Spatially Coherent Filtering of Multilead ECG via Laplace-Beltrami Eigenfunctions Decomposition

Ismael Hernández-Romero[1,2], Jorge Vicente-Puig[2], Ernesto Zacur[2], David Lundback[2], Andreu M Climent[1,2], Maria S Guillem[1,2]

[1]COR lab, ITACA Institute, Universitat Politècnica de València, Spain
[2]Corify Care SL, Spain

Abstract

Multilead ECG signals exhibit spatial coherence that is often disregarded when each lead is filtered independently. This work proposes a spatially coherent filtering approach based on Laplace-Beltrami eigenfunction decomposition, which projects the signal onto a basis derived from torso geometry. Filtering is then performed in the spectral domain using the same temporal operations as the lead-wise method. Evaluation on realistic synthetic BSPM data showed consistent improvements across noise levels. At 5 dB SNR, the proposed method reduced RMSE by 64% and increased correlation by 24% relative to the unfiltered signal and outperformed lead-wise filtering by 31% in RMSE and 5% in correlation. These results suggest that incorporating spatial structure into multilead ECG filtering enhances noise robustness without modifying the temporal pipeline.

1. Introduction

The acquisition and analysis of multilead electrocardiogram (ECG) recordings from the body surface, commonly employed in Body Surface Potential Mapping (BSPM) and non-invasive electrocardiographic imaging (ECGI), play a key role in characterizing cardiac electrical activity. However, the diagnostic utility of these recordings is frequently compromised by noise contamination, which can obscure relevant features. A significant limitation of conventional filtering approaches is that they typically operate on each ECG lead independently, disregarding the intrinsic spatial coherence across channels [1]. This coherence emerges because all leads capture projections of the same underlying electrical sources, shaped by the geometry and conductive properties of the torso [2]. Ignoring these spatial dependencies can distort the signal and degrade filtering performance.

Multivariate methods such as Principal Component Analysis (PCA) and Independent Component Analysis (ICA) attempt to address this limitation by processing multiple leads simultaneously [1]. However, these techniques rely on statistical criteria (e.g., variance maximization or independence) rather than explicitly incorporating the spatial structure defined by electrode placement.

To overcome this limitation, we propose a novel approach grounded in differential geometry. Inspired by recent advances in signal decomposition on non-Euclidean domains, particularly in neuroscience, we adopt the eigenfunctions of the Laplace-Beltrami (LB) operator as a spatial basis derived from the torso geometry [3–5]. These eigenfunctions form a natural set of spatial modes ordered by complexity, providing a principled framework to represent and process signals in a geometry-aware framework.

This study evaluates an LB-based filtering methodology that projects multilead ECG data on this intrinsic spatial basis. The preprocessing is then applied in the LB spectral domain before reconstructing the signals in the original electrode space. Using realistic synthetic ECG data, we compare this approach to standard lead-wise filtering techniques, demonstrating its ability to improve signal quality while preserving spatial integrity. This is an essential step toward more robust and accurate non-invasive cardiac assessment.

2. Methods

2.1. Laplace-Beltrami Eigenfunction Basis

The core of the proposed approach lies in representing the multilead ECG signal using a basis derived from the intrinsic geometry of the electrode arrangement on the torso surface. This is formalized through the Laplace-Beltrami (LB) operator, denoted as Δ, defined on the manifold $\mathcal{M}$ representing the torso surface.

The eigenfunctions ψ_j of the LB operator satisfy the Helmholtz equation:

 ISSN: 2325-887X DOI: 10.22489/CinC.2025.261

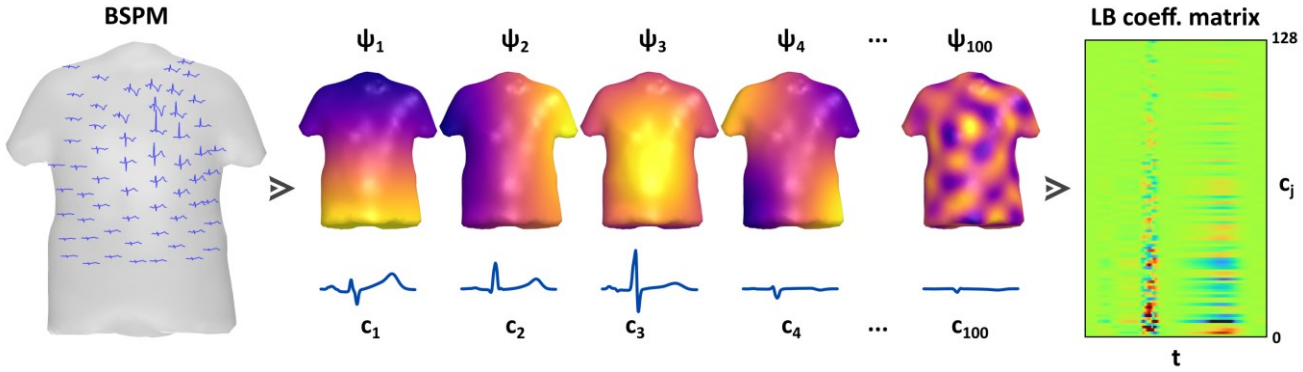

Figure 1. Laplace-Beltrami decomposition of a BSPM. Multilead ECG signals, selected eigenfunctions ψ_j, associated coefficients c_j, and full spectral LB coefficient matrix representation.

