

Temporal Shape Analysis via the Spectral Signature

Elena Bernardis¹, Ender Konukoglu², Yangming Ou¹, Dimitris N. Metaxas³, Benoit Desjardins¹, and Kilian M. Pohl¹

¹ Dept. of Radiology, University of Pennsylvania, Philadelphia, PA 19104, USA

² Microsoft Research, Cambridge, CB3 0FB, UK

³ Dept. of Computer Science, Rutgers University, Piscataway, NJ 08854, USA

Abstract. In this paper, we adapt spectral signatures for capturing morphological changes over time. Advanced techniques for capturing temporal shape changes frequently rely on first registering the sequence of shapes and then analyzing the corresponding set of high dimensional deformation maps. Instead, we propose a simple encoding motivated by the observation that small shape deformations lead to minor refinements in the spectral signature composed of the eigenvalues of the Laplace operator. The proposed encoding does not require registration, since spectral signatures are invariant to pose changes. We apply our representation to the shapes of the ventricles extracted from 22 cine MR scans of healthy controls and Tetralogy of Fallot patients. We then measure the accuracy score of our encoding by training a linear classifier, which outperforms the same classifier based on volumetric measurements.

1 Introduction

Capturing the shape and function of anatomy through volumetric measurements extracted from 4D medical scans has become of central importance in diagnosing diseases. For example, cardiologists rely on ejection fraction extracted from ultrasound or cine MR scans to assess patients. These volumetric measurements, however, are not sensitive enough to aid the diagnosis of many focal or diffuse cardiac diseases. In this paper, we introduce a new encoding of the shape and its temporal changes based on the spectral signature and show that this encoding is more sensitive for comparing two shapes and their temporal dynamics than volumetric measurements.

Advanced techniques for capturing the changes in shape over time frequently rely on registering the sequence of images and then analyzing the corresponding set of deformation maps [1,2]. Examples of this type of analysis, specific to the heart, include mappings motivated by biomechanical models [3] or mathematical properties [4], from which statistical models of the cardiac ventricles can be extrapolated [5]. While promising, these approaches are susceptible to negatively biasing the analysis due to the underlying assumptions and parameter settings of the registration framework, as well as the accuracy in reducing high dimensional deformation maps to a few features.

An alternative avenue is to characterize the shape of each structure with a low-dimensional set of parameters and use the latter for structure discrimination. Spectral signatures, i.e. eigensystems of the Laplace and Laplace-Beltrami operators, have recently gained popularity as powerful shape descriptors [6,7]. The eigenvalues of the

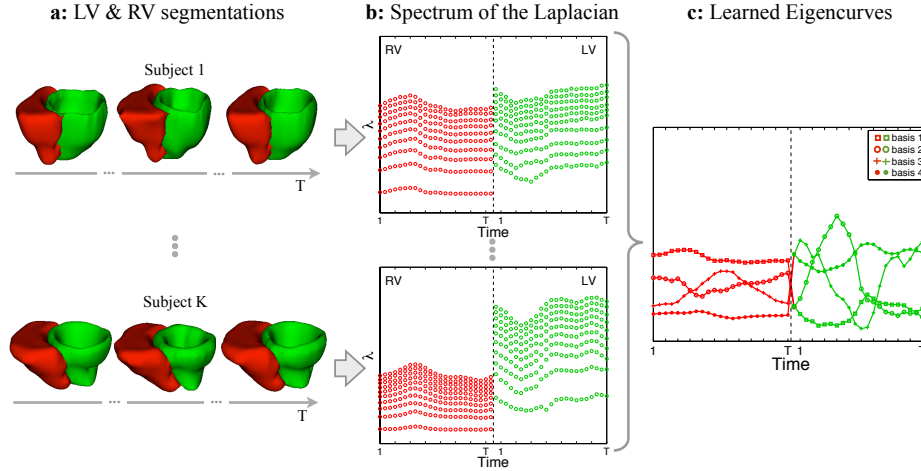


Fig. 1. Our method. **a:** Right (RV) and left (LV) ventricle segmentations for each heart sequence at three sample timepoints. **b:** Eigenvalue curves that encode temporal shape changes displayed as a function of time. **c:** RV and LV reduced shape signatures learned by a low dimensional embedding of these curves, while assuming no prior physical or statistical model of the data.

