Available online at www.eccomasproceedia.org
Eccomas Proceedia UNCECOMP (2025) 476-489

ECCOMAS
Proceedia

UNCECOMP 2025
6th ECCOMAS Thematic Conference on
Uncertainty Quantification in Computational Science and Engineering
M. Papadrakakis, V. Papadopoulos, G. Stefanou (eds.)
Rhodes Island, Greece, 15-18 June 2025

LABEL-EFFICIENT CARDIAC VOLUME ESTIMATION IN REAL-TIME MRI WITH SPATIAL CORRELATION

Felix Terhag[1,2], Philipp Knechtges[1], Anja Bach[4], Darius Gerlach[4], Jens Tank[4], Achim Basermann[1] and Raúl Tempone[2,3]

[1]Institute of Software Technology, German Aerospace Center (DLR)
Linder Hoehe, 51147, Cologne, Germany
e-mail: {felix.terhag, philipp.knechtges, achim.basermann}@dlr.de

[2]Chair of Mathematics for Uncertainty Quantificatio
RWTH Aachen University, Aachen, Germany
e-mail: tempone@uq.rwth-aachen.de

[3]Computer, Electrical and Mathematical Sciences and Engineering Division (CEMSE)
King Abdullah University of Science and Technology (KAUST), Thuwal, Saudi Arabia

[4]Institute of Aerospace Medicine, German Aerospace Center (DLR)
Linder Hoehe, 51147, Cologne, Germany
e-mail: {anja.bach, darius.gerlach, jens.tank}@dlr.de

Abstract. *Real-time magnetic resonance imaging (MRI) enables imaging under free-breathing conditions, allowing researchers to study the influence of breathing on the heartbeat. However, this advantage requires handling large amounts of data, exceeding 4,500 images per heart, making manual segmentation of the ventricular volume highly labor-intensive and costly.Hence, maximizing the information gained from each labeled image is essential to improve label efficiency.*

We propose a Bayesian model for predicting ventricular volumes with minimal manual labeling. Unlike standard approaches that treat each slice independently, the proposed model incorporates spatial correlations by enforcing smoothness via a discretized one-dimensional Laplace operator, avoiding strong geometric assumptions. Additionally, this study exploits conjugacy to compute the posterior analytically, eliminating the need for sampling or variational inference and allowing for efficient hyperparameter optimization via the type-II maximum likelihood. This approach enables the proposed model to employ spatial information effectively, improving predictions for slices with few labeled images.

This work compares the proposed model to a baseline model that treats each slice independently. In the model evaluation on real patient data, the proposed method outperformed the baseline, particularly in data-scarce scenarios, where spatial correlation considerably enhances prediction accuracy.

Keywords: Bayesian Modeling, Real-time Magnetic Resonance Imaging, Empirical Bayes, Label Efficien y, Maximum Likelihood Estimation

ISSN:2623-3339
Peer-review under responsibility of the organizing committee of UNCECOMP 2025.
doi: 10.7712/120225.12371.21170

1 INTRODUCTION

Understanding how breathing influences the heartbeat is of great interest in cardiac research. However, observing this relationship is challenging because traditional cine magnetic resonance imaging (MRI) is typically performed using a breath-holding protocol. With the emergence of real-time MRI, which can capture up to 50 frames per second (fps), breath-holding is no longer required, allowing for the direct observation of the effects of breathing on the heartbeat [14, 9, 13].

An important parameter in cardiac research is the ventricular volume over time, requiring ventricle segmentation in each image. However, real-time MRI is challenging because it generates significantly more data. For instance, the study dataset, captured at 30 fps, requires at least 10 s per slice to record about one to four breathing cycles, depending on the breathing speed. With a minimum of 15 slices per heart, this results in over 4,500 images requiring annotation, making manual segmentation highly labor-intensive.

Furthermore, the study dataset comprises patients with univentricular hearts, a rare congenital heart disease in which patients are born with a single functioning ventricle and who undergo multiple surgical interventions that significantly alter cardiac anatomy. This situation increases the difficulty of applying automatic segmentation. A more detailed description of the study that acquired the data is presented in [8].

The large volume of data and high segmentation cost require a label-efficient approach. Previous work has modeled the ventricular volume over time as a superposition of frequencies and employed sparse Bayesian learning to identify a sparse set of relevant frequencies [11]. Fewer labels are required to fit the model because only a few frequencies are necessary to describe the dynamics. Building upon this framework, we aim to incorporate spatial information while accounting for the complex and heterogeneous geometries of univentricular hearts.

Given this variability, strict geometric assumptions are minimized. Instead, this work applies the discretized one-dimensional (1D) Laplace operator to enforce smoothness in a flexible manner. This study constructs a Bayesian model incorporating these geometric correlations into the prior, ensuring that the likelihood dominates when sufficient data are available. This way, the proposed model enforces smoothness when necessary but gives more weight to the data when sufficient observations are available.

Furthermore, the proposed model utilizes conjugacy to compute the posterior analytically. This approach makes the procedure more efficient eliminating the need for sampling or variational inference to approximate the posterior. Moreover, the proposed model enables direct computation of the marginal likelihood gradient with respect to the hyperparameters. Thus, gradient-based methods, such as BFGS , can efficiently determine the maximum of the type-II likelihood.

In [11], a similar model is applied to real-time MRI without accounting for spatial correlations. Similar models are also widely applied in remote sensing applications [4, 5]. In [15], the authors proposed a related approach for signal recovery incorporating correlation, though in the time domain. However, their method assumes all sources, analogous to slices in our case, are mutually independent. While modeling correlations along the slice dimension is less relevant in remote sensing, it is crucial for real-time MRI applications.

