

©Copyright 2025
Cathleen M Nguyen

Bayesian Vector Flow Mapping (B-VFM)

Cathleen M Nguyen

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2025

Reading Committee:

Juan Carlos del Álamo de Pedro, Chair

Alberto Aliseda

Tiffany Chen

Program Authorized to Offer Degree:

Mechanical Engineering

University of Washington

Abstract

Bayesian Vector Flow Mapping (B-VFM)

Cathleen M Nguyen

Chair of the Supervisory Committee:
Juan Carlos del Álamo de Pedro
Department of Mechanical Engineering

Cardiovascular disease, the leading cause of death in the United States, underscores the need for improved diagnostic imaging tools. While current clinical assessments of heart function rely primarily on global metrics such as ejection fraction, chamber pressure, and flow rate, regional flow imaging offers complementary insight. Intracardiac flow properties such as vortex formation and blood residence time reveal physiologic patterns that are not captured by global measures alone and can provide predictive information on pathological remodeling and thrombus risk.

Echocardiography, a non-invasive, non-ionizing, portable, and relatively inexpensive modality, is widely used in clinical practice. Color-Doppler echocardiography, in particular, provides flow information along the ultrasound beam direction and serves as the foundation for vector flow mapping (VFM), a technique to reconstruct two-dimensional velocity fields in the left ventricle. However, existing VFM methods remain limited: they are highly sensitive to noise, rely on heuristic hyperparameter selection, and treat reconstruction as a deterministic problem without quantifying measurement uncertainty. These limitations hinder the reliability of VFM in challenging clinical conditions where imaging data are imperfect.

This thesis introduces Bayesian Vector Flow Mapping (B-VFM), a hierarchical probabilistic framework for reconstructing intracardiac velocity fields from color-Doppler data while explicitly modeling uncertainty. First, we perform a theoretical error analysis of ultrasound

acquisition to characterize sources of variability in Doppler measurements and segmentation. These uncertainties are then propagated through the B-VFM formulation, which model priors as Gaussian distributions. Unlike traditional approaches, B-VFM optimizes for hyperparameters and outputs both velocity fields using a probabilistic approach, taking into account local measurement uncertainties. Validation on synthetic cardiovascular flows demonstrates reduced reconstruction errors compared to state-of-the-art, 'vanilla', VFM, and a patient case study highlights the potential of using reconstructed methods with their error maps for further patient analysis.

Finally, we discuss the extensibility of this general Bayesian framework, including integration of multimodal imaging, incorporation of more complex priors, and future applications to three-dimensional flow reconstruction. Collectively, this work establishes a principled foundation for uncertainty-aware flow imaging, with the potential to enhance the clinical value of echocardiographic diagnostics in cardiovascular disease.

TABLE OF CONTENTS

	Page
List of Figures	iii
List of Tables	vii
Chapter 1: Introduction	1
Chapter 2: Bayesian Vector Flow Mapping	6
2.1 Methods	6
2.2 Results	19
2.3 Ill-posedness and Regularization of VFM	30
2.4 Discussion	37
Chapter 3: Theoretical Analysis	44
3.1 Understanding Sources of Uncertainty in Medical Imaging	45
3.2 MUST: MATLAB Ultrasound Toolbox	46
3.3 Results	50
3.4 Discussion	53
Chapter 4: Multi-Modality Fusion	55
4.1 Formulation	56
4.2 Uncertainty Modeling	62
4.3 Hyperparameter Selection and Wall Uncertainty Quantification	65
4.4 Results	67
4.5 Discussion	69
Chapter 5: 3D AI Vector Flow Mapping	74
5.1 Physics Informed Neural Networks (PINNs)	74
5.2 3D VFM from Triplane color-Doppler	74

5.3	3D AI-VFM	75
5.4	3D VFM from Triplane color-Doppler	77
5.5	Results	77
5.6	Residence Time	77
5.7	Discussion	78
5.8	Acknowledgments	81
Chapter 6:	Concluding Remarks	83
Bibliography	86