Laplace operator implicitly carry information on local shape invariants such as curvature, surface area and volume, and allow to encode shape information with few parameters without need of prior registration [8]. In this paper, we specifically explore the adaptation of this technology from capturing individuals shapes to capturing temporal morphologic changes.

Our work is motivated by the observation that small shape deformations lead to minor refinements in the spectral signature. Spectral signatures are thus well suited for capturing temporal shape changes of anatomical structures that vary slowly between measurements, such as cardiac ventricles from cine MR scans. Our representation captures temporal shape changes by first computing the spectral signature for each time point yielding a family of eigenvalue curves (Fig. 1b). We then encode temporal shape changes by a low dimensional embedding of the eigenvalue curves (Fig. 1c). By doing so, our simplistic representation assumes no prior physical or statistical model of the data and only depends on two parameters, namely, the number of eigenvalues and the dimension chosen for the lower dimension embedding.

We apply our representation to the shapes of the ventricles extracted from 22 cine MR scans of 11 healthy controls and 11 Tetralogy of Fallot (TOF) patients. We choose this specific scenario as there is no uncertainty, unlike with other diseases, about the diagnosis, so that the labeling of individual data sets can be viewed as ground truth. In addition, TOF is suitable for temporal analysis as it affects both RV shape and cardiac function. We then measure the accuracy score of our encoding by training a linear classifier and recording the leave-one-out cross-validation accuracy in distinguishing these two populations. Our representation outperforms the linear classifier based on volumetric measurements of the ventricles. Before we describe our encodings and experiments in further detail, we just note that we do not attempt to provide a representation tailored

towards cardiac disease detection but rather derive a shape representation for implicitly encoding temporal morphological changes that can be applied to the cardiac domain.

2 Temporal Shape Encoding

We now present in further detail the extension of the spectral signature of the Laplacian operator to the temporal domain. Laplace operators and their spectra have been studied in mathematics for a long time [9,10,11]. Their introduction in computational shape analysis is, however, rather recent [12]. We start with an introduction to Laplace operators and then describe our spectral shape encoding for temporal shape deformations.

2.1 Spectrum of Laplace Operator

Our brief overview of the Laplace operators is meant to give the necessary background to understand their role in temporal shape encoding. For a more thorough discussion, we refer the reader to [11,12].

We denote an object as a closed bounded domain $\Omega \subset \mathbb{R}^d$ with piecewise smooth boundaries. With respect to medical imaging domain, Ω corresponds to the volume extracted from the segmentation of an anatomical structure. The Laplace operator Δ_Ω on Ω is defined as $\Delta_\Omega f \triangleq \sum_{i=1}^d \frac{\partial^2}{\partial x_i^2} f$ for a twice differentiable function f , where $x \triangleq \{x_1, \dots, x_d\}$ are the spatial coordinates. The importance of this operator for shape analysis arises from its eigenvalues and eigenfunctions, which are the solutions of the Helmholtz equation with Dirichlet type boundary conditions, $\Delta_\Omega f + \lambda f = 0$, $\forall \mathbf{x} \in \Omega$ and $f(\mathbf{x}) = 0$, $\forall \mathbf{x} \in \partial\Omega$, where $\partial\Omega$ denotes the boundary of the object and $\lambda \in \mathbb{R}$ is a scalar [11]. There are infinite pairs of $\{(\lambda_j, f_j)\}_{j=1}^\infty$ that satisfy this equation and the ordered set of eigenvalues form a positive diverging sequence $0 < \lambda_1 \leq \lambda_2 \leq \dots$, called the Dirichlet spectrum of Δ_Ω , which we simply refer to as the ‘spectrum’.