2 METHODS

2.1 Modeling

In contrast to cine MRI, real-time MRI allows video sequences of the beating heart to be retrieved due to a faster capture rate. The heart is captured in L slices, where every slice $\ell \in 1, \dots, L$ comprises a video of N images. The images are captured with a constant capture rate at time steps $t_1, \dots, t_N$. This work is interested in the ventricle volume over time and assumes that it can be described as a superposition of $m+1$ frequencies $\mathcal{F} = \{f_0 = 0, f_1, f_2, \dots, f_m\}$, where the main contributing frequencies are the heart and respiratory rate. Limiting the number of frequencies m benefit parameter estimation, and $M := 2m+1$ is introduced to facilitate the notation. A sparse set of frequencies that adequately describes the data is determined using sparse Bayesian learning [11]. The proposed linear model is

$$Y = AX + N, \tag{1}$$

with

- $Y \in \mathbb{R}^{N \times L}$ the real-valued measurement points $y_{k,\ell}$,
- $X \in \mathbb{R}^{M \times L}$ the real and imaginary parts of the complex amplitudes,
- $N \in \mathbb{R}^{N \times L}$ additive noise with $n_{k,\ell} \sim \mathcal{N}(0, \sigma^2)$,
- $A \in \mathbb{R}^{N \times M}$ the transfer matrix $A = [\mathbf{1}_N \; Re(B) \; Im(B)]$, and
- with $\mathbf{1}_N$ an N-dimensional vector of ones and $B \in \mathbb{C}^{N \times m}$ define as,

$$B = \begin{bmatrix} \exp(-i2\pi f_1 t_1) & \dots & \exp(-i2\pi f_m t_1) \\ \vdots & \ddots & \vdots \\ \exp(-i2\pi f_1 t_N) & \dots & \exp(-i2\pi f_m t_N) \end{bmatrix}.$$

Prior distribution In [11], the prior for X is a zero-centered normal distribution, independent for all slices and frequencies. In contrast, this work assumes that the magnitudes of the frequencies depend on neighboring slices. Hence, the following matrix is introduced:

$$T = \begin{bmatrix} 2 & -1 & & & \\ -1 & 2 & -1 & & \\ & \ddots & \ddots & \ddots & \\ & & -1 & 2 & -1 \\ & & & -1 & 2 \end{bmatrix} \in \mathbb{R}^{L \times L},$$

which is a 1D equidistant discretization of the Laplace operator with Dirichlet-boundary conditions. The rationale is that without further geometric information, we prefer to select a prior that penalizes excessive curvature (i.e., the (discretized) second derivative) over the slices. The prior should be selected to resemble $X \sim ZT^{-1}$ with every column i of Z being normally distributed $Z_i \sim \mathcal{N}(0, sI_M)$ with variance $s > 0$ and I_M being the identity matrix of size M. With the vectorization operator $\mathrm{vec}(\cdot)$, we obtain

$$p(\mathrm{vec}(Z)|s) = \mathcal{N}(\mathrm{vec}(Z)|0, sI_{ML}),$$

where $\mathcal{N}(x|\mu,\Sigma)$ denotes the probability density function of a multivariate normal distribution at location x for the mean μ and covariance matrix Σ. The identity

$$\begin{aligned}\text{vec}(ZT^{-1}) =&(T^{-T}\otimes I_M)\text{vec}(Z)\\ =&(T^{-1}\otimes I_M)\text{vec}(Z),\end{aligned}$$

yields the prior $p(\text{vec}(X)|s)$, fulfillin the demands, where

$$\begin{aligned}p(\text{vec}(X)|s) &= \mathcal{N}(\text{vec}(X)|0,(T^{-1}\otimes I_M)sI_{ML}(T^{-T}\otimes I_M))\\ &= \mathcal{N}(\text{vec}(X)|0,T^{-2}\otimes sI_M)\end{aligned} \tag{2}$$

using the mixed product property of the Kronecker product $\otimes$ and the fact that $(T^{-1}\otimes I_M)$ is symmetric.

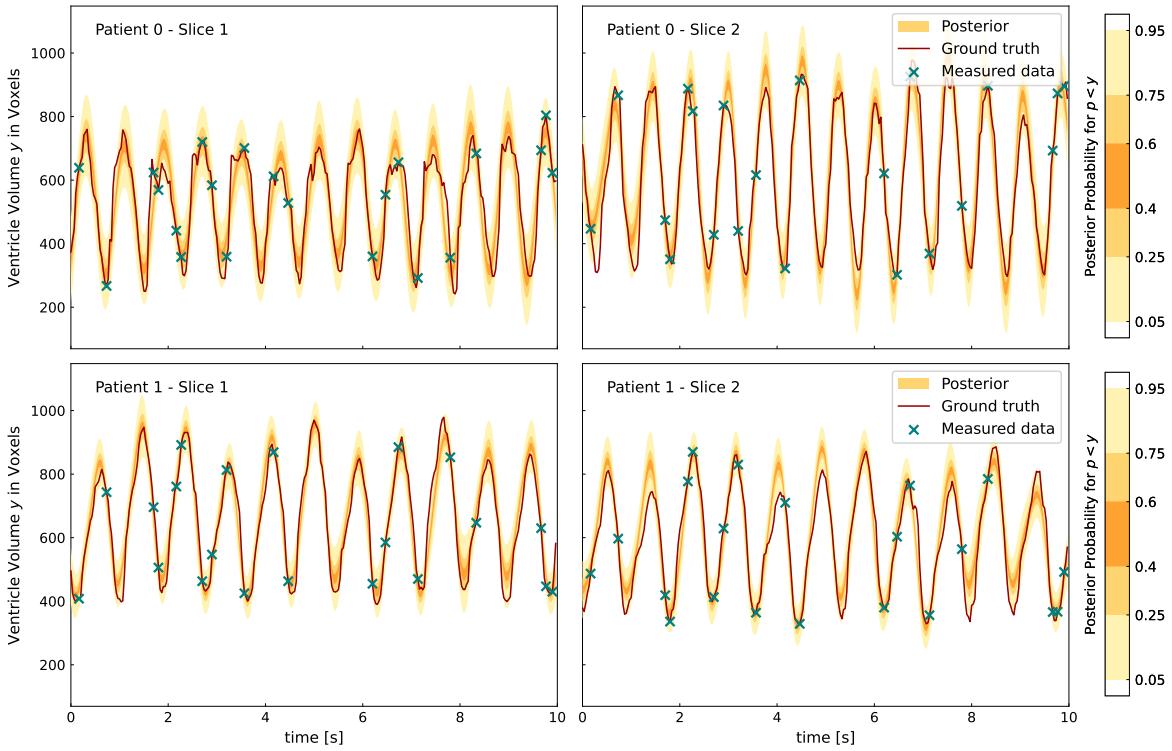

Figure 1: Ventricle volume over time for two consecutive slices in both patients. Images were captured at 30 fps, with each slice spanning 300 frames. The red curve represents the ground truth from expert-validated segmentations. The yellow curve marks the posterior predictive, based on 20 measurements per slice, indicated by a green x.

