Data-Driven Respiratory Signal Extraction for SPECT Imaging Using Laplacian Eigenmaps

James C. Sanders, *Student Member, IEEE*, Philipp Ritt, Torsten Kuwert, A. Hans Vija, *Member, IEEE*, Joachim Hornegger, *Member, IEEE*

***Abstract*— In Single Photon Emission-Computed Tomography (SPECT) imaging, respiratory motion may lead to artifacts and loss of quantitative accuracy. To compensate for this motion, a respiratory surrogate signal representing the patient's respiratory state over time is required. In practice, this surrogate signal is obtained via sensor-based approaches, but we seek to develop a data-driven solution that requires no external hardware. In this work, we compare two such methods, one linear and one non-linear, based on dimensionality reduction: Principle Component Analysis (PCA) and Laplacian Eigenmaps (LE). Our aim is to apply both to conventional SPECT and assess the feasibility of data-driven respiratory surrogate signal extraction for this modality. We expect that LE, which is less sensitive to outliers in data, will outperform PCA at high levels of image noise.**

Two phantom acquisitions were performed: one in which a sphere in cold background was translated axially by a piston actuator (dynamic), and a warm background with no sphere (static). Using binomial subsampling, both datasets were combined at various Signal-to-Noise Ratios (SNRs). LE and PCA surrogate signals were computed and compared via Pearson's correlation to the truth signal obtained from the actuator. As a follow-up, LE and PCA estimates from 27 cardiac SPECT acquisitions were compared to a simultaneously acquired signal from a pressure sensor embedded in an elastic belt.

In the phantom experiment, correlations between LE/PCA and truth were >0.9 for all SNR>5. For SNR<5, PCA deteriorated rapidly, whereas LE remained stable through SNR=2.5. For the patient validation, LE and PCA yielded average correlations of 0.86±0.14 and 0.37±0.26, respectively.

The phantom experiment indicated that LE outperforms PCA for low-SNR data. This conclusion was supported by the superior performance of LE for patient datasets, where noise may be high. However, the present work is limited by the simplistic motion present in the phantom experiment and the limited scope of the patient validation.

***Index Terms*—Laplacian Eigenmaps, manifold learning, respiratory gating, SPECT.**

Manuscript received November 23, 2015. This work was supported by Siemens Healthcare GmbH, Molecular Imaging.

James C. Sanders and Joachim Hornegger are with the Pattern Recognition Lab, University of Erlangen-Nuremberg, Martensstr. 3, 90108 Erlangen, Germany (e-mail: james.sanders@uk-erlangen.de).

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

A. Hans Vija is with Siemens Healthcare GmbH, Siemens Medical Solutions USA, Molecular Imaging, 2501 Barrington Rd., Hoffman Estates, IL 60169, USA.

I. Introduction

In Single Photon Emission-Computed Tomography (SPECT) imaging, patient respiratory motion occurring during the image formation process results in artifacts and quantitative errors in reconstructed images. To correct for this motion, a surrogate signal representing the motion state (e.g. a respiratory waveform) is required in order to subdivide the incoming data into segments, within which the patient is assumed approximately stationary.

Current methods for obtaining surrogate signals generally rely on external devices such as pressure belts or camera monitoring systems [1]. Data-driven solutions also exist that obviate the need for external hardware by estimating the surrogate signal from the data itself [2]. One family of such methods relies on the principle of dimensionality reduction (DR), which assumes that variations or structure in high-dimensional data can be summarized by a low-dimensional function. When applying this paradigm to emission tomography, it is assumed that if a stationary imaging system is acquiring a patient with static tracer kinetics, then the primary source of variation over time is due to respiratory motion. The DR algorithm then performs a mapping from the high-dimensional input space (SPECT projections in our case) to a low-dimensional surrogate signal where the variations over time are preserved.

One popular DR approach that yields a linear mapping is Principle Component Analysis (PCA), which has been applied with promising results to Positron Emission Tomography (PET) [3, 4]. However, PCA is based on an analysis of the data's covariance matrix which is sensitive to outliers and noise. For this reason it may be inappropriate for SPECT, where counts are lower and noise higher than in PET.