The spectrum has several advantageous properties for shape analysis in medical image analysis [11]. First, the spectrum encodes information regarding the intrinsic geometry of the object. This information content is due to an identity called *heat-trace* and its equivalent polynomial expansion $Z(\tau) \triangleq \sum_{j=1}^\infty e^{-\lambda_j \tau} = \sum_{m=0}^\infty a_{m/2} \tau^{-d/2+m/2}$, with $\tau > 0$. The coefficients $a_{m/2}$ are given as sums of volume and boundary integrals of some local invariants of Ω [13,11], such as its volume or its surface mean curvature. The functional relation between $\{\lambda_j\}_{j=1}^\infty$ and $\{a_{m/2}\}_{m=0}^\infty$ links the geometry of an object to its spectrum, and is the ingredient that makes the Laplace operator interesting for shape analysis. Second, the eigenvalues are invariant to isometric transformations. This invariance can be even extended to scaling [12], thus eliminating the need to align the shapes to a common coordinate system for further analysis. Finally, the eigenvalues change continuously with the deformations applied to the object’s boundary, i.e. there is a continuous link between the differences in eigenvalues and the difference in shape. This continuous link is a critical component for our encoding of temporal shape deformation. For a deeper intuition, let us consider an object that changes its shape with time. We represent the temporal dependence of the object’s shape with $\Omega(t)$ and the temporal dependence of its spectra with $\Lambda(t)$. Now, if we interpret the motion of an object as the deformation between two time points, $\Lambda(\cdot)$ not only then captures the geometry of $\Omega(t)$

at a given t but it also provides information regarding the motion of $\Omega(\cdot)$. In this paper, we make use of both information when encoding 3D+t(ime) objects.

2.2 Learning Temporal Shape Changes

We now outline the framework for encoding 3D + t. Given a set of K 3D+t volumes defined over T time points, we first compute for each volume $i \in \{1, \dots, K\}$ and time point $t \in \{1, \dots, T\}$ the first N eigenvalues of the spectrum (see also Fig. 1b), which we denote with $\{\lambda_1^i(t), \lambda_2^i(t), \dots, \lambda_N^i(t)\}$. We then construct temporal curves from the eigenvalues across the N signatures of each volume, i.e. $\lambda_j^i := \{\lambda_j^i(1), \dots, \lambda_j^i(T)\}$, for $j \in \{1, \dots, N\}$, is the j^{th} temporal (eigenvalue) curve of volume i . The final step encodes the temporal shape deformations by learning a low dimensional embedding of those curves across all K volumes. By doing so, further analysis of our encoding always has to consider the subject specific matrix of eigenvalues defining curves across time instead of viewing λ_j^i individually.

For dimensionality reduction, we apply Non-negative Matrix Factorization (NMF) [14] to the $N \cdot K$ temporal curves. In general, the dimension of the data matrix V is $M \times O$ where M is the number of measurements and O the number of objects. In our temporal shape encoding, each column of the data matrix V represents a temporal curve λ_j^i . Thus, the dimension of V becomes the (number of timepoints T) \times ($K \cdot N$). NMF then factorizes matrix V into a basis matrix H and coefficient matrix W so that

$$V \approx WH \text{ subject to minimizing } F(W, H) \triangleq \sum_{i=1}^{N \cdot K} \sum_{\mu=1}^T [V_{i\mu} \log(WH)_{i\mu} - (WH)_{i\mu}],$$

where V refers to the entries of the corresponding matrix. Setting the number of basis vectors to B and $b \in \{1, \dots, B\}$, the optimal H and W are obtained by finding the local minimum of $F(\cdot, \cdot)$ via the following iterative algorithm:

$$\mathbf{1:} W'_{ib} \leftarrow W_{ib} \sum_{\mu=1}^T \frac{V_{i\mu}}{(WH)_{i\mu}} H_{b\mu} \quad \mathbf{2:} W_{ib} \leftarrow \frac{W'_{ib}}{\sum_{j=1}^{N \cdot K} W'_{jb}} \quad \mathbf{3:} H_{b\mu} \leftarrow H_{b\mu} \sum_{i=1}^{N \cdot K} W_{ib} \frac{V_{i\mu}}{(WH)_{i\mu}}$$

All temporal curves can now be represented by the basis matrix H and the corresponding B coefficients provided by W . These B coefficients across all the sequences of N eigenvalues are then the temporal shape encodings of our 3D+t objects.

We note that one could have chosen any other dimensionality reduction method. We simply choose NMF as it does not make any assumption of the underlying distribution, unlike for example Principle Component Analysis as motivated by our experiments.