Likelihood For partially observed data Y, let $J \subset \{1,\dots,NL\}$ be the index set of the observed data in vectorized form and $P \in \mathbb{R}^{|J|\times NL}$ be the projection matrix, with entries $P_{i,J_i} = 1$ for $i \in 1,\dots,|J|$ and zero elsewhere. For the model define in (1), the likelihood of data Y given X and σ^2 is

$$p(\text{vec}(Y)|X,\sigma^2) = \frac{\exp\left(-\frac{1}{2\sigma^2}\,\|P(I_L\otimes A)\text{vec}(X)-\text{vec}(Y)\|_2^2\right)}{(2\pi\sigma^2)^{|J|/2}}.$$

We defin $G := P(I_L \otimes A) \in \mathbb{R}^{|J| \times ML}$ to simplify this formulation, yielding

$$p(\text{vec}(Y)|X, \sigma^2) = \frac{\exp\left(-\frac{1}{2\sigma^2} \|G\text{vec}(X) - \text{vec}(Y)\|_2^2\right)}{(2\pi\sigma^2)^{|J|/2}}. \tag{3}$$

Posterior distribution Given the likelihood (3) and prior distribution (2), assuming the parameters s and σ^2 are known the posterior of X can be calculated using the Bayes rule

$$p(X|Y, s, \sigma^2) := \frac{p(\text{vec}(Y)|X, \sigma^2)p(\text{vec}(X)|s)}{p(Y|s, \sigma^2)}. \tag{4}$$

The likelihood and prior distribution are normal distributions, allowing the analytical computation of the normal posterior distribution of X with mean μ_X and covariance Σ_X given by

$$\begin{aligned} \Sigma_X &= \left(\frac{1}{\sigma^2} G^T G + (T^{-2} \otimes sI_M)^{-1}\right)^{-1} \\ \mu_X &= \frac{1}{\sigma^2} \Sigma_X G^T \text{vec}(Y). \end{aligned} \tag{5}$$

The posterior predictive distribution of $\text{vec}(Y)$ can be directly derived from the vectorized form of (1).Thus, the distribution is normal distribution with mean μ_Y and covariance Σ_Y

$$\begin{aligned} \Sigma_Y &= (I_L \otimes A)\Sigma_X (I_L \otimes A)^T + \sigma^2 I_{NL} \\ \mu_Y &= (I_L \otimes A)\mu_X. \end{aligned} \tag{6}$$

Figure 1 presents the posterior distribution given 20 measurement points per slice. While it does not perfectly capture the ground truth, it performs well considering that only 20 out of 300 measurements per slice are used, demonstrating that the model is effective in this setting.

2.2 Empirical Bayes

The proposed model treats the parameters σ^2 and s as hyperparameters. Instead of computing the joint posterior $p(X, \sigma^2, s|Y)$, the conditional posterior (4) is used to estimate $\hat{\sigma}^2$ and $\hat{s}$. The hyperparameters are determined by maximizing the marginal likelihood

$$\hat{\sigma}^2, \hat{s} = \arg\max_{\sigma^2, s} p(Y|\sigma^2, s), \tag{7}$$

a procedure known as type-II-maximum likelihood optimization or empirical Bayes [2, 7]. In this case, the marginal likelihood is a product of the prior (2) and likelihood (3) integrated over X. Applying the data covariance $\Sigma_p := \sigma^2 I_{|J|} + G(T^{-2} \otimes sI_M)G^T$ and the Woodbury formula yields

$$p(Y|\sigma^2, s) = \int p(Y|X, \sigma^2)p(X|s)dX \propto \frac{\exp\left(-\frac{1}{2}\text{vec}(Y)^T \Sigma_p^{-1} \text{vec}(Y)\right)}{\sqrt{\det(\Sigma_p)}}. \tag{8}$$

Hence, the log-marginal likelihood for some constant value c is given by

$$\log p(Y|\sigma^2, s) = -\text{vec}(Y)^T \Sigma_p^{-1} \text{vec}(Y) - \log\det(\Sigma_p) + c.$$

Next, we can estimate $\hat{\sigma}^2$ and $\hat{s}$ in (7) by minimizing the following loss function

$$\mathcal{L}(\sigma^2, s) = \text{vec}(Y)^T \Sigma_p^{-1} \text{vec}(Y) + \log\det(\Sigma_p). \tag{9}$$

To minimize $\mathcal{L}$ using a gradient based approach, we need to calculate the derivative of $\mathcal{L}$. Using the formula for the derivative of the matrix inverse, we can calculate

$$\begin{aligned}\frac{\partial \Sigma_p^{-1}}{\partial s} &= -\Sigma_p^{-1}\frac{\partial \Sigma_p}{\partial s}\Sigma_p^{-1} \\ &= -\Sigma_p^{-1}G(T^{-2}\otimes I_M)G^T\Sigma_p^{-1}\end{aligned}$$

and

$$\frac{\partial \Sigma_p^{-1}}{\partial \sigma^2} = -\Sigma_p^{-2}.$$

Furthermore, Jacobi's formula results in

$$\frac{\partial \log\det(\Sigma_p)}{\partial s} = \operatorname{tr}\left(\Sigma_p^{-1}\frac{\partial \Sigma_p}{\partial s}\right) = \operatorname{tr}\left(\Sigma_p^{-1}G(T^{-2}\otimes I_M)G^T\right)$$

and analogously,

$$\frac{\partial \log\det(\Sigma_p)}{\partial \sigma^2} = \operatorname{tr}\left(\Sigma_p^{-1}\right).$$

Combining these results leads to the following derivatives

$$\begin{aligned}\frac{\partial \mathcal{L}(\sigma^2, s)}{\partial s} &= \operatorname{tr}\left(\Sigma_p^{-1}G(T^{-2}\otimes I_M)G^T\right) - \operatorname{vec}(Y)^T\Sigma_p^{-1}G(T^{-2}\otimes I_M)G^T\Sigma_p^{-1}\operatorname{vec}(Y) \\ \frac{\partial \mathcal{L}(\sigma^2, s)}{\partial \sigma^2} &= \operatorname{tr}\left(\Sigma_p^{-1}\right) - \operatorname{vec}(Y)^T\Sigma_p^{-2}\operatorname{vec}(Y),\end{aligned} \tag{10}$$

allowing the use of gradient-based methods, such as the BFGS-algorithm [1, 3, 6, 10].