Laplacian Eigenmaps (LE), first proposed by Belkin and Niyogi in [5], is another DR method that has recently been successfully applied for the purpose of respiratory surrogate signal extraction to magnetic resonance and ultrasound imaging [6]. As opposed to PCA, LE performs a nonlinear mapping which makes no assumptions about the nature of the data's structure, which could foreseeably be beneficial in some cases. Furthermore, LE employs a kernel to only consider local relationships between input data points and is therefore potentially more robust to noise and outliers.

The aim of this study is to assess the feasibility of data-driven surrogate signal extraction in SPECT imaging using

PCA and LE. We expect LE, which focuses on local rather than global variations in the data and is not limited to a linear mapping, will outperform PCA, particularly at higher noise levels. We present results from a phantom experiment and cardiac patients that serve as a preliminary evaluation.

II. Methods

In order to apply DR to SPECT and extract a surrogate signal, list mode data from a two-headed Symbia T2 SPECT/CT (Siemens Healthcare GmbH, Molecular Imaging) system is first binned at each camera position into a series of T 300ms frames. Each image frame consists of data from both detectors, which are acquired simultaneously. All detected counts are utilized, as even scattered photons may contain information about patient movement.

Next, frames at each time bin i are reformatted as column vectors representing points in a high-dimensional space $\vec{\mathbf{x}}_i \in \mathbb{R}^{NM}$, where N and M are the number of frame rows and columns, respectively. A matrix $\mathbf{X} = [\vec{\mathbf{x}}_1, \dots, \vec{\mathbf{x}}_T]$ is then assembled and passed to the DR method of choice. The surrogate signal $\vec{\mathbf{y}} \in \mathbb{R}^{T\times K}$ is then estimated, where $K \ll NM$. In this work, we restrict ourselves to a one-dimensional surrogate signal and set $K := 1$.

PCA seeks to identify vectors in the input space along which the variance in the data is maximized. This is accomplished via an eigenvalue decomposition of the data's covariance matrix. In our case of a one dimensional surrogate signal, one simply chooses the eigenvector $\vec{\mathbf{v}}$ corresponding to the largest eigenvalue. The surrogate signal $\vec{\mathbf{y}}$ is then computed element-by-element by projecting the input data for each time bin onto this vector: $\mathrm{y}_i = \vec{\mathbf{v}}^{\mathrm{T}}\vec{\mathbf{x}}_i$. For the study presented here, we relied on the PCA implementation in the Dimensionality Reduction Toolbox [7].

LE, on the other hand, begins by first computing an affinity matrix $\mathbf{W} \in \mathbb{R}^{T\times T}$ whose elements $\mathrm{w}_{i,j}$ represent the similarity between two input frames at times i and j:

$$\mathrm{w}_{i,j} = e^{-\frac{\left\|\vec{\mathbf{x}}_i - \vec{\mathbf{x}}_j\right\|_2^2}{\alpha}}. \qquad (1)$$

Classically, the parameter α is set to some constant, and a pruning of the elements of $\mathbf{W}$ is performed based on thresholding or nearest neighbor rules [5, 6]. However, in our preliminary work we found α and pruning rules difficult to define in such a way that could guarantee good results across all of our datasets. We found that setting α to the mean of all Euclidean distances between the columns of $\mathbf{X}$ and forgoing pruning overcomes this difficulty.

After computation of the affinity matrix, the following objective function is minimized:

$$\underset{\vec{\mathbf{y}}}{\operatorname{argmin}} \sum_{i,j} \mathrm{w}_{i,j} \left\|\mathrm{y}_i - \mathrm{y}_j\right\|_2^2 \quad s.t. \quad \vec{\mathbf{y}}^{\mathrm{T}}\mathbf{D}\vec{\mathbf{y}} = 1, \qquad (2)$$

where $\mathbf{D}$ is a diagonal matrix with elements $\mathbf{D}_{i,i} = \sum_j \mathrm{w}_{i,j}$. This objective function seeks to ensure that time bins with similar image frames and, hence, larger $\mathrm{w}_{i,j}$'s, will produce similar surrogate signal amplitudes. The constraint prevents an arbitrary scaling of the solution.