3 Experimental Setup

We evaluate our temporal shape encoding by applying it on ventricle segmentations from short-axis cardiac MRI scans and comparing its accuracy score in a classification problem to discriminate healthy individuals from patients who had Tetralogy of Fallot (TOF) corrected surgery in infancy. Our dataset includes the cine MR scans of 11 TOF

cases and 11 healthy volunteers ($K=22$). TOF patients are post repair and age-matched to normal controls. The cine MR scans of each case are normalized to 23 timepoints. A medical expert then semi-automatically segmented the blood pool of the right ventricle and myocardium of the left ventricle at the end-diastole (ED) timepoint using ‘Segment’ [15] with manual corrections of the results. For simplicity, we refer to the blood pool of the right ventricle just as right ventricle or RV, and to the myocardium of the left ventricle as left ventricle or LV. We then propagate this segmentation to the other time points via the non-rigid registration of [16]. We refine those segmentations via dilation and erosion to remove possible holes, disconnections or topology changes caused by registration errors. Sample segmentation sequences are shown in Fig. 1a.

For computing the Laplace spectrum, we define the shapes of LV and RV based on the corresponding segmentations. We then compute the Dirichlet eigenvalues using finite differences on the regular grid. While using a high order FEM method might provide more accurate eigenvalues, this is outside the scope of this article as our focus is to apply the signature to the temporal domain, hence we use the simplest implementation. Furthermore, one can also consider using Neumann eigenvalues however they would be more sensitive to noisy segmentations e.g. isolated islands.

We analyze the accuracy of our new temporal shape encoding in describing the changes in shape describing the changes in shape of the RV and LV by first applying it as well as related representations to our sequences of segmentations. We then measure the accuracy of a linear Support Vector Machine (SVM) [17] in correctly labeling images based on those features via the leave-one-out principle. Table 1 records the outcome when applying the classifier to LV and RV separately as well as when we combine the eigencurves of both ventricles.

We start by extracting the volumes and computing the first eigenvalue of the RV and LV. We view each volume or first eigenvalue as a feature so that classification of every RV and LV is based on 23 features (given from the original 23 timepoints). When concatenating the ventricles (which we will denote by RV-LV), classification is based on 46 features instead. Classification results based on volume and based on the first eigenvalue are shown in Table 1 (a) and (b) respectively. Simply using volume for distinguishing shape changes between TOF and normal subjects only leads to an accuracy of approximately 60% for LV and RV separately and of 68% when combined. As mentioned in the previous section, the functional relation between $\{\lambda_j\}_{j=1}^{\infty}$ and $\{a_{m/2}\}_{m=0}^{\infty}$ is the link between the geometry of an object and its spectrum. In particular, the first eigenvalue is directly proportional to the volume of object j since $a_0 = (4\pi)^{-3/2}V_j$. As predicted, the results obtained using the first eigenvalue alone are comparable to the volume ones.

As shown in the rest of the experiments, increasing the number of eigenvalues adds shape information, improving the overall recall. We start by considering the eigenvalue sequences λ_j^i individually. We do so by computing the eigenvalues for each of the hearts, setting the spectrum cutoff at $N = 100$. For comparison to our new encoding, the results of Table 1 (c) do not explicitly model temporal dependencies as we first combine the $23 \cdot N$ features for each ventricle (or $46 \cdot N$ for RV-LV). We then reduce the $M = 23 \cdot N$ eigenvalues (or $46 \cdot N$ for RV-LV) for the $O = K$ hearts via NMF (see Sec. 2.2) to $B = 22$ features. Parameters N and B are chosen empirically to be

	LV Prec.	Recall	Accuracy	RV Prec.	Recall	Accuracy	RV-LV Prec.	Recall	Accuracy
a: volume	0.7727	0.5862	0.6136	0.9394	0.5741	0.6212	0.9301	0.6274	0.6888
b: λ_1	0.8182	0.5625	0.5909	0.8636	0.6441	0.6932	0.8864	0.6290	0.6818
c: $\lambda_{1:1:100}$	0.7143	0.5914	0.6104	0.7557	0.6425	0.6676	0.8182	0.7674	0.7851
d: $\lambda_{1:1:100}$ (3D+t)	0.6970	0.8214	0.7727	0.7576	0.8065	0.7879	0.8788	0.8286	0.8485
e: $\lambda_{1:10:100}$ (3D+t)	0.7879	0.8667	0.8333	0.8182	0.8438	0.8333	0.9545	0.8400	0.8864
f: PCA $\lambda_{1:1:100}$	0.7273	0.6957	0.7045	0.7273	0.7500	0.7424	0.7879	0.7647	0.7727

Table 1. TOF classification results. For each experimental setup, we record precision (prec.), recall and accuracy for LV, RV and RV-LV, obtained from our leave-one-out cross validation. The first eigenvalue (b) alone performs similarly to the volume feature (a). Increasing number of eigenvalues (c) improves the performance, especially recall. d-e) Our 3D+t shape signatures method yields increased scores. f) choosing PCA instead of NMF on the same eigencurves used for (d) results in lower scores.

