating the need to set these parameters while still providing robust performance.

Equation (1) can be optimized with the help of the graph Laplacian $\mathbf{L} = \mathbf{D} - \mathbf{W}$, where $\mathbf{D}$ is a diagonal matrix whose elements represent the degree of connectivity of each high-dimensional point: $\mathbf{D}_{i,i} = \sum_j \mathrm{w}_{i,j}$. As shown in [34], the optimization can be reduced to a generalized eigenvalue problem:

$$\mathbf{L}\mathbf{y}_{n,*} = \lambda \mathbf{D}\mathbf{y}_{n,*}, \tag{3}$$

where the $\mathbf{y}_{n,*}$ is the eigenvector corresponding to the *n*-th non-zero eigenvalue λ and represents the *n*-th row of $\mathbf{Y}$. To obtain an *N*-dimensional representation of the original data, one must simply extract the *N* eigenvectors corresponding to the *N* smallest non-zero eigenvalues.

B. *Determining the Scale Parameter α*

In preliminary investigations, *k*, ϵ, and α proved difficult to manually set, providing highly inconsistent results over our patient collective for a given value. In an effort to determine α adaptively, we relied on work by Coifman *et al*. on Diffusion Maps, a DR approach related to LE [39]. The authors noted that for small α, all non-diagonal elements of $\mathbf{W}$ go to zero, and for large α, all elements approach unity. Neither one of these cases is capable of yielding a useful mapping, and α must therefore be chosen from the transition region between these two asymptotes where it has a magnitude relevant to the distances between points contained in $\mathbf{W}$. To accomplish this, we dispense with thresholding and nearest-neighbor rules and simply set α equal to the mean over all squared Euclidean distances between time bins:

$$\alpha = \frac{1}{T^2}\sum_{i,j}\left\|\mathbf{x}_i - \mathbf{x}_j\right\|_2^2 \tag{4}$$

The properties of this method are illustrated in Fig. 1, where $Q(\alpha) = \sum_{i,j} \mathrm{w}_{i,j}(\alpha)$ is plotted against α for one set of cardiac patient data. Each curve represents a different noise level generated via binomial subsampling of the counts in the original projection data. This is accomplished by performing *Z* Bernoulli trials with probability *f* at each pixel, where *Z* and *f* are the number of counts in the original data and the subsampling fraction, respectively. It can be seen that for all cases, the lower asymptote of $Q(\alpha)$ is at T, the number of input points, and the upper asymptote lies correspondingly at T^2. As image noise increases (indicated in the legend by reduced percentage of total counts), the transition region for a viable α is shifted, partially explaining the difficulty in choosing a universal value for all acquisitions *a priori*. However, as shown by the vertical lines, setting the scale parameter according to (4) tracks the transition region.

It should be noted that the local aspect of the LE method is thus no longer explicitly enforced in this implementation. However, the use of the exponential kernel in tandem with an appropriately chosen α causes $\mathrm{w}_{i,j}$ to naturally decay to zero for distant pairs of input points, mimicking the paring behavior of ϵ-thresholding and *k*-nearest-neighbor rules.

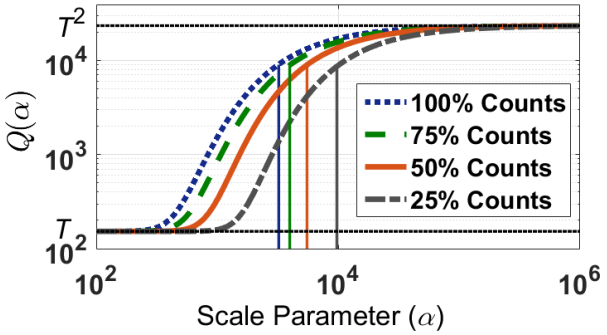

Fig. 1. The sum of all elements in the affinity matrix $Q(\alpha) = \sum_{i,j} \mathrm{w}_{i,j}(\alpha)$ as a function of scale parameter α is shown (log-log plot). Each curve represents data from a particular noise level of the same cardiac acquisition. The vertical lines indicate the value of α chosen adaptively using (4) at each noise level.

C. *Data and Preprocessing*

In our implementation of the LE method for SPECT, we begin by temporally binning incoming list-mode data into 200 ms frames. Detected photons are spatially binned into 256x256 pixel matrices with 2.4 mm isotropic pixels. During preliminary work, we observed that the quality of estimated surrogate signals was best when all detected photons were used, rather than just those from the photopeak. Thus, we bin all acquired counts, regardless of energy.

Following binning, frames are smoothed spatially with a 32×32 pixel rectangular window and temporally at each pixel with a 2 time bin moving average window, thus accomplishing 3D smoothing similar to the 4D smoothing in [25]. Temporal smoothing with a moving average filter of even length effectively shifts the centers of the surrogate signal's time bins by a uniform amount. This was not important for the validation study presented here but should be taken into account for applications requiring absolute synchronization with the SPECT data. Our SPECT/CT system is outfitted with two detectors with an adjustable angular separation. Data from both detectors is acquired simultaneously at each rotation increment in a step-and-shoot mode, thus yielding two projections for each camera stop p. For this reason, frames from both detectors are concatenated to create a 256×512×T dataset for each stop. After vectorization of each time bin, input data to the LE method takes the form $\mathbf{X}_p \in \mathbb{R}^{2^{17}\times T}$, where each column vector $\mathbf{x}_{*,i}$ represents an input point at time bin i.

As a final step for cases where the camera sensitivity is uniform across the field of view (e.g. parallel-hole collimation), data in $\mathbf{X}_p$ is normalized such that the sum of each column is equal. This is done under the assumption that, given uniform camera sensitivity and motion that is confined to the field of view, variations in total counts will be due to

statistical fluctuations rather than motion. As the LE method responds to all variations in the data over time, this normalization serves to reduce sensitivity to those purely due to the randomness of counting statistics. This assumption holds for, e.g., liver and lung scanning with a parallel-hole collimator, but not in cardiac scanning with astigmatic collimation, which includes a complex pattern of collimator bore angulations to increase sensitivity in the center of the field of view [40]. In this case, no input normalization is performed, as spatially varying sensitivity results in count rates that are a function of the motion itself and hence capable of contributing to the estimate.

Acquired data from each of the camera stops is disjoint in time and is thus treated as a series of independent measurements to yield one $\mathbf{Y}_p \in \mathbb{R}^{N\times T}$ at each camera stop p. In this study, we restricted the dimensionality of the LE output to one (i.e. $N = 1$), allowing $\mathbf{Y}_p$ to be represented as a vector $\mathbf{y}_p = [\mathrm{y}_1, \ldots \mathrm{y}_T]$, where each element corresponds to the estimated respiratory amplitude of the patient at the corresponding time bin.