$$\Delta\psi_j(x) = \lambda_j \psi_j(x), \qquad x \in \mathcal{M} \qquad (1)$$

where λ_j are real, non-negative eigenvalues ordered as $0 = \lambda_0 \leq \lambda_1 \leq \ldots$. These eigenfunctions form a complete orthonormal basis for square-integrable functions on $\mathcal{M}$, ordered by increasing spatial frequency.

Any scalar field f(x) defined on the surface, such as the BSPM, can be represented as a linear combination of the LB eigenfunctions:

$$f(x) = \Sigma\, c_j\, \psi_j(x) \qquad (2)$$

where the coefficients $c_j = \langle f, \psi_j \rangle$ denote the projection of f onto each eigenfunction ψ_j.

For practical implementation, the continuous LB operator Δ is approximated by a discrete graph Laplacian matrix $L \in \mathbb{R}^{N \times N}$, constructed using the cotangent-weight scheme in the torso surface with N nodes. To approximate the continuous eigenfunctions defined in Eq. (1), we solve the discrete eigenvalue problem $LV = V\Lambda$, where Λ is the diagonal matrix of eigenvalues λ_j, and the columns of V are the discrete eigenvectors approximating the values of ψ_j at the electrode locations. The number of eigenvectors is set to match the number of active electrodes (i.e., 128 channels).

The resulting matrix $\Psi = V \in \mathbb{R}^{N \times N}$ defines an orthonormal basis. Interpolating the multilead ECG signals to the entire torso [6], the matrix $BSPM \in \mathbb{R}^{N \times T}$ is obtained, where T is the number of time samples. The signals can be projected onto the LB basis to get a coefficient matrix $C \in \mathbb{R}^{N \times T}$:

$$C = \Psi\dagger\, BSPM \qquad (3)$$

where $\Psi\dagger$ denotes the Moore-Penrose pseudoinverse of Ψ. Figure 1 illustrates this decomposition, showing an example BSPM, selected LB basis functions, and the resulting coefficient matrix C.

The original signals can be reconstructed from the coefficient domain as:

$$BSPM = \Psi\, C \qquad (4)$$

This framework allows signal processing to be carried out in the LB spectral domain (C), exploiting spatial regularity, before mapping back to the body surface space (BSPM).

2.2. Filtering Strategies for BSPM

We compared the proposed spatially coherent filtering approach with a conventional lead-wise method. Both applied the same preprocessing steps—baseline removal and bandpass filtering—but in different domains.

In the LB-based method, the BSPM was projected onto the LB eigenfunction basis. Each resulting spectral coefficient was processed with a zero-phase 4th-order Butterworth filter (0.5–70 Hz) and corrected for baseline wander using a moving median filter (0.5-second window, 50% overlap). The denoised signal was then reconstructed in the electrode space via the inverse projection.

In the lead-wise method, each signal was filtered independently, applying the same temporal filtering and baseline correction directly in the electrode domain. This setup allows isolating the effect of operating in the LB spectral domain, while keeping the filtering logic identical.

2.3. Synthetic Dataset Generation

Synthetic BSPM signals were generated using the FECGSYN framework [7], originally developed for maternal-fetal ECG modeling, and adapted here to simulate torso-projected cardiac activity. A dynamic dipole was projected onto a realistic torso mesh to synthesize a 128-lead BSPM over 10 seconds at 1000 Hz, with the clean signals used as ground truth.