3 RESULTS

Real-time MRI is an emerging technology that captures images at up to 50 fps, allowing the observation of respiratory effects on the heartbeat by removing the need for the breath-holding protocol in traditional cine MRI [14, 9, 13]. The volume of the ventricle over time - ventricular fillin - is of clinical relevance and is used to estimate importantparameters like stroke volume. An expert must segment the ventricle to derive the ventricular filling which is expensive and tedious. Real-time MRI produces considerable data; thus, more images must be segmented. The study dataset comprises data from two patients, with four consecutive slices of 300 images completely labeled. This work applies empirical Bayes to estimate hyperparameters s and σ^2 by maximizing the marginal likelihood (8), corresponding to minimizing the loss function (9). We employed gradient-based optimization methods using the gradient (10). The resulting optimization identifie the minimum of the loss function, as illustrated in the loss landscape in Figure 2.

To assess the stability of the proposed method, we employed a Laplace approximation, which approximates the log-marginal likelihood with a normal distribution centered at its maximum $\hat{\sigma}^2, \hat{s}$. The covariance of this normal distribution is given by the inverse Hessian H^{-1} of the log-marginal likelihood evaluated at $(\hat{\sigma}^2, \hat{s})$. This method results in the following normal approximation:

$$\mathcal{N}\left((\sigma^2, s) \mid (\hat{\sigma}^2, \hat{s}), H^{-1}\right). \tag{11}$$

We numerically approximated the Hessian and applied it to compute the 95% confidenc region in Figure 2. To further evaluate the stability of the proposed method, we sampled 50 pairs

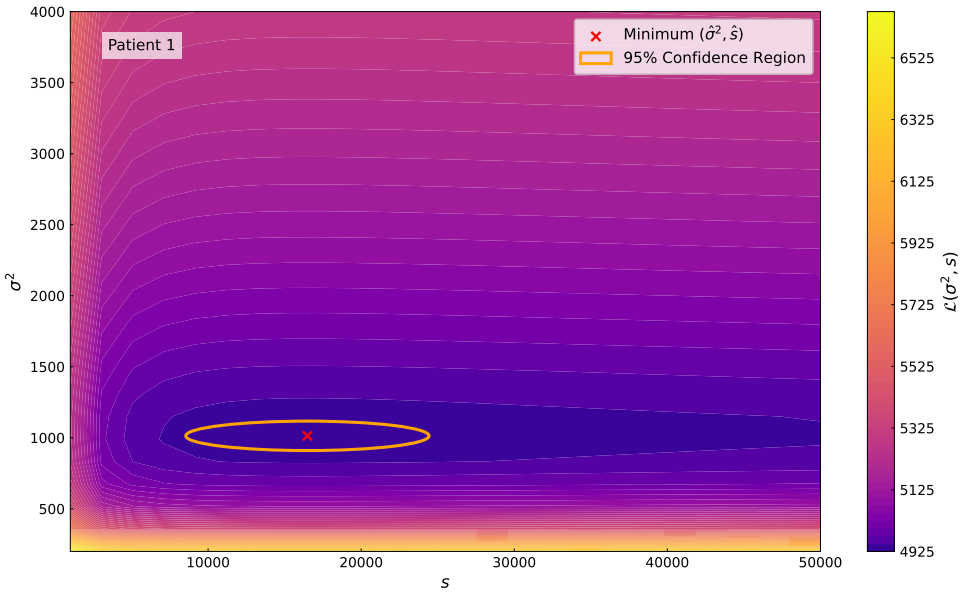

Figure 2: Loss landscape of (9) for given s and σ^2. The loss is equivalent to the negative log-marginal likelihood (8) up to a constant. The red x marks the minimum determined using the BFGS algorithm. The ellipsoid corresponds to the 95% confidenc region.

(σ^2, s) from the approximated normal distribution and analyzed the information variation (12). The observed variance in information loss was consistently low, suggesting that the method remains stable in the most likely region of (σ^2, s), see Figure 3.

The proposed method employs spatial correlation between slices to improve the labeling efficien y, which is achieved by incorporating the 1D Laplace discretization T into the prior. To evaluate this approach, we compared the model to one that disregards spatial correlation, treating slices independently, corresponding to setting $T = I_L$ as the identity matrix. We call this independent model the *baseline model*. Performance is measured using the *information loss*, define as the discrepancy in the negative log-likelihood of the ground-truthdata between the posterior with full data and that with only partially observed data. Let index set $J \subset 1, ..., NL$ be the observed data and μ_{Y_L}, Σ_{Y_L} the mean and covariance of the posterior (6) given this data. If μ_Y, Σ_Y are the mean and variance of the posterior with all data, $J = 1, ..., NL$, then the information loss is define as

$$D(J, Y) = \log(\mathcal{N}(Y|\mu_Y, \Sigma_Y)) - \log(\mathcal{N}(Y|\mu_{Y_L}, \Sigma_{Y_L})). \tag{12}$$

For better comparability, we used the same hyperparameters, s and σ^2, for both models. The difference in negative log-likelihood for the complete data is negligible for the spatial and baseline models. The difference is smaller than 0.02, a relative difference of $3 \cdot 10^{-6}$. This result suggests that when data are numerous, the influenc of the different prior is negligible because the likelihood is dominant.

Two scenarios were tested to assess the performance of the proposed approach. Using neural networks for segmentation can produce reliable results; however, it could fail for single slices [11, 12], requiring experts to label these slices manually. To replicate this behavior in the firs scenario, we applied jackknife resampling, omitting one slice while assuming the remaining