D. *Post Processing*

A number of post processing steps are required to assemble a complete surrogate signal for the entire acquisition from the individual $\mathbf{y}_p$'s. These are illustrated in Fig. 2 and include the following:

1) *Smoothing:* First off, at each camera stop, the output signal of the LE estimation step is smoothed using a Savitzky-Golay filter with a span and order of nine and three, respectively. This choice of filtering was motivated by the filter's tendency to preserve the magnitude of peaks in a noisy signal [41].
2) *Projection-Wise Polarity Correction:* Each smoothed $\mathbf{y}_p$ is the output of an eigenvalue decomposition and is hence scaled by an arbitrary value. All $\mathbf{y}_p$'s must be polarity corrected and normalized such that they are consistent in sign and magnitude with others across the acquisition. To attain consistent polarity, two types of physical references computed from $\mathbf{X}_p$ were used depending on collimation: For acquisitions with spatially varying camera sensitivity (astigmatic collimation), the total detected counts in each time bin served as a reference, similar to [22]. For parallel-hole collimation, the craniocaudal component of the image CoM was used in a similar way to [26].[1] While too noisy to provide a respiratory signal themselves, these references (shown schematically as the orange dashed line in Fig. 2) serve as an anchor to enforce a consistent polarity throughout the acquisition. Each $\mathbf{y}_p$ is compared to the reference via Pearson's correlation and flipped if the correlation is negative. Note that no assumption is made at this point as to whether end-expiration corresponds to high or low values.
3) *Baseline Correction:* After projection-wise polarity correction, each $\mathbf{y}_p$ is subjected to a robust baseline correction scheme. First, outliers (defined here as estimated surrogate signal values more than two standard deviations away from the mean) are stripped from $\mathbf{y}_p$ to yield an outlier-free set $\mathrm{Y}_{\mathrm{olf}} = \{\mathrm{y}_i \mid |\mathrm{y}_i - \bar{\mathbf{y}}_p| \leq 2\sigma_{\mathbf{y}_p}\ \forall\ i = 1 \ldots T\}$, where $\bar{\mathbf{y}}_p$ and $\sigma_{\mathbf{y}_p}$ are the mean and standard deviation over the elements of $\mathbf{y}_p$, respectively. Then, the values of $\mathrm{Y}_{\mathrm{olf}}$ are sorted in ascending order, and the first decile is extracted to yield a subset $\mathrm{Y}_{10\%} \subseteq \mathrm{Y}_{\mathrm{olf}}$ containing the lowest values of the signal. The mean of this subset $\bar{\mathrm{Y}}_{10\%}$ is then subtracted from the original estimate to yield $\mathbf{y}'_p$ with elements $\mathrm{y}'_{p,i} = \mathrm{y}_{p,i} - \bar{\mathrm{Y}}_{10\%}\ \forall\ i \in \{1 \ldots T\}$, thus encouraging a uniform baseline across all camera stops.
4) *Normalization:* To ensure the curves have a uniform magnitude as well, the outlier rule is again applied to $\mathbf{y}'_p$ to obtain $\mathrm{Y}'_{\mathrm{olf}} = \{\mathrm{y}'_i \mid |\mathrm{y}'_i - \bar{\mathbf{y}}'_p| \leq 2\sigma_{\mathbf{y}'_p}\ \forall\ i = 1 \ldots T\}$, where $\bar{\mathbf{y}}'_p$ and $\sigma_{\mathbf{y}'_p}$ are again the mean and standard deviation over the elements of $\mathbf{y}'_p$, respectively. Each $\mathbf{y}'_p$ is then normalized by the standard deviation σ_{olf} over $\mathrm{Y}'_{\mathrm{olf}}$ to yield $\mathbf{y}''_p = \mathbf{y}'_p/\sqrt{2}\sigma_{\mathrm{olf}}$. Basing the normalization off of statistics from percentiles and outlier-free portions of the data provides more consistent performance in the face of noisy data and irregular breathing. If the respiratory waveform is approximately sinusoidal, typical respiratory cycles will have a peak-to-peak amplitude of two and a baseline of roughly zero.
5) *Concatenation:* To assemble a global respiratory surrogate signal for the entire acquisition, all $\mathbf{y}''_p$'s are then concatenated to yield a vector $\mathbf{y} \in \mathbb{R}^{PT\times 1}$.
6) *Global Polarity Correction (not shown in figure):* The projection-wise polarity correction and normalization has ensured that each $\mathbf{y}''_p$ is consistent to the others, but the global polarity may still be inverted. To correct this, $\mathbf{y}$ is subjected to a check using knowledge that, as most time is spent in end expiration, histograms of respiratory amplitude tend to be bottom heavy [42]. We enforce this constraint by constructing a histogram of $\mathbf{y}$ and inverting the signal if it is negatively skewed. The resulting vector represents a surrogate signal for respiratory amplitude throughout the entire acquisition as a function of time.

E. *Implementation and Runtime*

The binning routine was implemented in C# as a compilable program. All other processing steps were implemented in-house using Matlab (The MathWorks, Inc., Natick, MA, USA). Binning and processing were performed on a workstation laptop with a 2.6 GHz Intel i7-3720QM processor. Run time for binning increased with the total number of counts in each acquisition as well as number of camera stops and ranged from roughly 70 to 140 s. Time required to process the binned data increased with number of views and dwell time and ranged from ca. 240 to 400 s.

[1]Bruyant, King, and Pretorius apply a threshold prior to CoM calculation that is optimized for cardiac imaging. As our method is more general, we do not threshold the data.

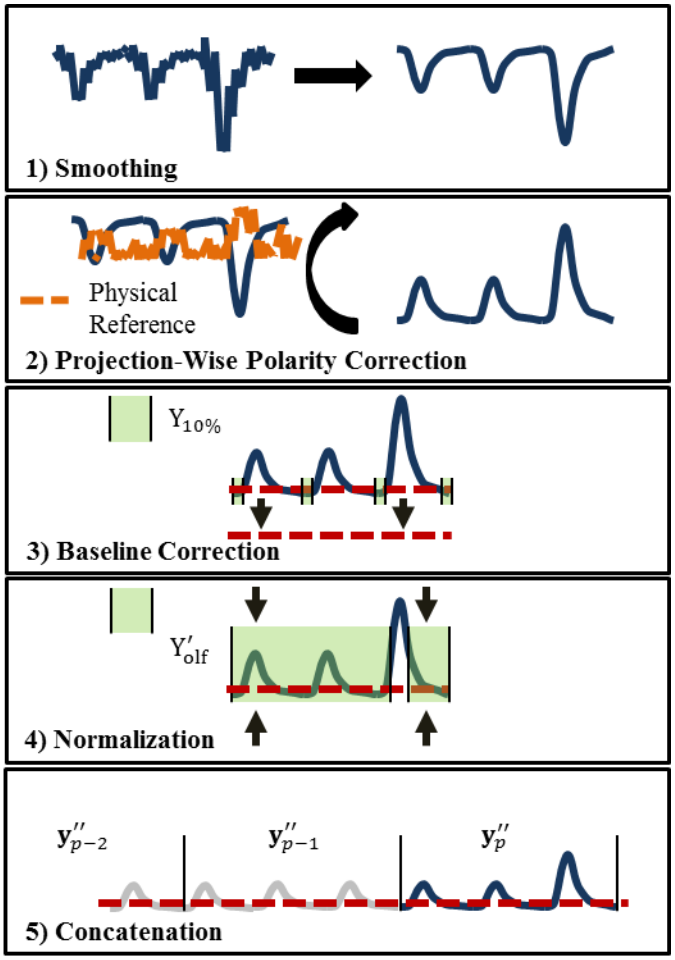

Fig. 2. Schematic description of post processing steps. Final global polarity correction not shown.