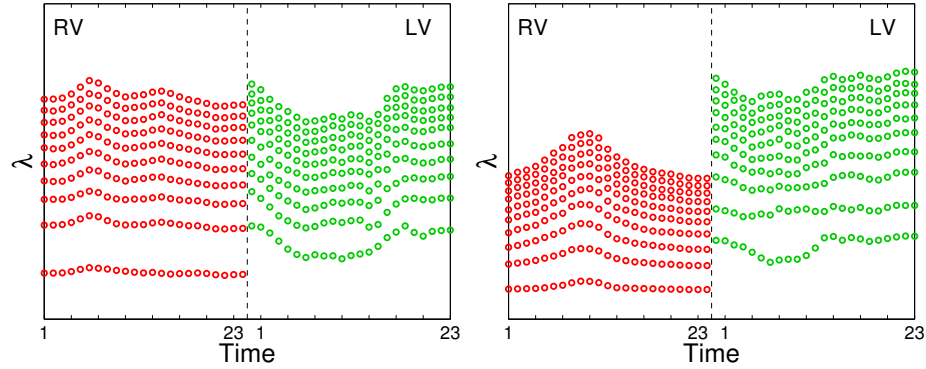
minimum while maximizing classification accuracy. We experimented for N from 1 to 200 and for B from 2 to 200. Furthermore, B is set to the same for the LV, RV, and RV-LV scenarios. The resulting classification scores (c) are higher compared to (a,b).

We further improve the classification score by using our new 3D+t temporal shape representation of Sec. 2.2 (Table 1 (d,e)). Here, we reduced the dimensionality of the $N \cdot K$ eigencurves from $T = 23$ entries (respectively 46 for the RV-LV) to $B = 4$ after exploring the entire range from 2 to 23. Fig. 2a illustrates sample eigencurves. We then feed $B \cdot C$ features to the classifier, where C is the number of curves selected for the classification step. In line (d), we use all the eigencurves for classification, hence setting $C = N (= 100)$, while we subsample the eigencurves in (e) setting $C = 10$. The 3D+t encoding yields a higher precision and significantly improves recall, hence giving an overall improved accuracy. As shown from the classification results in (e), selecting a subset of the eigencurves, once all the eigencurves have already been used for the learning step, allows to improve the outcome of the classifier, indicating that another method for data compression or classification could improve even further the results of our temporal shape descriptor as well as that of the other scores.


To motivate the use of NMF, we also applied PCA (f) to learn the temporal eigencurves using the same setup of (d). The accuracy scores drop to the level we measured for the volumes scores indicating that the Gaussian assumption of PCA is violated by these sets of temporal eigencurves. Our quantitative findings are also reflected in the visual comparison of the basis vectors for $B = 4$ of the two dimensionality reduction methods in Fig. 2b. While the bases obtained by PCA are very noisy, the NMF ones are cleaner and better describe the smooth temporal changes observed in the eigencurves.

4 Conclusions

In this paper, we exploit the implicit local shape properties captured by spectral signatures, i.e. eigensystems of the Laplace and Laplace-Beltrami operators, and adapt them to capture morphological changes over time. We propose a fairly simple encoding based on the observation that small shape deformations lead to minor refinements in



a: Sample eigencurves for healthy (left) and TOF (right) cases.



b: Reduced spectral signatures

Fig. 2. Shape signatures. **a:** Sample eigencurves before reduction. **b:** Learned spectral signatures reduced via NMF and PCA for RV and LV (over 23 timepoints), and for RV-LV (over 46 timepoints). In the absence of the correct initial statistical model, PCA results in a high level of noise and does not capture the smooth temporal changes of the eigencurves.

the spectral signature. The shape analysis is independent of the original segmentation used and, given initial segmentations, computing the spectral signature is independent of registration. The accuracy obtained in the classifications demonstrate that that our temporal-shape representation can be successfully used to classify TOF cardiac disease patients. Our results indicate that our new temporal shape representation better incorporates the temporal relation of the data. Thus, we are able to better capture RV and LV deformations in this population than scores capturing shape changes over time by separately measuring the 3D shape at each timepoint.