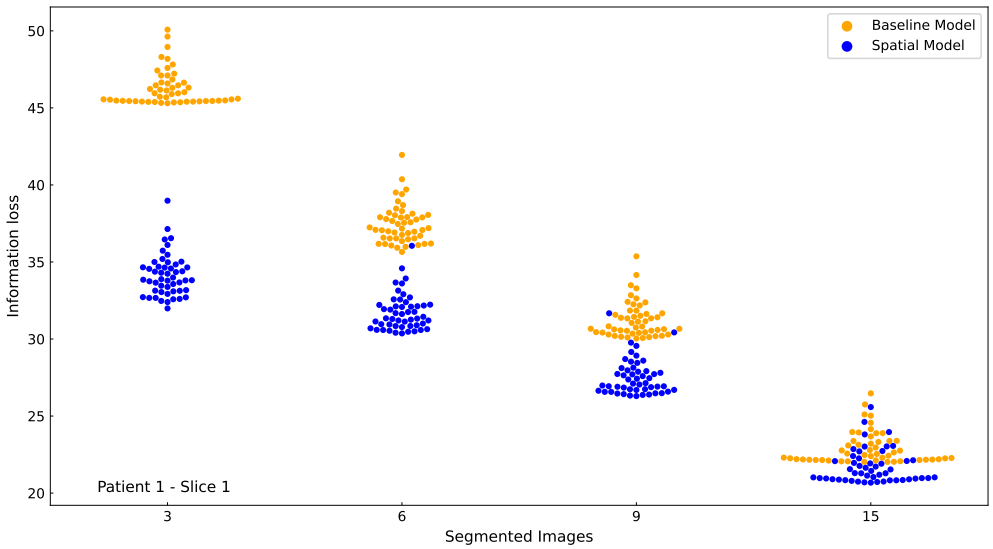

Figure 3: Information loss between the model trained on all data and each of the baseline and spatial models, evaluated for 50 sampled pairs (s, σ^2) from the normal approximation of the marginal log-likelihood around $(\hat{\sigma}^2, \hat{s})$. Both models were evaluated for varying numbers of segmented images for Patient 1 and Slice 1. The sample variance is low, indicating the stability of the proposed modeling approach with respect to σ^2 and s.

slices were fully segmented. Then, we measured the information loss on the left-out slice and iteratively added newly segmented images. The segmentation order was determined by greedily selecting indices that minimized the trace of the posterior covariance, following [11]. Figure 4 illustrates that the spatial model consistently outperforms the baseline across all slices. The spatial model retains more information than the baseline when most measurements from a slice are omitted and quickly recovers information from just a few labeled examples. This result suggests that the built-in spatial prior effectively transfers information from neighboring slices. Notably, this effect weakens after around 15 segmented images, as the likelihood (identical for both models) begins to dominate the prior.

To further illustrate our method, we present the posterior distribution for both patients using three and seven measurements to illustrate the method. In Figure 5, the advantage of the proposed approach is evident. Even with only three measurements, the posterior qualitatively aligns with the ground truth, whereas the baseline model fails to capture it. This result demonstrates how the spatial model effectively uses information from neighboring slices. This trend is consistent across all slices, with full results in the appendix in Figure 7.

In the second scenario, where this modeling approach is beneficial no images are labeled, and a practitioner must label these images individually. A reliable prediction with as little data as possible is beneficial We started with no labeled images for this scenario and randomly measured up to 42. Thus, in the end, 42 of the 1,200 images of the full dataset were labeled. We employed the same random seeds for the baseline and spatial models to maintain comparability. The procedure was repeated 300 times (Figure 6).

Incorporating the spatial correlation into the model is beneficia in this comparison, especially for a few labeled images. Each measurement point can influenc the posterior in all

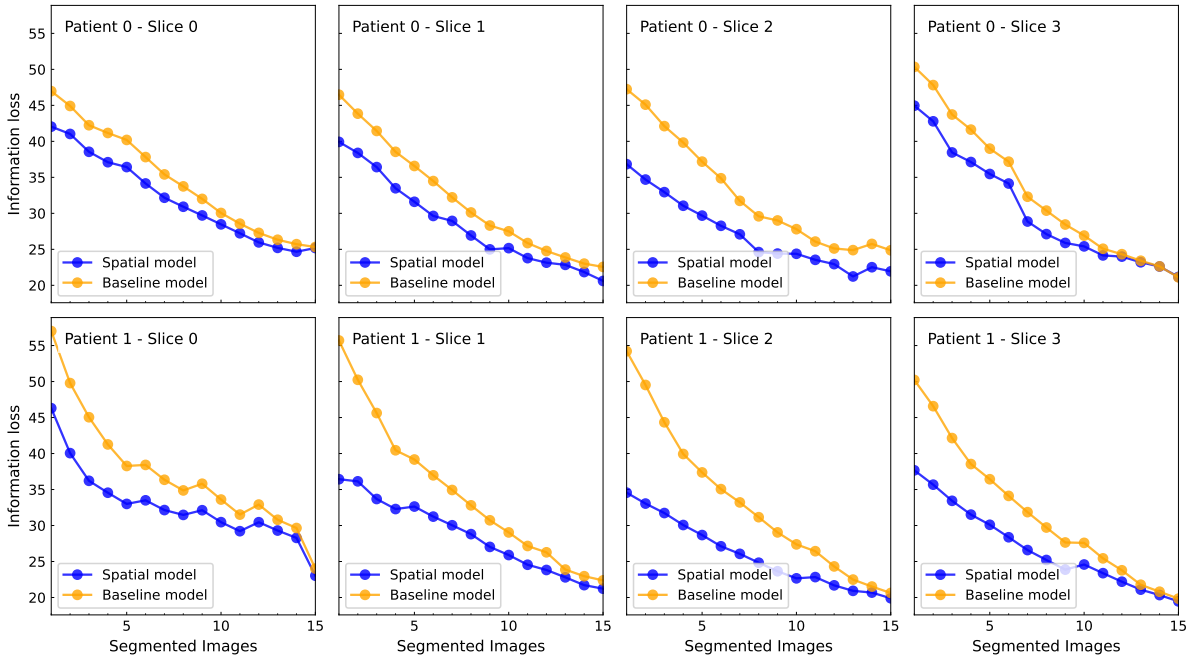

Figure 4: Information loss comparison of the baseline and spatial models on all slices and patients. In all plots, the corresponding slice is left out. Image-by-image measurements are included, up to 15 images per slice. The lower line for the spatial model indicates that it can transfer information from the other segmented slices.

slices, which is impossible for the baseline model. Especially for the more regular heartbeat of Patient 1, the spatial model vastly outperforms the baseline model. Although the influenc of spatial correlation in the prior diminishes as more data are incorporated, its effect remains observable even after 42 measurements for both patients.

4 CONCLUSION

In cardiac real-time MRI, segmenting the ventricle is essential to obtaining ventricular volume information, which is used to derive critical clinical parameters. However, maximizing the information from each labeled image is crucial, given the high segmentation cost. We incorporated prior knowledge into the proposed model to achieve this maximization , representing the heartbeat over time as a superposition of frequencies. We assumed that the magnitudes of each frequency are correlated across neighboring slices and integrated this assumption into the prior using the discretized 1D Laplace operator to enforce smoothness.