F. *Patient Validation*

In a validation study approved by our hospital's institutional review board, data was collected at the University of Erlangen-Nuremberg Clinic of Nuclear Medicine from patients undergoing routine exams between May, 2014 and August, 2015. For each acquisition, following granting of informed consent, an Anzai AZ-733V belt with pressure sensor (Anzai Medical Co., Ltd., Tokyo, Japan) was affixed to the patient. After confirmation that a proper respiratory trace was visible (and adjustment of the sensor when needed), the Anzai acquisition was commenced, followed shortly thereafter by the planned SPECT scan. For this study, data from three types of acquisitions were collected: myocardial perfusion (stress and rest), shunt diagnostic scans for radioembolization planning, and lung inhalation/perfusion. All acquisitions were carried out using a Siemens Symbia T2 SPECT/CT system (Siemens Molecular Imaging, Inc., Hoffman Estates, USA). Table I summarizes the acquisition parameters involved for each scan type.

Following data acquisition, the Anzai signal, which is acquired with a sampling frequency of 40 Hz, was resampled via b-spline interpolation to the temporal resolution of the LE estimate (5 Hz). To eliminate low-frequency baseline shifts such as those observed in [43] and [44], we detrended the Anzai signal by subtracting a moving average with a window width of one minute. To facilitate visual comparison, the baseline correction and normalization steps applied to each camera stop of the LE signal were then applied to the entire Anzai signal.

As there was no way to synchronize the Anzai sensor and our SPECT system via hardware, we instead synchronized the Anzai signal to the LE estimate by searching for the maximum cross-correlation between the two and shifting accordingly. The sync was further refined at sub-time bin resolution by searching within a ±1 bin neighborhood in 20 ms increments via b-spline interpolation. Synchronization in this way is insensitive to phase shifts or time delays that are constant throughout the entire acquisition but may be error prone if delays are not constant.

To compare the Anzai gold standard signal with the LE estimate, we computed Pearson's correlation coefficient between the two signals for each dataset and termed it ρ_{Anzai}. As this metric assesses the degree of linear dependence between the two signals, it should be a good predictor of how similar gating results would be based on Anzai and LE using prevailing amplitude-based gating methods [45, 46]. In addition to correlation, we manually compared each projection's surrogate signal estimate with the corresponding excerpt from the Anzai signal and assessed whether or not the polarity correction failed. The same was done at the global, acquisition-wide scale.

TABLE I
ACQUISITION DETAILS FOR PATIENT VALIDATION

Scan Type	Tracer	Injected Activity at Acquisition Time (MBq)	Collimator	Detector Angular Separation	Camera Stops (Dwell Time)	Number of Datasets
Cardiac (Stress/Rest)	Tc-99m-Sestamibi	237 ± 22/ 665 ± 56	SMARTZOOM (Astigmatic)	104°	17(30 s)	14 / 13
Shunt diagnostic for liver radioembolization planning	Tc-99m-MAA[a]	122 ± 37	LEHR[b]	180°	60(15 s)[c]	26
Lung (Inhalation/Perfusion)	Tc-99m-DTPA[d]/ -MAA[a]	760 ± 117[e]/ 158 ± 31	LEHR[b]	180°	60(25 s)/ 60(15 s)	7 / 7

[a]Macroaggregated albumin
[b]Low Energy-High Resolution
[c]Due to a low injected dose, one patient was acquired with a dwell time of 30 seconds.
[d]Diethylene-triamine-pentaacetate
[e]Lung inhalation activity values represent total amount prepared, decay corrected to acquisition time. Due to high variability in the inhalation process, the actual amount of radioactivity in the lungs may be much less.

G. *Probing Effect of Count Rate*

SPECT imaging typically operates in a low count regime, and it is thus critical that any algorithm applied to it be assessed for robustness against noise. As we are seeking to estimate a function over time, we recorded the average detected counts per second in each of the patient acquisitions to serve as a measure of noise level in the data and enable further analysis.

Although our patient collective provides a wide range of count rates, it is difficult to separate variations in results due strictly to noise from those due to inter-patient differences. To overcome this, we performed a separate analysis on three acquisitions with above average ρ_{Anzai} (one from each scan type) and progressively increased noise artificially via binomial subsampling on the full-count data. At each increment we then computed correlations between the reduced-count LE estimate and the Anzai signal. We repeated this analysis on one cardiac scan with low ρ_{Anzai} as well.

III. Results

A. *Surrogate Signal Estimate*

Fig. 3(a) shows the estimated and Anzai surrogate signals from the first three projections of a representative cardiac rest acquisition. Black bars denote the time spent during detector positioning. It can be seen that the LE estimate consistently represents all of the respiratory cycles present in the data, albeit with slight underestimations at cycle minima. The normalization provides a consistent baseline and peak-to-peak amplitude across data from each of the camera stops.

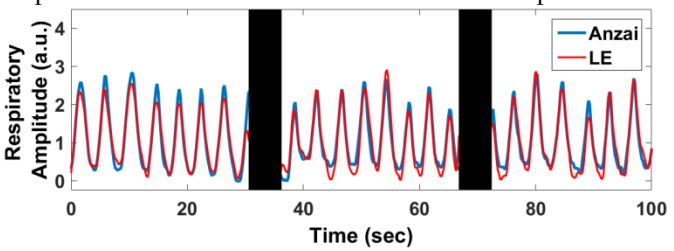

(a)

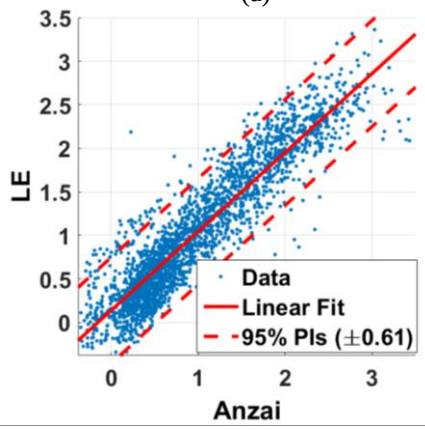

(b)

Fig. 3. Estimated and Anzai surrogate signals plotted against each other for three projections of a representative cardiac rest scan (a). Respiratory cycles are well represented, and normalization provides consistency between camera stops. Black bars cover time during detector rotation between camera stops. The scatter plot in (b) shows results for each time bin over entire acquisition.

The correlation of the LE estimate to the Anzai signal for this scan was 0.91, and data from the entire 8.5 minute scan is shown in the scatter plot in Fig. 3(b), where each point represents data from one 200 ms time bin. There is a clear linear relationship with a slope close to unity and 95% prediction intervals roughly 25% of the total peak-to-peak amplitude, according to a linear fit.

Table II shows mean ρ_{Anzai} ± standard deviation across all acquisitions for each scan group. Median ρ_{Anzai} is also reported. All correlations were significant with a p-value<0.001. Cardiac and lung subgroups are listed separately due to their strongly differing count levels. All of the acquisition types yielded median correlations ≥0.85 except for the lung inhalation scans, where the mean across the seven patients was lower at 0.54. According to Wilcoxon rank-sum testing, median correlations for all patient groups were statistically equal at the 10% confidence level with the exception of the lung inhalation scans, which were significantly lower than those for the stress, rest, and shunt diagnostic scans.