Belkin and Niyogi showed in [5] that the problem can be minimized by means of a generalized eigenvector problem $(\mathbf{D} - \mathbf{W})\vec{\mathbf{v}} = \lambda\mathbf{D}\vec{\mathbf{v}}$. For a one-dimensional surrogate signal, $\vec{\mathbf{y}}$ is simply the eigenvector $\vec{\mathbf{v}}$ corresponding to the smallest non-zero eigenvalue. For a detailed discussion of the algorithm, see [5]. For the work here, we used an in-house implementation of LE written in Matlab (The MathWorks, Inc.).

This process is repeated across all camera positions, and results from each are concatenated after baseline correction and normalization to yield a respiratory surrogate for the complete acquisition. In general, due to the nature of eigenvector-based methods, the polarity of the estimates at each camera position may be inverted with respect to others. This issue was not considered here, and polarity was manually corrected. Estimates from both methods were lowpass filtered with a Savitzky-Golay kernel prior to analysis.

To assess performance of LE and PCA, a phantom experiment was carried out consisting of two acquisitions. In the first, we filled an elliptical IEC phantom with cold water and a 16 mL sphere with 90 MBq Tc-99m. The sphere was then mounted on a dynamic piston phantom and placed inside the center of the background via a hollow cylinder insert with an external access port. The piston provided a ±1cm axial translation of the sphere with a respiratory-like waveform with period of 4 s that was available for subsequent analysis. For the second acquisition, the sphere was removed, and 492 MBq Tc-99m was added to the 8.9L background. This yielded a dynamic dataset with the sphere in motion and a static dataset of the warm background. Both scans consisted of 60 views of 20 seconds each.

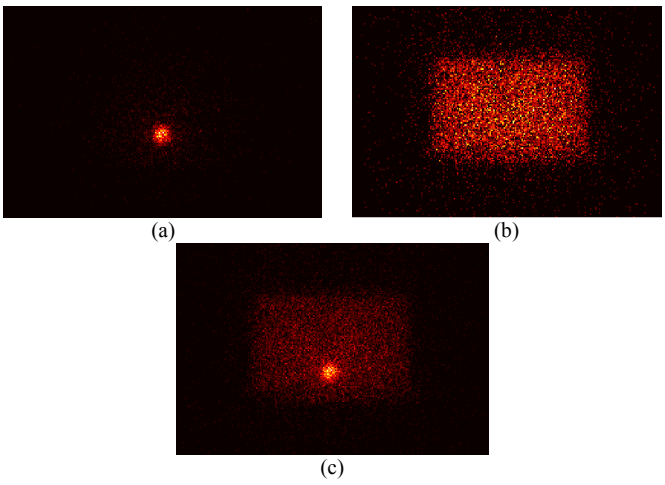

Fig. 1. Example projection data from phantom acquisition. Dynamic dataset (a) with hot sphere mounted on respiratory piston phantom was registered with and added to static background data (b) at various contrast levels derived via binomial subsampling. Surrogate signals were estimated from aggregate data (c) using PCA and LE.

Following the acquisitions, nuclear list mode data from both were binned into 300 ms time bins and registered using x-ray CTs acquired after each SPECT scan. Registered bins were then summed after binomial subsampling to yield datasets

with a range of Signal to Noise Ratios (SNRs, defined as ratio of mean dynamic count density to square root of mean static count density). This process is illustrated in Figure 1. Ten realizations of the subsampled datasets were computed at each SNR level. LE and PCA surrogate signal estimates were computed and compared against the truth via Pearson's correlation.

In a follow-up validation, SPECT data from 14 patients undergoing a one-day stress/rest cardiac study were collected (27 scans total). For stress and rest scans, respectively, patients were injected with 236.9±21.6 and 698.6±35.9 MBq Tc-99m-MIBI before an IQ·SPECT acquisition consisting of 17 views at 30 s each with a converging collimator [8]. A respiratory surrogate signal was simultaneously recorded using an elastic belt with embedded pressure sensor (Anzai Medical Co.). LE and PCA estimates were similarly computed from the list mode data and correlated with the pressure sensor signal, which served as truth.