A key advantage of our model is its efficien y. By applying conjugacy, we computed the posterior analytically, eliminating the need for sampling or variational inference. This outcome simplifie inference and allows the direct computation of the marginal likelihood gradient with respect to the hyperparameters. Thus, we can efficientl optimize the type-II likelihood using gradient-based methods, such as BFGS. Moreover, the proposed model is stable with respect to the choice of s and σ^2.

We evaluated the proposed approach in two scenarios: one with a missing slice but complete information from the remaining slices and one with no segmented images for any slice. In both cases, the proposed method outperformed the baseline, which analyzes each slice in-

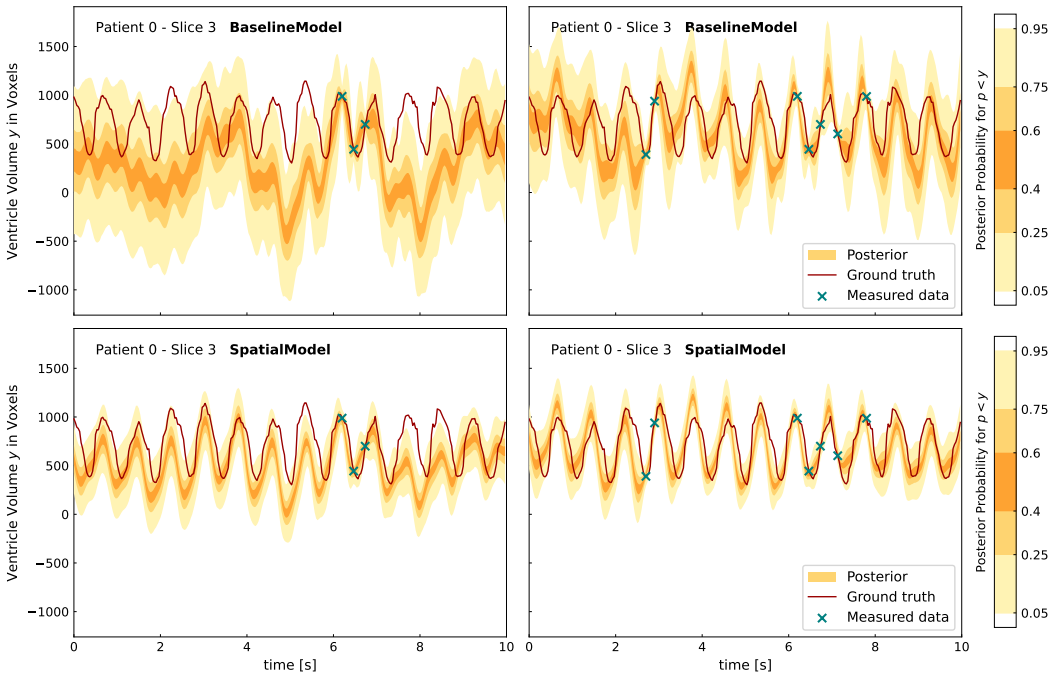

(a) Patient 0

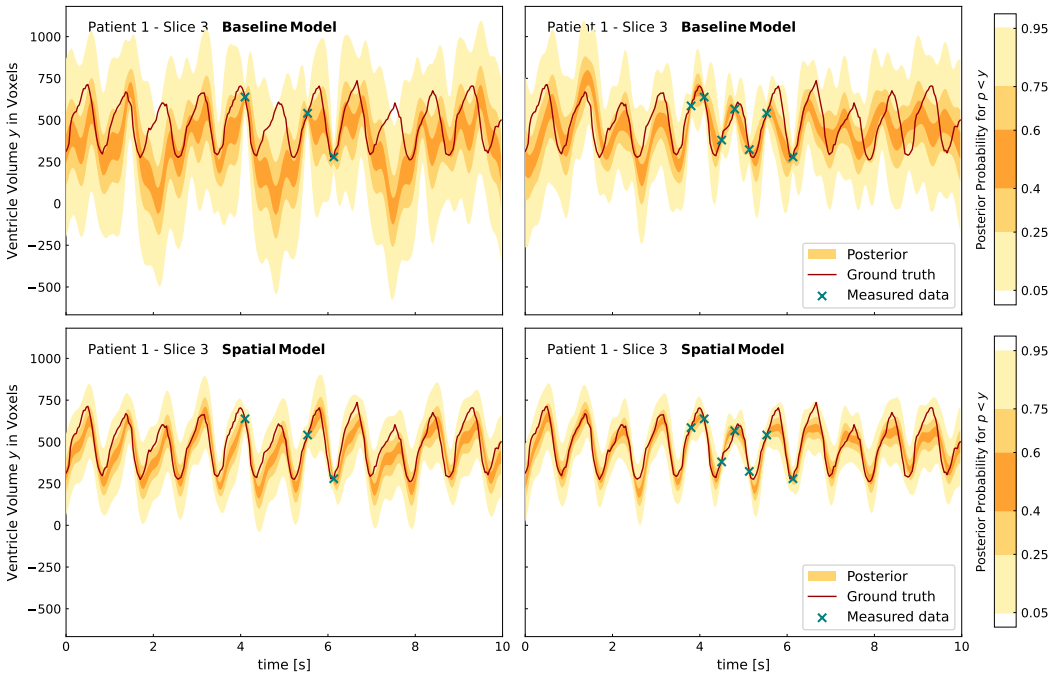

(b) Patient 1

Figure 5: Ventricle volume over time for Slice 3 for both patients. The red curve represents the ground truth from expert-validated segmentations; the yellow curve presents the posterior predictive based on (left) three and (right) seven measurements, assuming the remaining slices are fully segmented. The proposed model produces a posterior that closely aligns with the ground truth using information from neighboring slices. In contrast, the baseline model, which treats each slice independently, fails to capture this structure.