TABLE II
CORRELATIONS BETWEEN ANZAI SIGNAL AND LAPLACIAN EIGENMAPS ESTIMATE AND COUNT RATES ACROSS PATIENT COLLECTIVE

Scan Type	Pearson's ρ_{Anzai}[a] ± SD	Median ρ_{Anzai}	kCounts/s ± SD
Cardiac	0.87 ± 0.10	0.90	50.3 ± 31.3
Stress	0.84 ± 0.13	0.89	21.4 ± 2.9
Rest	0.90 ± 0.03	0.90	81.4 ± 12.7
Shunt diagnostic	0.75 ± 0.36	0.85	15.7 ± 5.5
Lung	0.68 ± 0.24	0.71	14.5 ± 11.1
Inhalation	0.54 ± 0.24	0.54	4.4 ± 2.9
Perfusion	0.83 ± 0.12	0.86	24.7 ± 5.5

[a]Correlations for all acquisitions were significant with p-value < 0.001.

The divergence between the median and mean ρ_{Anzai}, as well as the relatively high standard deviations, for cardiac stress and shunt diagnostic scans are indicative of outliers within each collective having disproportionately poor results. These groups had one and two patients each, respectively, with ρ_{Anzai} more than three standard deviations lower than the mean within the corresponding group.

B. *Polarity Correction*

Table III shows the absolute number of failures and failure rates for the projection-wise and global polarity correction. For cardiac and shunt diagnostic scans, projection-wise correction only failed in ca. 1% of cases. These errors were confined to three acquisitions, two of which were outliers with poor overall correlation. The projection-wise failure rate for the lung patients was higher at 6.3% and spread across eight of the 14 acquisitions. In each of these cases, either the surrogate estimate or the physical reference was too noisy to provide a meaningful correlation with each other. Global polarity correction failed in a single patient whose natural breathing violated the bottom heavy constraint.

TABLE III
NUMBER OF PROJECTIONS AND ACQUISITIONS IN WHICH POLARITY CORRECTION FAILED. PERCENTAGE OF TOTAL PROJECTIONS/ACQUISITIONS AFFECTED SHOWN IN PARENTHESES

Errors	Cardiac	Shunt Diagnostic	Lung
Projection-Wise	4(0.9%)	16(1.0%)	53(6.3%)
Global	0(0%)	1(3.9%)	0(0%)

C. *Count Rates*

Fig. 4 shows ρ_{Anzai} for the cardiac stress/rest (a), shunt diagnostic (b), and the lung inhalation/perfusion (c) scans plotted against their respective count rates. It can be seen that, with the exception of the shunt diagnostic acquisitions, scans with poor surrogate signal estimates tend to be in the low count regime for their respective protocols. However, despite this visual impression, mean count rates for each scan group (reported in Table II) were not significantly correlated with ρ_{Anzai}. For subgroups, only results from the lung inhalation acquisitions were significantly correlated with count rate (Pearson's ρ=0.86, p-value=0.024).

Results of the second count rate analysis are presented in Fig. 5. For the three above average scans, the structure of the correlations as a function of count rate is well-fit by a logistic function. Points from each dataset show a plateau at higher count levels and a breakdown region as count rate is reduced. The poor quality cardiac stress scan, however, exhibits no plateau, and ρ_{Anzai} decreases even at mild levels of subsampling.

IV. DISCUSSION

A. *Effect of Irregular Respiration*

In general, the post-processing steps were able to provide a common baseline, polarity, and magnitude for signals where breathing is either uniform throughout the acquisition or slowly varying. The baseline correction scheme makes no assumptions about whether each $\mathbf{y}_p$ is polarity-aligned with end-expiration at high or low values (both cases were observed in our collective). However, it does work under the assumption that for the case where end-expiration is at high values, any large variations in the magnitude of respiratory peaks will be of limited temporal duration. If this is satisfied, transients will be ignored by discarding outliers from Y_{olf}, and the influence of small variations in respiratory cycle magnitude will be mitigated by taking the mean of the lowest decile rather than simply taking, e.g., the minimum.

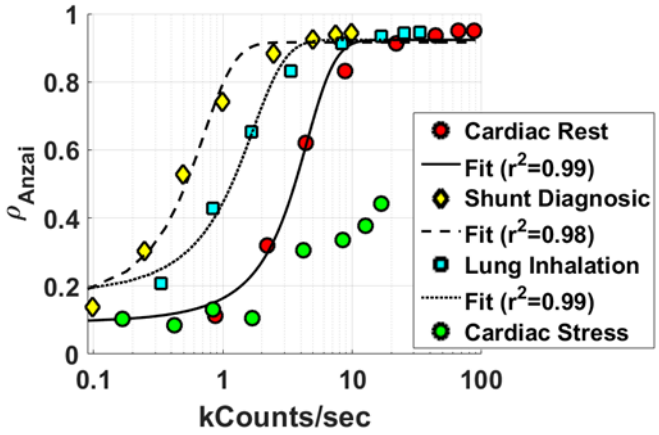

Fig. 5. Correlation to Anzai plotted against count rates derived from above average cardiac rest, shunt diagnostic, and lung inhalation scans. Lower count rate data generated using binomial subsampling of full count data. Fits of logistic function to each set of points shown in black. Correlations for poor quality cardiac stress signal shown in green circles.

Fig. 6(a) shows three projections from a more difficult scenario where a cardiac patient exhibited periodic breathing, whereby the amplitude of the respiratory cycle appears modulated by a low frequency component. This pathology is related to cardiac disease and is not uncommon for patients receiving myocardial perfusion scans [47]. The baseline between projections is not as consistent as Fig. 3, and the magnitude of peaks is frequently underestimated. It can be seen particularly from the larger spread in Fig. 6(b) that the modulation of the breathing cycle degrades the results of the surrogate signal estimation somewhat (ρ_{Anzai} was below average at 0.81). Nevertheless, the overall structure of the signal was well captured.

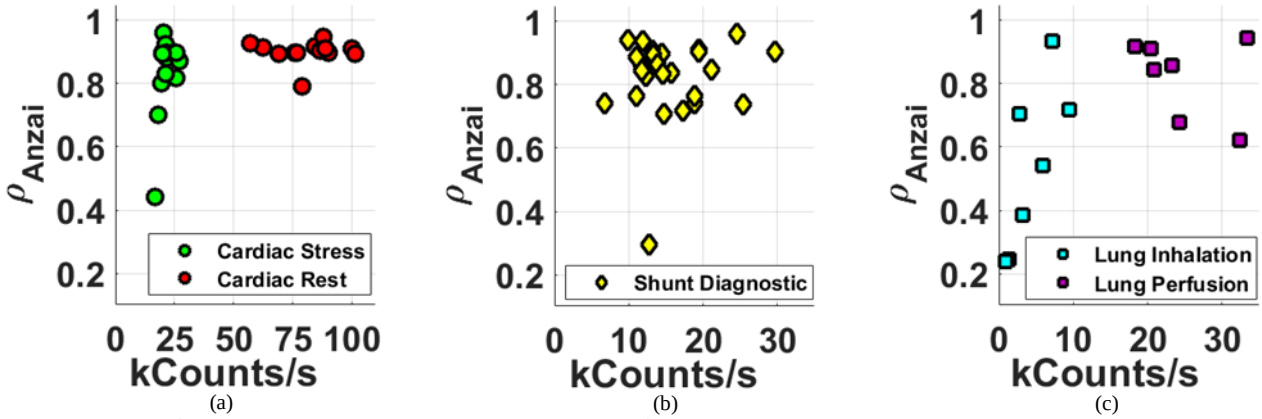

(a) (b) (c)

Fig. 4. Correlation of LE to Anzai plotted against count rates in the data for cardiac (a), shunt diagnostic (b), and lung (c) scan types. For clarity, the shunt diagnostic scan with inverted polarity (ρ_{Anzai}=-0.92, 19.5 kCounts/s) is not shown.