III. Results

Figure 2 shows results of the phantom experiment. Both methods show similar behavior across the range of SNRs, and the data are well fit by a logistic function, yielding r^2 coefficients of 0.99 and 0.98 for LE and PCA, respectively. In the asymptotic regime (SNR > 5), both approaches behave similarly. However, the performance of PCA degrades rapidly for SNR < 5, while that of LE remains stable through higher noise levels.

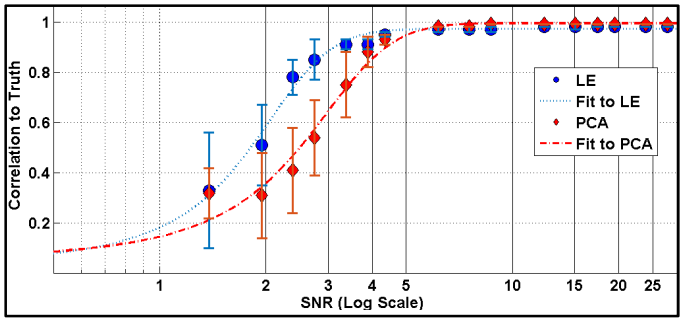

Fig 2. Correlation to truth versus SNR for LE (blue circles) and PCA (red diamonds). Error bars represent ±1 standard deviation across the ten realizations. Fits were parameterized by a logistic function. Note how LE outperforms PCA at low SNR levels.

The patient results are summarized in Table I. Overall, LE outperformed PCA with a mean correlation (±standard deviation) of 0.86±0.14 compared to 0.37±0.26. The difference is starker when comparing stress and rest scans separately. Although rest acquisitions have roughly three and a half times as many counts as stress, for LE, the median correlations for both were statistically equal. For PCA, however, the median correlation the stress data was significantly worse than that of the rest scans ($p<0.001$). This is illustrated qualitatively in Figure 3, which shows estimates for one camera position of a representative patient stress acquisition.

LE performed consistently well across all rest scans and all but two stress scans. These scans, which were classified as outliers via a Tukey outlier filter [9] are responsible for the skew between the median and mean in the LE data and the larger standard deviation in this group. PCA, on the other hand, failed in numerous patients, and, in general, yielded surrogate signals of much more inconsistent quality.

TABLE I. Patient Results – Correlations with Anzai Belt

	LE		PCA	
Scan Type	*Mean±SD*	*Median*	*Mean±SD*	*Median*
Stress	0.84±0.17	0.87	0.18±0.10	0.17
Rest	0.91±0.04	0.92	0.57±0.23	0.63
All	0.86±0.14	0.90	0.37±0.26	0.30

Correlations from all scans are significant with p-value < 0.001 except for one PCA scan, whose p-value = 0.02.

IV. Discussion and Conclusion

Results from the phantom experiment show that, as a function of SNR, the performance of PCA and LE follow a logistic curve. However, although PCA and LE perform similarly for high-SNR data, LE appears to be more robust at higher noise levels. This behavior was also observed in the patient validation, where, with the exception of two outliers, LE remained consistent for lower count data, while those of PCA deteriorated significantly.

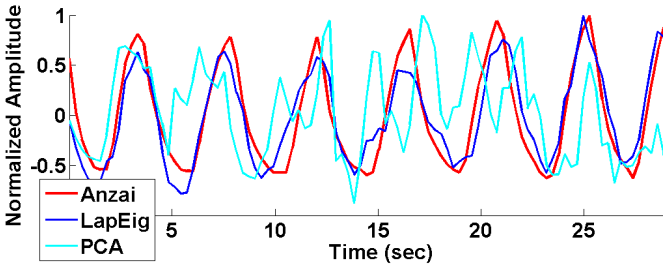

Fig. 3. Representative results from one camera position of a patient stress acquisition. Although somewhat noisy, there is a good correspondence between the Anzai (red), and LE (dark blue) respiratory surrogates. The PCA-derived surrogate (cyan) shows only some structure in common with Anzai. The correlations between the LE and PCA surrogates and the Anzai signal were 0.85 and 0.19, respectively, over the entire scan duration.