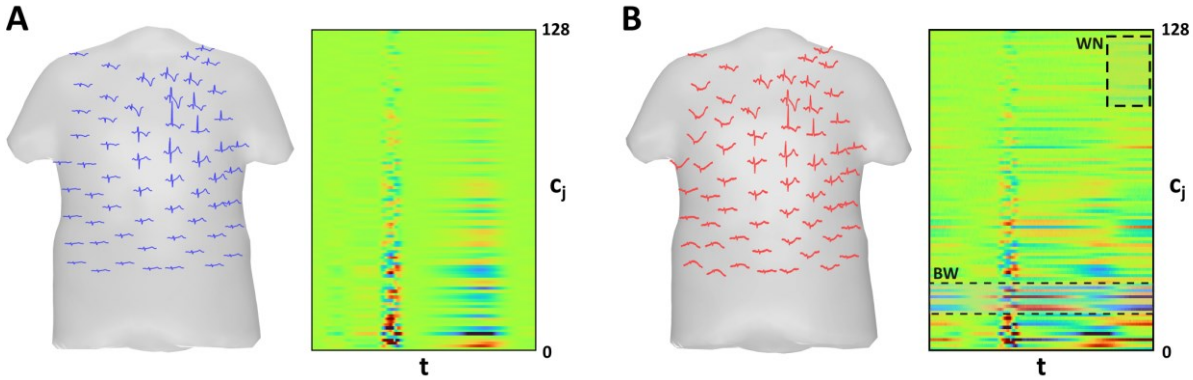

Figure 2. LB decomposition of a BSPM under two conditions. (A) Clean signals yield structured coefficients concentrated in low-order modes. (B) Noisy signals show disrupted patterns due to baseline wander (BW) and white noise (WN), affecting low and high modes, respectively.

To emulate realistic recording conditions, spatially coherent baseline wander (BW) and uncorrelated white Gaussian noise (WN) were superimposed. These noise components were combined and added to the clean signals to generate three noisy datasets at SNR levels of –5, 5, and 15 dB, enabling controlled comparisons between filtering methods.

2.4. Evaluation Metrics

Filtering performance was quantitatively evaluated using Root Mean Square Error (RMSE) and Pearson Correlation Coefficient (CC), which respectively measure the magnitude of reconstruction error and the temporal similarity between the filtered and clean signals across channels. Both metrics were computed over all the signals and averaged across all leads.

3. Results

3.1. Analysis of LB Coefficients and Reconstructed Signals

To illustrate how spatial coherence and noise manifest in the LB spectral domain, Figure 2 shows a representative BSPM signal and its corresponding LB coefficients under two conditions. Panel A depicts the clean signal, exhibiting high spatial coherence and well-structured coefficients with energy concentrated in low-frequency modes. Panel B shows the same signal corrupted with spatially coherent BW and uncorrelated WN, resulting in dispersed and temporally irregular coefficients. Highlighted regions illustrate the impact of BW on low-order modes and WN on higher-order components.

Figure 3 illustrates BSPM signals from all leads under SNR = 5 dB to compare reconstruction quality across methods. Panel A shows the noisy signals (light blue) and the clean reference (dark blue), highlighting the impact of added noise. Panel B displays lead-wise filtered signals (orange), while Panel C shows LB-filtered signals (green), both overlaid with the clean reference. Visually, LB filtering better preserves the waveform morphology and amplitude consistency across leads.

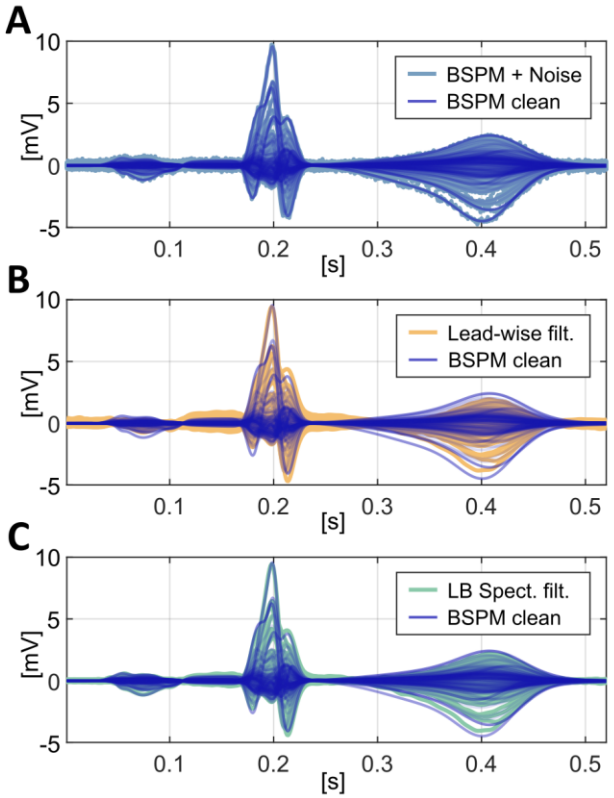

Figure 3. Visual comparison of BSPM signals across all leads under SNR = 5 dB. (A) Noisy input, (B) Lead-wise filtered signals, and (C) LB spectral filtered signals vs. clean references.

3.2. Quantitative Evaluation

Filtering performance was assessed at three representative SNR levels (–5, 5 and 15 dB) using RMSE and Pearson CC. The results are shown in Figure 4.