Acknowledgments

We would like to thank DongHye Ye for his help on generating the cardiac dataset. This project was supported in part by Grant Number UL1RR024134 and by the Institute for Translational Medicine and Therapeutics' (ITMAT) Transdisciplinary Program.

References

1. Ardekani, S., Weiss, R.G., Lardo, A.C., George, R.T., Lima, J.A.C., Wu, K.C., Miller, M.I., Winslow, R.L., Younes, L.: Cardiac motion analysis in ischemic and non-ischemic cardiomyopathy using parallel transport. In: ISBI. (2009)
2. Beg, M.F., Helm, P.A., Mcveigh, E., Miller, M.I., Winslow, R.L.: Computational cardiac anatomy using mri. magnetic resonance in medicine (2004)
3. Wang, X., Chen, T., Zhang, S., Metaxas, D., Axel, L.: Lv motion and strain computation from tmri based on meshless deformable models. In: MICCAI. (2008) 636–644
4. Helm, P., Beg, M.F., Miller, M.I., Winslow, R.L.: Measuring and mapping cardiac fiber and laminar architecture using diffusion tensor mr imaging. *Annals of the New York Academy of Sciences* **1047**(1) (2005) 296–307
5. Mansi, T., Durrleman, S., Bernhardt, B., Sermesant, M., Delingette, H., Voigt, I., Lurz, P., Taylor, A.M., Blanc, J., Boudjemline, Y., Pennec, X., Ayache, N.: A statistical model of right ventricle in tetralogy of fallot for prediction of remodeling and therapy planning. In: MICCAI. (2009) 214–221
6. Reuter, M., Niethammer, M., Wolter, F.E., Bouix, S., Shenton, M.: Global medical shape analysis using the volumetric laplace spectrum. In: Proceedings of the 2007 International Conference on Cyberworlds, NASA-GEM Workshop, Los Alamitos, CA, USA, IEEE Computer Society (2007) 417–426
7. Reuter, M.: Laplace Spectra for Shape Recognition. Books on Demand GmbH (2006)
8. Niethammer, M., Reuter, M., erich Wolter, F., Peinecke, N., seong Koo, M., Shenton, M.E.: Global medical shape analysis using the laplace-beltrami spectrum. In: MICCAI. (2007)
9. Weyl, H.: Das asymptotische verteilungsgesetz der eigenwerte linearer partieller differentialgleichungen. *Math Ann* (1912) 441–69
10. Kac, M.: Can one hear the shape of a drum? *The American Mathematical Monthly* **73**(7) (1966) 1–23
11. Courant, R., Hilbert, D.: *Method of Mathematical Physics, vol I*. Interscience Publishers (1966)
12. Reuter, M., Wolter, F.E., Peinecke, N.: Laplace-Beltrami spectra as 'Shape-DNA' of surfaces and solids. *Computer-Aided Design* **38** (2006) 342–66
13. Protter, M.: Can one hear the shape of a drum? Revisited. *SIAM Review* **29**(2) (June 1987) 185 – 197
14. Lee, D.D., Seung, H.S.: Algorithms for non-negative matrix factorization. In: NIPS. (2000) 556–562
15. Heiberg, E., Sjogren, J., Ugander, M., Carlsson, M., Engblom, H., Arheden, H.: Design and validation of segment - freely available software for cardiovascular image analysis. *BMC Medical Imaging* **10**(1) (2010) 1
16. Ou, Y., Sotiras, A., Paragios, N., Davatzikos, C.: Dramms: Deformable registration via attribute matching and mutual-saliency weighting. *Medical Image Analysis* **15**(4) (2011) 622–639
17. Burges, C.J.C.: A tutorial on support vector machines for pattern recognition. *Data Mining and Knowledge Discovery* **2** (1998) 121–167