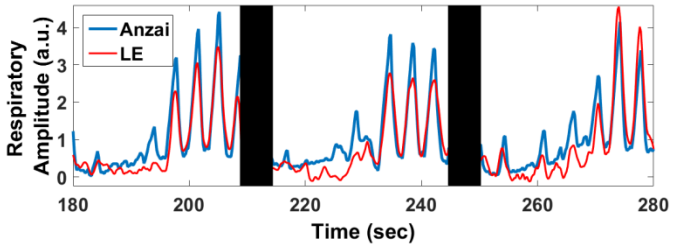

(a)

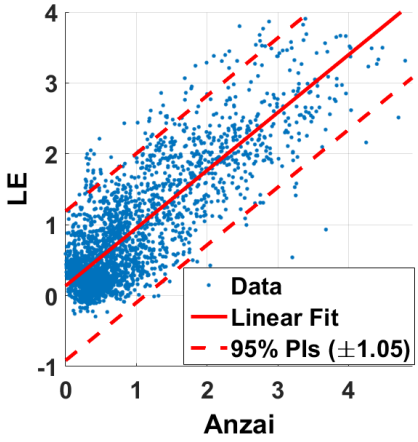

(b)

Fig. 6. Estimated and Anzai surrogate signals plotted against each other for three projections of a below average cardiac rest scan (a). High-magnitude cycles are well recovered, but some low-magnitude cycles are lost in noise. Normalization shows some inconsistency in magnitude due to irregular breathing pattern that is reflected in the spread of the scatter plot in (b). Black bars cover time during detector rotation between camera stops.

It is possible for abnormal respiratory waveforms to induce a failure on our global polarity correction as well. This occurred in one shunt diagnostic scan, where the patient spent more time at full-inspiration than full-expiration/beginning-inspiration, violating the bottom heavy constraint enforced by the method. The resulting ρ_{Anzai} was -0.92, indicating a high quality estimate that was incorrectly inverted during post processing. If the polarity had been correct, the results for the shunt diagnostic patients would have been improved, with a ρ_{Anzai} of 0.83±0.13 and a median of 0.86. Furthermore, if the surrogate signal is being used to gate data as an input to a motion-correction method, the global polarity may be inconsequential as long as the relative separation between respiratory gates is maintained.

Normalization is also affected by irregular breathing patterns, where transients of extended duration with amplitudes greatly exceeding that of other respiratory cycles can be problematic. In such cases, LE maps samples highly non-linearly into the low-dimensional space, and places outliers disproportionately far away from other points. For short duration transients, such as coughs, outliers will be excluded from Y'_{olf}, and the normalization step will accurately match the amplitude of normal respiratory cycles to those in other projections. However, for longer duration transients, such as rogue deep inhalations, the normalization might fail, leaving one peak matching its true value and effectively suppressing the others. Fig. 7 depicts an example of this behavior from a lung inhalation scan.

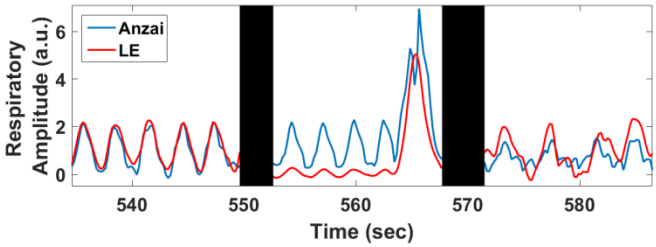

Fig. 7. Respiratory surrogates from three projections of a lung inhalation scan. Failure of normalization due to rogue deep inhalation is visible in central projection. Black bars indicate time spent during detector positioning, the duration of which may vary due to autocontouring.

The nonlinearity of LE is also visible in patients with regular breathing, as indicated by the slight concavity of the scattered points in Fig. 3(b). We observed this characteristic in several patients, but in such cases, a more fundamental disagreement between internal motion magnitudes and those appearing at the Anzai sensor cannot be ruled out.

B. *Effect of Count Rate*

As noted in the methods section, during preliminary work we observed a benefit from using all acquired counts rather than just those in the photopeak. The resulting frames are low contrast and useless for reconstruction, but the increase in counts – even if coming from scattered photons – stabilized our surrogate signal estimates. Table IV provides a trio of examples illustrating the effect of using all counts. For datasets with high count rates to begin with, there is little or no benefit, as can be seen in the first row. However, for patients with lower count rates, ρ_{Anzai} is higher when using all photons (fourth column) than when using either the non-photopeak or only photopeak counts alone (second and third column, respectively). This indicates that both scattered and non-scattered photons carry some information related to respiratory motion that is useful for the LE algorithm.

TABLE IV
EXEMPLARY RESULTS ILLUSTRATING BENEFIT OF UTILIZING ALL DETECTED PHOTON COUNTS

Scan Type	ρ_{Anzai} Non-Photopeak	ρ_{Anzai} Photopeak Only	ρ_{Anzai} All Counts	kCounts/s Non-Photopeak (Photopeak)
Cardiac Rest	0.91	0.89	0.91	74.0(25.7)
Cardiac Stress	0.52	0.17	0.70	14.1(3.8)
Cardiac Stress	0.75	0.60	0.80	14.1(5.2)

The increase in noise that comes with lower counts should somehow degrade algorithm performance, yet only the results of the lung inhalation scans were significantly correlated with count rate. This makes sense in light of the logistic characteristic for the high quality scans in Fig. 5, which we also observed in a phantom experiment in [38]. This behavior indicates that there is a range of count rates in the plateau region that deliver similar results and a breakdown point at which performance degrades rapidly. This nonlinear relationship would not be captured by a Pearson correlation and points to the fact that, with the exception of the lung

inhalation scans and the outlier cardiac scan (cardiac stress scan in Fig. 5), LE was operating in the plateau regime for the majority of datasets available.

Although it is possible to explain some of the variation in results using count rates and the logistic characteristic, the presence of datasets with poor results in the shunt diagnostic and lung perfusion groups despite average count rates indicates that, at the very least, the location of the breakdown point is patient-specific. Furthermore, other factors such as the regularity of the breathing cycle discussed above certainly also play a role. Ultimately, more patient data, particularly for lung scans, are needed to thoroughly evaluate the robustness of the algorithm.

C. *Towards a Performance Predictor*

One challenge in applying an automated data-driven method in clinical practice is providing the user with an indicator as to whether or not it has failed. We can take advantage of the logistic breakdown behavior shown in Fig. 5 to design a predictor for situations in which the method is likely to fail due to insufficient photon statistics. In practice a ground truth signal is not available, but by subsampling acquired data and comparing an LE surrogate signal derived from this data with one derived from the full-count data, it is possible to roughly determine whether the original estimate was in the plateau or breakdown regime. A good estimate from the plateau should be well-correlated with its reduced-count counterpart, and an estimate from the breakdown region will devolve into noise if counts are reduced. We tested this by reducing the counts in the first projection of each acquisition by an empirically chosen factor of 50% and computing a surrogate signal. Correlating this signal with the corresponding one from the full-count data yielded a value $\rho_{50\%}$ for each acquisition that serves as a performance predictor. Across all patient datasets, the predictors correlated significantly with the corresponding ρ_{Anzai} (ρ=0.77, $p<0.001$). This relationship between predictor and actual performance is of similar strength to the "Quality Factor" predictor based on a frequency domain analysis of projection data proposed in [48] for PET.