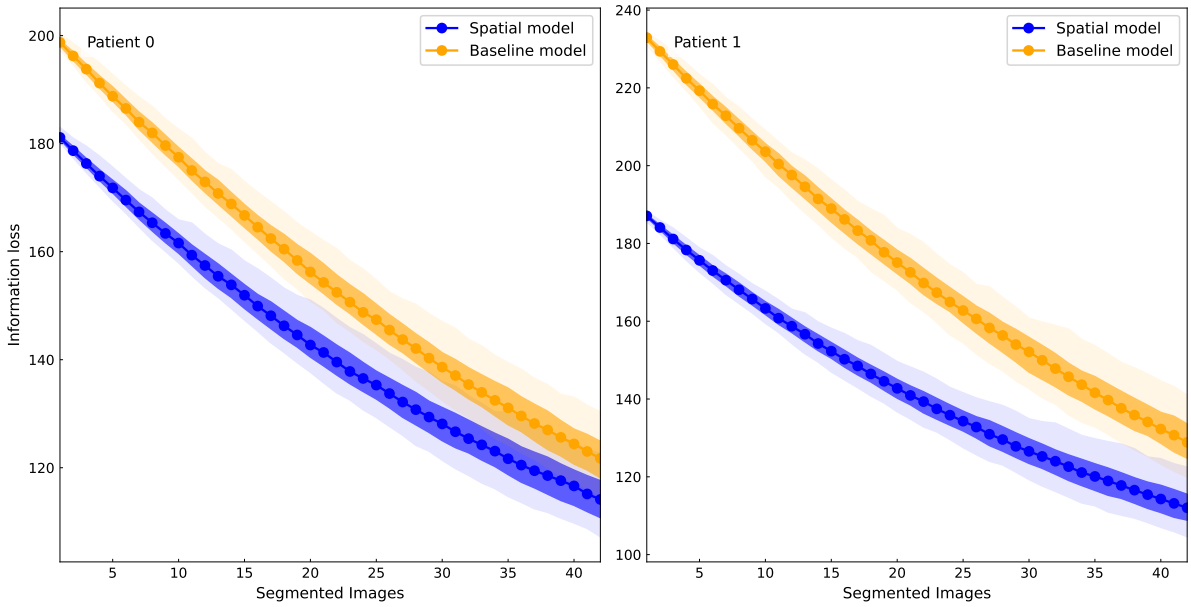

Figure 6: Information loss comparison between the baseline and spatial models for both patients. Up to 42 images were randomly selected and segmented throughout the four slices. The 300 identical random sequences were drawn for each patient and method. The solid line represents the median, whereas the shaded area represents the inner 50% and inner 90%. The lower line for the spatial model indicates that it can transfer information from other segmented slices.

dependently. The proposed method performed exceptionally well in data-scarce environments, which is critical in real-time MRI, where the high data volume limits manual segmentation, and complex cardiac geometries pose challenges for automatic methods. Furthermore, the results demonstrate that the proposed model produces a posterior that qualitatively aligns with the ground truth even with as few as three labeled frames per slice. In contrast, the baseline model fails to capture the underlying dynamics.

Acknowledgments

This work was performed as part of the Helmholtz School for Data Science in Life, Earth and Energy (HDS-LEE) and received funding from the Helmholtz Association of German Research Centres. This publication is partly supported by the Alexander von Humboldt Foundation and the King Abdullah University of Science and Technology (KAUST) Offic of Sponsored Research (OSR) under Award No. OSR-2019-CRG8-4033.

We would like to thank the whole Hypofon study team: Nicole Müller, Julian Alexander Härtel, Jan Schmitz, Ute Baur, Melanie von der Wiesche, Iris Rieger, Jon von Stritzky, Christopher Hart, Janina Bros, Benedikt Seeger, Emily Zollmann, Marijke Grau, Boris Dragutinovic, Laura-Maria de Boni, Jan-Niklas Hönemann, Wilhelm Bloch, Daniel Aeschbach, Eva-Maria Elmenhorst, Ulrike Herberg, Alena Hess, Moritz Schumann, Tobias Kratz, Jens Jordan and Johannes Breuer and the participants for the unique study data and Stiftung KinderHerz for making this study possible.