This seems to imply that the aspect of locality in the input space imposed by the kernel in (1) might indeed succeed at mitigating the influence of noise and outliers for LE. Only very similar frames will result in weights large enough to affect the estimated surrogate amplitude at a particular time bin, and outliers will have very little effect at all. For PCA, on the other hand, the image frame at each time bin affects the resulting surrogate value at all other times, allowing even a single outlier to redirect the principal component in an unwanted direction not truly related to respiratory motion. It is possible that the nonlinear nature of LE provides an added benefit as well, but this cannot be concluded from our available data.

In [10], Thielemans et al. observed that PCA actually outperformed LE at low counts in a collective of eleven FDG-PET patients. However, the differences were small, and neither method completely broke down. The differing conclusion could be explained by the fact that the authors of that study used default parameters from the Dimensionality

Reduction Toolbox when processing their data. In our preliminary work, we found that LE is very sensitive to parameter selection (hence our decision to set α adaptively). A further, and perhaps more significant, difference between the two studies is the fundamental difference in data coverage between PET and conventional SPECT. The former provides full tomographic coverage for the entire acquisition duration, whereas the latter must cope with disjoint measurements from different projection angles. More work must be done to investigate the effect differing data coverage has on these methods.

The results presented here indicate that data-driven respiratory surrogate signal extraction using Laplacian Eigenmaps seems well suited for conventional SPECT imaging. However, limitations of our study include the simplistic type of motion examined in the phantom experiment and limited scope of the patient validation. Future work should focus on validating the LE method for more complicated motion scenarios where truth is known (e.g. dynamic anthropomorphic phantoms) and a broader patient collective to evaluate LE's performance for tracers with differing distribution characteristics.

ACKNOWLEDGMENT

The authors would like to thank the technologists at the Clinic of Nuclear Medicine, University Hospital Erlangen who expended valuable time and effort to obtain the cardiac patient data.

J. C. S. thanks X. Rong for valuable theoretical discussions of the material presented above, as well as his review of this manuscript.

REFERENCES

[1] S. A. Nehmeh and Y. E. Erdi, "Respiratory Motion in Positron Emission Tomography/Computed Tomography: A Review," *Semin. Nucl. Med.,* vol. 38, pp. 167-176, 2008.

[2] F. Büther, M. Dawood, L. Stegger, F. Wübbeling, M. Schäfers, O. Schober, and K. P. Schäfers, "List Mode-Driven Cardiac and Respiratory Gating in PET," *J. Nucl. Med.,* vol. 50, pp. 674-681, 2009.

[3] K. Thielemans, S. Rathore, F. Engbrant, and P. Razifar, "Device-less gating for PET/CT using PCA," in *IEEE NSS/MIC*, Valencia, Spain, 2011, pp. 3904-3910.

[4] S. Fürst, R. Grimm, I. Hong, M. Souvatzoglou, M. E. Casey, M. Schwaiger, S. G. Nekolla, and S. I. Ziegler, "Motion Correction Strategies for Integrated PET/MR," *J. Nucl. Med.,* vol. 56, pp. 261-269, 2015.

[5] M. Belkin and P. Niyogi, "Laplacian Eigenmaps for Dimensionality Reduction and Data Representation," *Neural Comput.,* vol. 15, pp. 1376-1396, 2003.

[6] 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, pp. 806-818, 2012.

[7] L. J. P. Van der Maaten, E. O. Postma, and H. J. Van Den Herik, "Dimensionality reduction: A comparative review, " *J. Mach. Learn. Res.*, vol. 10, 2009.

[8] J. Zeintl, T. D. Rempel, M. Bhattacharya, R. E. Malmin, and A. H. Vija, "Performance Characteristics of the SMARTZOOM® Collimator," in *IEEE NSS/MIC*, Valencia, Spain, 2011, pp. 2426-2429.

[9] J. W. Tukey, *Exploratory Data Analysis*, Reading, MA: Addison-Wesley, 1977, pp. 43-44.

[10] K. Thielemans, P. 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 *IEEE NSS/MIC*, Seoul, South Korea, 2013, pp. 1-4.