Going one step further, we performed a receiver operating characteristic (ROC) analysis using all of the available patient scans, where those with $\rho_{Anzai} \leq 0.6$ were classified as "true failures". The predictor's $\rho_{50\%}$ served as the discriminating variable, and a threshold θ applied to $\rho_{50\%}$ was swept from zero to unity. At each threshold evaluation, datasets with $\rho_{50\%} \leq \theta$ were classified as "predicted failures", allowing sensitivity and specificity values to the calculated at each point. The results can be seen in Fig. 8, where the ROC's area under the curve (AUC) of 0.92 indicate that $\rho_{50\%}$ is indeed a useful predictor of performance for failures caused by noise. The annotated operating point (θ=0.7) is sensitive enough to detect many failures and unlikely to misclassify good estimates.

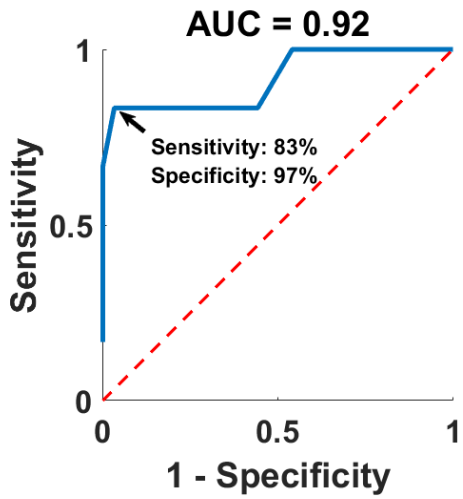

Fig. 8. ROC curve showing ability of performance predictor to classify poor surrogate signal estimates.

D. *Comparison to Other Respiratory Surrogate Signal Extraction Methods*

Table V contains a collection of results from other respiratory surrogate signal extraction methods reported in the literature. It can be seen that hardware-based methods agree well with each other, with radar and time of flight cameras providing very similar surrogate signals to those from an Anzai belt [21, 49]. A successful data-driven method should be capable of providing a similar level of consistency.

One data-driven method that produced comparable results utilized MR navigator images [50], which offer higher resolution and lower noise than projections from emission tomography. Büther *et al.* in [23] and Kesner *et al.* in [48] proposed data-driven methods for PET using non-DR, image-based metrics but reported somewhat poorer results.[2] Despite lower correlations to truth, both authors were able to show significant lesion motion in gated reconstructions. In [37], Thielemans *et al.* compared data-driven methods based on sensitivity [22] and spectral analysis [25] to PCA and LE for a collective of 14 FDG and four NH_3 cardiac PET patients. The authors did not report specific average correlations, but indicated in figures that both DR-based methods and the spectral analysis method performed well, with an average correlation greater than 0.8 against an infrared camera-based gold standard.

This body of work indicates that data-driven methods can approach the level of agreement of hardware-based techniques. Nevertheless, their application within emission tomography has been primarily limited to PET, a modality with superior angular coverage, resolution, and noise characteristics to traditional SPECT. Despite these hurdles, the proposed method is on par with data-driven techniques from PET, and even hardware-based approaches, provided sufficient counts are present.

[2] Büther *et al.* reported Spearman rank correlations rather than Pearson values, and the results are thus not directly comparable to others in the table. However, for data with a roughly linear relationship the two will be approximately equal, and for a nonlinear dataset Spearman should actually be higher.

TABLE V
SURVEY OF RESULTS OF RESPIRATORY SURROGATE SIGNAL EXTRACTION METHODS REPORTED IN THE LITERATURE

Paper	Application Modality	Proposed Method vs. Truth	Number of Test Subjects	Correlation
Pfanner *et al.* [49]	CT	Radar vs. Anzai	10	0.92
Schaller *et al.* [21]	ToF[a] Camera	ToF[a] vs. Anzai (Abdomen/Thorax)	13	0.91 / 0.85
Martinez-Möller *et al.*[50]	PET/MR	MR (diaphragm tracking) vs. Anzai	12	0.91
Büther *et al.* [23]	PET	PET Sensitivity vs. Video	29	0.65 (Spearman)
Büther *et al.* [23]	PET	PET Center of Mass vs. Video	29	0.68 (Spearman)
Kesner *et al.* [48]	PET	PET Voxel TACs[b] vs. Anzai	24	0.61
Proposed	SPECT	SPECT LE vs. Anzai	67	0.81

[a]Time of Flight
[b]Time Activity Curves

The evaluation in this paper and others is primarily based on comparison to the Anzai belt, which is a popular choice for a gold standard, as it is widely available and approved for clinical use. However, the Anzai does not provide a perfect ground truth. The pressure sensor outputs a one-dimensional measurement, whereas camera-based methods (and LE) utilize information from the entire torso and/or abdomen. When the type of breathing (abdominal or thoracic) is mismatched relative to the location of the Anzai sensor, the "gold standard" signal may therefore not correspond to one based on a large field of view, even though the latter is possibly more valid. This disagreement was noted by in [21] by Schaller *et al.*, who explicitly investigated the effect, and cannot be ruled out in all of our patients. Another potential source of error is time delay between the occurrence of internal motion and the arrival of surface pressure at the Anzai sensor. Due to our synchronization methodology, this time shift would cancel out if it were constant, but it is known to be variable [51] and could thus unpredictably degrade the correlation between the two signals.

V. CONCLUSION

In this work we proposed a fully-automated, data-driven respiratory surrogate signal extraction method for SPECT imaging with a traditional dual-headed system. We based the method on Laplacian Eigenmaps and added an adaptive scale parameter to improve usability despite wide variations in the properties of SPECT input data. We also laid out a series of post-processing steps that overcome the problems caused by SPECT's disjoint projections. In a subsequent patient validation, the proposed method was well-correlated with a clinically accepted gold standard. In a follow-up analysis, we proposed a predictor for identifying scans where the method is likely to fail.

Although promising, the validation is limited by the accuracy of the Anzai belt, which is an imperfect ground truth. Furthermore, although the respiratory surrogate signal is an enabler for subsequent respiratory gating or motion correction, the ultimate clinical utility of the work presented here must be established by future studies. Nevertheless, to the best of our knowledge, no such method for traditional SPECT has been described in the literature previously with an accompanying large patient validation. Our results are on par with those reported for other approaches and modalities, and the work presented here will facilitate future efforts toward respiratory motion management for SPECT imaging.

VI. ACKNOWLEDGMENT

The authors would like to thank our clinic's technologists who provided valuable assistance during collection of patient data. J.C. Sanders also extends his thanks to Jakob Wasza for the valuable discussions at the beginning of the work presented here.

VII. REFERENCES

[1] Y. Wang, S. J. Riederer, and R. L. Ehman, "Repiratory motion of the heart: kinematics and the implications for the spatial resolution in coronary imaging," *Magn. Reson. Med.*, vol. 33, no. 5, pp. 713-179, 1995.