Appendix

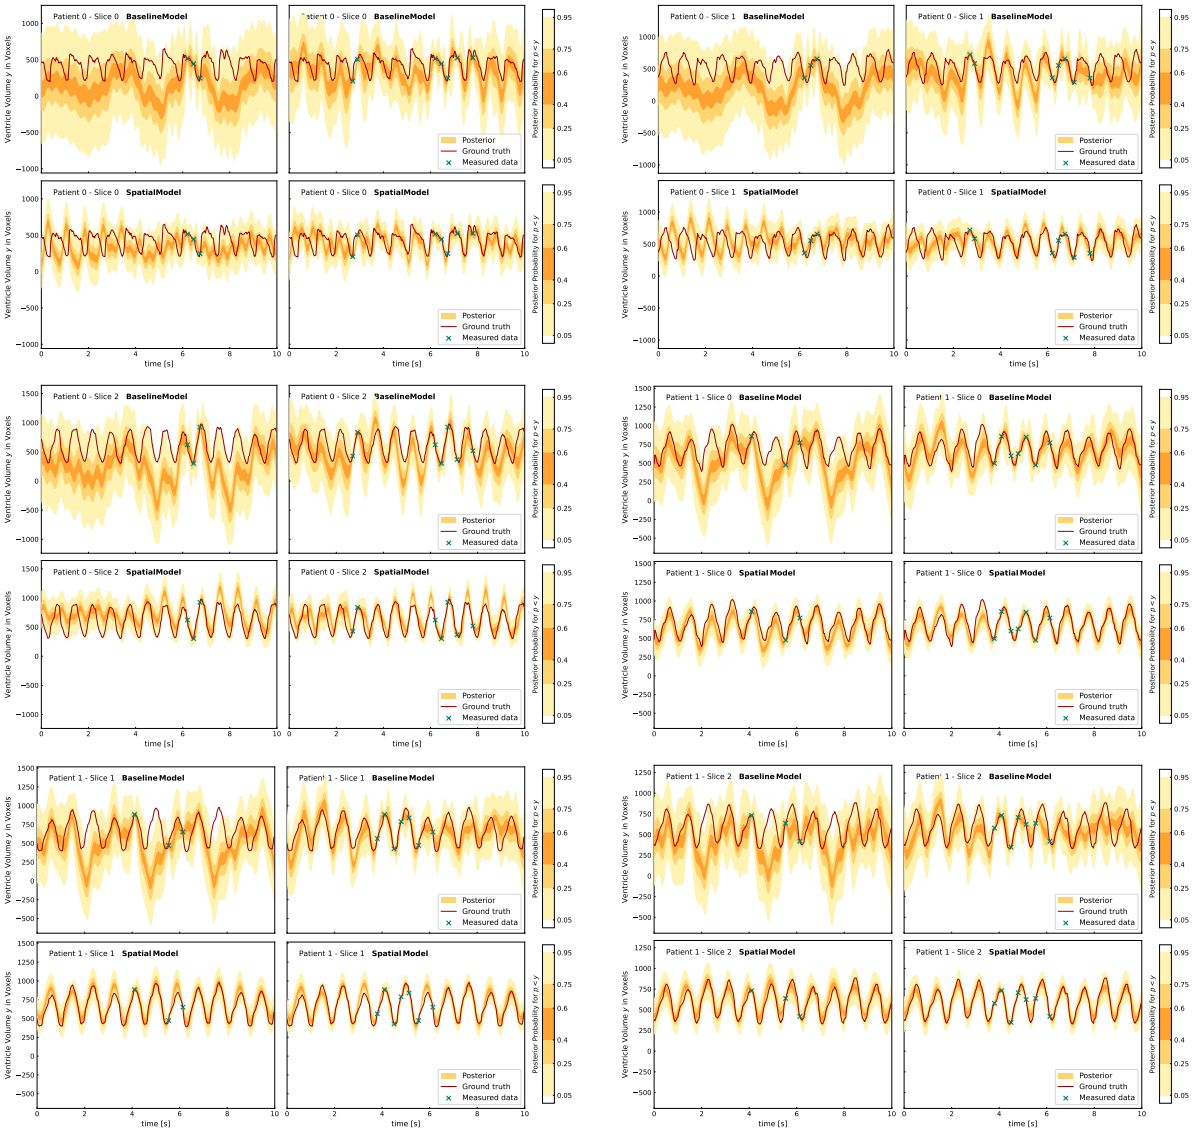

Figure 7: Ventricle volume over time for Slices 1, 2, and 3 in both patients. The red curve represents the ground truth from expert-validated segmentations; the yellow curve marks the posterior predictive, based on three and seven measurements, assuming the remaining slices are fully segmented. The proposed model produces a posterior that closely aligns with the ground truth using information from neighboring slices. In contrast, the baseline model, which treats each slice independently, fails to capture this structure.

REFERENCES

[1] C. G. Broyden. The convergence of a class of double-rank minimization algorithms 1. general considerations. *IMA Journal of Applied Mathematics*, 6(1):76–90, 03 1970.

[2] Bradley P Carlin and Thomas A Louis. *Bayes and Empirical Bayes Methods for Data Analysis*. Chapman and Hall, 1996.

[3] R. Fletcher. A new approach to variable metric algorithms. *The Computer Journal*, 13(3):317–322, 01 1970.

[4] Kay L. Gemba, Santosh Nannuru, Peter Gerstoft, and William S. Hodgkiss. Multi-frequency sparse Bayesian learning for robust matched fiel processing. *The Journal of the Acoustical Society of America*, 141(5):3411–3420, 05 2017.

[5] Peter Gerstoft, Christoph F. Mecklenbräuker, Angeliki Xenaki, and Santosh Nannuru. Multisnapshot Sparse Bayesian Learning for DOA. *IEEE Signal Processing Letters*, 23(10):1469–1473, 2016.

[6] Donald Goldfarb. A family of variable-metric methods derived by variational means. *Mathematics of computation*, 24(109):23–26, 1970.

[7] Kevin P. Murphy. *Probabilistic Machine Learning: Advanced Topics*. MIT Press, 2023.

[8] Nicole Müller, Julian Alexander Härtel, Jan Schmitz, Ute Baur, Melanie von der Wiesche, Iris Rieger, Darius Gerlach, Jon von Stritzky, Anja Bach, Christopher Hart, Janina Bros, Benedikt Seeger, Emily Zollmann, Marijke Grau, Boris Dragutinovic, Laura-Maria de Boni, Jan-Niklas Hönemann, Wilhelm Bloch, Daniel Aeschbach, Eva-Maria Elmenhorst, Ulrike Herberg, Alena Hess, Moritz Schumann, Tobias Kratz, Jens Jordan, Johannes Breuer, and Jens Tank. Peripheral oxygenation and pulmonary hemodynamics in individuals with fontan circulation during 24-hour high-altitude exposure simulation. *Circulation*, 149(18):1466–1468, 2024.

[9] Michael Salerno, Behzad Sharif, Håkan Arheden, Andreas Kumar, Leon Axel, Debiao Li, and Stefan Neubauer. Recent advances in cardiovascular magnetic resonance. *Circulation: Cardiovascular Imaging*, 10(6):e003951, 2017.

[10] David F Shanno. Conditioning of quasi-newton methods for function minimization. *Mathematics of computation*, 24(111):647–656, 1970.

[11] Felix Terhag, Philipp Knechtges, Achim Basermann, and Raúl Tempone. Sparse bayesian learning for label efficien y in cardiac real-time MRI. 2025. (preprint).

[12] Felix Terhag, Philipp Knechtges, Achim Basermann, and Raúl Tempone. Uncertainty quantificatio in machine learning based segmentation: A post-hoc approach for left ventricle volume estimation in MRI. *SIAM/ASA Journal on Uncertainty Quantification*, 13(1):90–113, 2025.

[13] Shuo Zhang, Arun A. Joseph, Dirk Voit, Sebastian Schaetz, Klaus-Dietmar Merboldt, Christina Unterberg-Buchwald, Anja Hennemuth, Joachim Lotz, and Jens Frahm. Real-time magnetic resonance imaging of cardiac function and fl w—recent progress. *Quantitative Imaging in Medicine and Surgery*, 4(5):313, October 2014.

[14] Shuo Zhang, Martin Uecker, Dirk Voit, Klaus-Dietmar Merboldt, and Jens Frahm. Real-time cardiovascular magnetic resonance at high temporal resolution: radial flas with non-linear inverse reconstruction. *Journal of Cardiovascular Magnetic Resonance*, 12(1):39, 2010.

[15] Zhilin Zhang and Bhaskar D. Rao. Sparse signal recovery with temporally correlated source vectors using sparse bayesian learning. *IEEE Journal of Selected Topics in Signal Processing*, 5(5):912–926, 2011.