Both the LB-based and the benchmark lead-wise methods share similar preprocessing steps—baseline removal and bandpass filtering—but operate in different domains. Across all tested noise levels, the LB-based method consistently achieved lower RMSE and higher correlation, although the difference was moderate.

At –5 dB, RMSE was 3.53 for the lead-wise method and 3.07 for the LB-based approach, while CC increased from 0.72 to 0.78. At 5 dB, RMSE was 1.73 vs. 1.20, and CC improved from 0.91 to 0.96. At 15 dB, RMSE was 1.39 vs. 0.69, and CC increased from 0.94 to 0.99.

These results indicate that, while both methods apply similar temporal filtering strategies, the use of the LB basis provides consistent improvements in denoising performance, particularly under moderate and low SNR conditions.

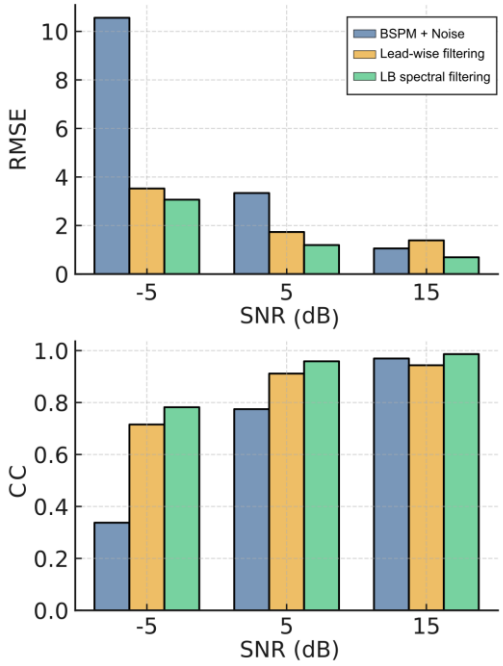

Figure 4. Filtering performance at three SNR levels (-5, 5 and 15 dB) using RMSE and Pearson CC.

4. Discussion

This study explored a spatially coherent filtering strategy for multilead ECG using Laplace-Beltrami eigenfunction decomposition. By projecting the signal onto a basis derived from torso geometry, the method captures spatial coherence across leads, which is often overlooked by standard filtering approaches. The LB-based method showed lower error and better correlation with the clean signal, especially under low SNR conditions. Compared to lead-wise filtering and methods like PCA or ICA, the LB basis aligns with electrode placement, providing a more structured representation. Beyond denoising, similar to its use in neuroimaging, LB decomposition could be employed to extract dominant spatial modes, characterize activation gradients, or reduce dimensionality while preserving spatial organization [4-5]. Future work will focus on clinical validation and adaptive strategies for mode selection and spectral filtering, offering a robust and anatomically grounded approach for multilead ECG processing.

5. Conclusion

A spatially coherent filtering method for multilead ECG was presented using Laplace-Beltrami eigenfunctions. Filtering on this spatial basis preserves inter-lead coherence and consistently outperformed lead-wise approaches, especially at low SNR. The framework also offers a structured representation of BSPM signals that may benefit further analysis. Future work will address clinical validation and broader applications.

Acknowledgments

This work was supported by grants CPP2021-008562, CPP2023-01050, PID2023-149812OB-I00, CNS2022-135512, Torres Quevedo PTQ2023-013018 funded by MCIN/AEI/10.13039/501100011033 and by "European Union NextGenerationEU/PRTR"; and funded by Generalitat Valenciana CIAICO/2022/020 and by "ESF Investing in your future".

References

[1] Romero I. PCA and ICA applied to noise reduction in multi-lead ECG. *Comput Cardiol* 2011;38:613-616.

[2] Sameni R, et al. Multichannel ECG and noise modeling: Application to maternal and fetal ECG signals. *EURASIP J Adv Signal Process* 2007;2007:43407.

[3] Atasoy S, et al. Human brain networks function in connectome-specific harmonic waves. *Nat Commun* 2016;7:10340.

[4] Lioi G, et al. Gradients of connectivity as graph Fourier bases of brain activity. *Netw Neurosci* 2021;5(2):322-336.

[5] Huang SG, et al. Revisiting convolutional neural network on graphs with polynomial approximations of Laplace-Beltrami spectral filtering. *Neural Comput Appl* 2021;33:13693-13704.

[6] Oostendorp T, et al. Interpolation on a triangulated 3D surface. *J Comput Phys* 1989;80:331-343.

[7] Behar J, et al. An ECG simulator for generating maternal-foetal activity mixtures on abdominal ECG recordings. *Physiol Meas*. 2014;35(8):1537.

Address for correspondence: ismael.hernandez@corifycare.com