[2] L. Le Meunier, R. Maass-Moreno, J. A. Carrasquillo, W. Dieckmann, and S. L. Bacharach, "PET/CT imaging: effect of respiratory motion on apparent myocardial uptake," *J. Nucl. Cardiol.*, vol. 13, no. 6, pp. 821-830, 2006.

[3] P. G. Danias, M. Stuber, R. M. Botnar, K. V. Kissinger, R. R. Edelman, and W. J. Manning, "Relationship between Motion of Coronary Arteries and Diaphragm During Free Breathing: Lessons from Real-Time MR Imaging," *Am. J. Roentgenol.*, vol. 172, no. 4, pp. 1061-1065, 1999.

[4] K. McLeish, D. L. G. Hill, D. Atkinson, J. M. Blackall, and R. Razavi, "A Study of the Motion and Deformation of the Heart due to Respiration," *IEEE Trans. Med. Imaging*, vol. 21, no. 9, pp. 1142-1150, 2002.

[5] B. M. W. Tsui, W. P. Segars, and D. S. Lalush, "Effects of Upward Creep and Respiratory Motion in Myocardial SPECT," *IEEE Trans. Nucl. Sci.*, vol. 47, no. 3, pp. 1192-1195, 2000.

[6] S. J. McQuaid and B. F. Hutton, "Sources of attenuation-correction artefacts in cardiac PET/CT and SPECT/CT," *Eur. J. Nucl. Med. Mol. Imaging*, vol. 35, no. 6, pp. 1117-1123, 2008.

[7] M. A. Clifford, F. Banovac, E. Levy, and K. Cleary, "Assessment of Hepatic Motion Secondary to Respiration for Computer Assisted Interventions," *Comput. Aided Surg.*, vol. 7, no. 5, pp. 291-299, 2002.

[8] E. Garin, Y. Rolland, J. Edeline, N. Icard, L. Lenoir, S. Laffont, *et al.*, "Personalized Dosimetry with Intensification Using ^{90}Y-Loaded Glass Microsphere Radioembolization Induces Prolonged Overall Survival in Hepatocellular Carcinoma Patients with Portal Vein Thrombosis," *J. Nucl. Med.*, vol. 56, no. 3, pp. 339-346, 2015.

[9] K. Suga, Y. Kawakami, M. Zaki, T. Yamashita, K. Shimizu, and N. Matsunaga, "Clinical utility of co-registered respiratory-gated ^{99m}Tc-Technegas/MAA SPECT-CT images in the assessment of regional lung functional impairment in patients with lung cancer,"

Eur. J. Nucl. Med. Mol. Imaging, vol. 31, no. 9, pp. 1280-1290, 2004.

[10] C. Chan, M. Harris, M. Le, J. Biondi, Y. Grobshtein, Y.-H. Liu, *et al.*, "End-expiration respiratory gating for a high-resolution stationary cardiac SPECT system," *Phys. Med. Biol.*, vol. 59, no. 20, pp. 6267-6287, 2014.

[11] J. Dey and M. A. King, "Non-Rigid Full Torso Respiratory Motion Correction of SPECT Studies," in *Nuclear Science Symposium and Medical Imaging Conference Record (NSS/MIC)*, Knoxville, TN, USA, 2010, pp. 2356-2358.

[12] G. Kovalski, O. Israel, Z. Keidar, A. Frenkel, J. Sachs, and H. Azhari, "Correction of Heart Motion Due to Respiration in Clinical Myocardial Perfusion SPECT Scans Using Respiratory Gating," *J. Nucl. Med.*, vol. 48, no. 4, pp. 630-636, 2007.

[13] H. Schumacher, J. Modersitzki, and B. Fischer, "Combined Reconstruction and Motion Correction in SPECT Imaging," *IEEE Trans. Nucl. Sci.*, vol. 56, no. 1, pp. 73-80, 2009.

[14] J. G. Parker, B. A. Mair, and D. R. Gilland, "Respiratory motion correction in gated cardiac SPECT using quaternion-based, rigid-body registration," *Med. Phys.*, vol. 36, no. 10, pp. 4742-4754, 2009.

[15] K. Müller, A. K. Maier, C. Schwemmer, G. Lauritsch, S. de Buck, J.-Y. Wielandts, *et al.*, "Image artefact propagation in motion estimation and reconstruction in interventional cardiac C-arm CT," *Phys. Med. Biol.*, vol. 59, no. 12, pp. 3121-3138, 2014.

[16] X. A. Li, C. Stepaniak, and E. Gore, "Technical and dosimetric aspects of respiratory gating using a pressure-sensor motion monitoring system," *Med. Phys.*, vol. 33, no. 1, pp. 145-154, 2006.

[17] T. Zhang, H. Keller, M. J. O'Brien, T. R. Mackie, and B. Paliwal, "Application of the spirometer in respiratory gated radiotherapy," *Med. Phys.*, vol. 30, no. 12, 2003.

[18] G. Chang, T. Chang, J. W. C. Jr., and O. R. Mawlawi, "Design and performance of a respiratory amplitude gating device for PET/CT imaging," *Med. Phys.*, vol. 37, no. 4, pp. 1408-1412, 2010.

[19] R. Beach, D., P. H. Pretorius, G. Boening, P. P. Bruyant, B. Feng, R. R. Fulton, *et al.*, "Feasibility of Stereo-Infrared Tracking to Monitor Patient Motion During Cardiac SPECT Imaging," *IEEE Trans. Nucl. Sci.*, vol. 51, no. 5, pp. 2693-2698, 2004.

[20] J. M. Mukherjee, B. F. Hutton, K. L. Johnson, P. H. Pretorius, and M. A. King, "An evaluation of data-driven motion estimation in comparison to the usage of external-surrogates in cardiac SPECT imaging," *Phys. Med. Biol.*, vol. 58, no. 21, pp. 7625-7646, 2013.

[21] C. Schaller, J. Penne, and J. Hornegger, "Time-of-Flight Sensor for Respiratory Motion Gating," *Med. Phys.*, vol. 35, no. 7, pp. 3090-3093, 2008.

[22] J. He, G. J. O'Keefe, T. Ackerly, and M. Geso, "Respiratory Motion Gating Based on List-mode Data in 3D PET: A Simulation Study Using the Dynamic NCAT Phantom," in *International Conference on Information Science and Engineering*, Nanjing, China, 2009, pp. 3697-3700.

[23] F. Büther, M. Dawood, L. Stegger, F. Wübbeling, M. Schäfers, O. Schober, *et al.*, "List Mode-Driven Cardiac and Respiratory Gating in PET," *J. Nucl. Med.*, vol. 50, no. 5, pp. 674-681, 2009.

[24] G. J. Klein, B. W. Reutter, E. H. Botvinivk, T. F. Budinger, and R. H. Huesman, "Fine-Scale Motion Detection Using Intrinsic List Mode PET Information," in *Workshop on Mathematical Methods in Biomedical Image Analysis*, Los Alamitos, CA, USA, 2001, pp. 71-78.

[25] P. J. Schleyer, M. J. O'Doherty, S. F. Barrington, and P. K. Marsden, "Retrospective data-driven respiratory gating for PET/CT," *Phys. Med. Biol.*, vol. 54, no. 7, pp. 1935-1950, 2009.

[26] P. P. Bruyant, M. A. King, and P. H. Pretorius, "Correction of the respiratory motion of the heart by tracking of the center of mass of theresholded projections: A simulation study using the dynamic MCAT phantom," *IEEE Trans. Nucl. Sci.*, vol. 49, no. 5, pp. 2159-2166, 2002.

[27] C.-L. Ko, Y.-W. Wu, M.-F. Cheng, R.-F. Yen, W.-C. Wu, and K.-Y. Tzen, "Data-driven respiratory motion tracking and compensation in CZT cameras: A comprehensive analysis of phantom and human images," *J. Nucl. Cardiol.*, vol. 22, no. 2, pp. 308-318, 2015.

[28] M. Bögel, H. G. Hofmann, J. Hornegger, R. Fahrig, S. Britzen, and A. Maier, "Respiratory Motion Compensation Using Diaphragm Tracking for Cone-Beam C-Arm CT: A Simulation and Phantom Study," *International Journal of Biomedical Imaging*, vol. 2013, pp. 1-10, 2013.

[29] L. Zijp, J.-J. Sonke, and M. van Herk, "Extraction of the respiratory signal from sequential thorax Cone-Beam X-ray images," in *14th ICCR*, Seoul, South Korea, 2004, pp. 507-509.

[30] J. B. Tenenbaum, V. d. Silva, and J. C. Langford, "A Global Geometric Framework for Nonlinear Dimensionality Reduction," *Science*, vol. 290, pp. 2319-2323, 2000.

[31] P. Fischer, T. Pohl, and J. Hornegger, "Real-Time Respiratory Signal Extraction for X-Ray Sequences Using Incremental Manifold Learning " in *IEEE 11th International Symposium on Biomedical Imaging (ISBI)*, Beijing, 2014, pp. 915-918.

[32] K. Thielemans, S. Rathore, F. Engbrant, and P. Razifar, "Device-less gating for PET/CT using PCA," in *Nuclear Science Symposium and Medical Imaging Conference Record (NSS/MIC)*, Valencia, Spain, 2011.

[33] H. Yan, X. Wang, W. Yin, T. Pan, M. Ahmad, X. Mou, *et al.*, "Extracting respiratory signals from thoracic cone beam CT projections," *Phys. Med. Biol.*, vol. 58, no. 5, pp. 1447-1464, 2013.

[34] M. Belkin and P. Niyogi, "Laplacian Eigenmaps for Dimensionality Reduction and Data Representation," *Neural Comput.*, vol. 15, no. 6, pp. 1373-1396, 2003.

[35] K. K. Bhatia, A. Rao, A. N. Price, R. Wolz, J. Hajnal, and D. Rueckert, "Hierarchical Manifold Learning," in *Medical Image Computing and Computer-Assisted Intervention-MICCAI 2012*, Nice, 2012, pp. 512-519.

[36] 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.

[37] K. Thielemans, P. J. Schleyer, P. K. Marsden, R. M. Manjeshwar, S. D. Wollenweber, and A. Ganin, "Comparison of different methods for data-driven respiratory gating of PET data," in *Nuclear Science Symposium and Medical Imaging Conference Record (NSS/MIC)*, Seoul, South Korea, 2013.

[38] J. C. Sanders, P. Ritt, T. Kuwert, A. H. Vija, and J. Hornegger, "Data-Driven Respiratory Signal Extraction for SPECT Imaging Using Laplacian Eigenmaps," in *Nuclear Science Symposium and Medical Imaging Conference Record (NSS/MIC)*, San Diego, CA, 2015.

[39] R. R. Coifman, Y. Shkolnisky, F. J. Sigworth, and A. Singer, "Graph Laplacian Tomography From Unknown Random Projections," *IEEE Transactions on Image Processing*, vol. 17, no. 10, pp. 1891-1899, 2008.

[40] J. Zeintl, T. D. Rempel, M. Bhattacharya, R. Malmin, and A. H. Vija, "Performance Characteristics of the SMARTZOOM® Collimator," in *Nuclear Science Symposium and Medical Imaging Conference Record (NSS/MIC)*, Valencia, Spain, 2011, pp. 2426-2429.

[41] M. U. A. Bromba and H. Ziegler, "Application Hints for Savitzky-Golay Digital Smoothing Filters," *Anal. Chem.*, vol. 53, no. 11, pp. 1583-1586, 1981.

[42] C. Liu, L. A. Pierce II, A. M. Alessio, and P. E. Kinahan, "The impact of respiratory motion on tumor quantification and delineation in static PET/CT imaging," *Phys. Med. Biol.*, vol. 54, no. 24, pp. 7345-7362, 2009.

[43] G. Chang, T. Chang, T. Pan, J. W. Clark Jr., and O. R. Mawlawi, "Implementation of an Automated Respiratory Amplitude Gating Technique for PET/CT: Clinical Evaluation," *J. Nucl. Med.*, vol. 51, no. 1, pp. 16-24, 2010.

[44] S. Quirk, N. Becker, and W. L. Smith, "External respiratory motion analysis and statistics for patients and volunteers," *J. Appl. Clin. Med. Phys.*, vol. 14, no. 2, pp. 90-101, 2012.

[45] W. Lu, P. J. Parikh, J. P. Hubenschmidt, J. D. Bradley, and D. A. Low, "A comparison between amplitude sorting and phase-ange sorting using external respiratory measurement for 4D CT," *Med. Phys.*, vol. 33, no. 8, pp. 2964-2974, 2006.

[46] M. Dawood, F. Büther, N. Lang, O. Schober, and K. P. Schäfers, "Respiratory gating in positron emission tomography: A quantitative comparision of different gating schemes," *Med. Phys.*, vol. 34, no. 7, pp. 3067-3076, 2007.

[47] M. C. K. Khoo, R. E. Kronauer, K. P. Strohl, and A. S. Slutsky, "Factors inducing periodic breathing in humans: a general model," *J. Appl. Physiol.*, vol. 53, no. 3, pp. 644-659, 1982.

[48] A. L. Kesner, R. A. Bundschuh, N. C. Detorie, M. Dahlbom, S. I. Ziegler, J. Czernin, *et al.*, "Respiratory Gated PET Derived in a Fully Automated Manner From Raw PET Data," *IEEE Trans. Nucl. Sci.*, vol. 56, no. 3, pp. 677-686, 2009.

[49] F. Pfanner, J. Maier, T. Allmenddinger, T. Flohr, and M. Kachelrieß, "Monitoring internal organ motion with continuous wave radar in CT," *Med. Phys.*, vol. 40, no. 9, pp. 091915-1-091915-16, 2013.

[50] A. Martinez-Möller, D. Zikic, R. M. Botnar, R. A. Bundschuh, W. Howe, S. I. Ziegler, *et al.*, "Dual cardiac-respiratory gated PET: implementation and results from a feasibility study," *Eur. J. Nucl. Med. Mol. Imaging*, vol. 34, pp. 1447-1454, 2007.

[51] D. Ionascu, S. B. Jiang, S. Nishioka, H. Shirato, and R. I. Berbeco, "Internal-external correlation investigations of respiratoy induced motion of lung tumors," *Med. Phys.*, vol. 34, no. 10, pp. 3893-3903, 2007.