LIST OF FIGURES

Figure Number	Page
1.1 Cylindrical Coordinate System	1
1.2 Vanilla VFM Pipeline	3
2.1 Bayesian VFM Pipeline	7
2.2 Ideal: Hyperprior Markov Monte Carlo Chains (iterations log scale)	19
2.3 Ideal: A) Vanilla-VFM and Bayesian VFM Inputs and Outputs, B) Vanilla-VFM vs Ground Truth Correlation Plots, C) Bayesian-VFM vs Ground Truth Correlation Plots	20
2.4 Doppler Hole: Hyperprior Markov Monte Carlo Chains (iterations log scale)	21
2.5 Doppler Hole: A) Vanilla-VFM and Bayesian VFM Inputs and Outputs, B) Vanilla-VFM vs Ground Truth Correlation Plots, C) Bayesian-VFM vs Ground Truth Correlation Plots	22
2.6 Vanilla-VFM: Impact of boundary condition enforcement with poor wall segmentation. A) Poor side wall segmentation, B) Poor "apex" wall segmentation.	24
2.7 Poor Segmentation Scenario with Gamma Approach: Hyperprior Markov Monte Carlo Chains (iterations in log scale) Algorithm 1	25
2.8 Poor Segmentation Scenario with Gamma Approach: A) Vanilla-VFM and Bayesian VFM Inputs and Outputs, B) Vanilla-VFM vs Ground Truth Correlation Plots, C) Bayesian-VFM vs Ground Truth Correlation Plots (Algorithm 1)	25
2.9 Poor Segmentation Scenario with Metropolis Hastings Approach: Hyperprior Markov Monte Carlo Chains (iterations in log scale) Algorithm 2	26
2.10 Poor Segmentation Scenario with Metropolis Hastings Approach: A) Vanilla-VFM and Bayesian VFM Inputs and Outputs, B) Vanilla-VFM vs Ground Truth Correlation Plots, C) Bayesian-VFM vs Ground Truth Correlation Plots (Algorithm 2)	26
2.11 Error in 2D planar flow assumption (log scale) Equ. 2.31. The blue histogram shows the divergence error from Vanilla VFM, while the red histogram corresponds to Bayesian VFM divergence error.	28

2.12	Clinical Case (Ejection): A) Ejection color-Doppler acquisitions and segmentations, B) B-VFM Input color-Doppler and data errors, C) Vanilla VFM reconstructed field and input deterministic mask, D) B-VFM reconstructed velocity field and mask with variance maps.	31
2.13	Vanilla-VFM challenged with unmet geometric constraints, producing reasonably accurate velocity field reconstructions (v_r, v_θ but exhibiting high uncertainty along the θ arc (Σ_{v_θ} , with N_r, N_θ in red).	37
3.1	A) Streamfunction of Hill's Vortex B) Streamfunction of Lamb-Chaplygin Dipole C) Generated synthetic ultrasound contrast echo using MUST	49
3.2	Particle coordinates were defined within the MATLAB Ultrasound Toolbox to simulate flow patterns of synthetic velocity fields at different transducer angles. Simulated ultrasound wave propagation and scatterer echoes generate realistic radiofrequency data, which is then processed to create synthetic the corresponding Doppler images.	50
3.3	Measurement of mean variance in synthetic doppler images when angle between the transducer and the flow field is varied.	51
3.4	Measurement of mean variance in synthetic doppler images r distance between transducer and flow field is varied for $\alpha = 0$ and $\alpha = 0.5$ dB/cm/MHz. . . .	52
4.1	Fusion Flowchart: 1) Echo-PIV and Doppler Fusion, Data 1 would contain both u_r and u_θ components provided by PIV's 2D field and Data 2 would contain u_r and $u_\theta = 0$ from color-Doppler acquisition. 2) Doppler Multiscale Fusion has $u_{r,1}, u_{\theta,1} = 0, u_{r,2}, u_{\theta,2} = 0$ as inputs. Both fusion modalities have corresponding error measurements and wall segmentation acquisitions as inputs. 57	57
4.2	PIV compares corresponding interrogation windows between two frames via cross-correlation. The resulting correlation surface is approximated by a paraboloid. The location of its maximum gives the displacement between frames, while the curvature coefficients (b,c) reflect measurement confidence. 58	58
4.3	Bootstrapping is used to estimate uncertainty in PIV measurements by resampling and replacing random pixels (denoted by red patterned boxes in image) in each input data interrogation window. PIV is rerun to obtain results for each new sample, providing an empirical distribution of the measurements for uncertainty quantification.	63

4.4	Examples of left ventricle ultrasound acquisitions and their corresponding segmentations. Figure A shows a higher-quality ultrasound acquisition with clear imaging of the left ventricle walls, enabling more confident and accurate segmentation of the side walls. However, some uncertainty remains at the apex due to intrinsic imaging limitations. Figure B illustrates a noisier ultrasound acquisition, which results in a more challenging segmentation process with greater uncertainty and variability. These examples highlight the subjective nature of manual segmentation, demonstrating how image quality can directly impacts segmentation reliability and downstream analyses.	67
4.5	CNN Segmentation. Panel A shows the input ultrasound image fed into the convolutional neural network. Panel B displays the resulting probabilistic atlas, where values range from 0 to 1 — with 0 indicating no confidence in the wall location and 1 representing full confidence in the detected wall location.	68
4.6	Echo-PIV Color-Doppler Fusion: Row 1 illustrates the comparison of obtaining echo-PIV velocity fields from the standard PIV method and the improved Multigrid-Ensemble method. Multigrid-Ensemble significantly improved the PIV result, providing higher dynamic range of the velocity field within the left ventricle. Row 2 illustrates the multimodality fusion results. After obtaining the 2D PIV velocity field, the result is projected into the polar coordinate system (ie the Doppler home coordinate system) to allow for fusion with Doppler acquisition. The final fused velocity field captures richer and more robust flow information by combining the strengths of both modalities.	69
4.7	Comparison of PIV, Vector Flow Mapping, and Echo-PIV Color-Doppler Fusion: Three left ventricular (LV) flow quantification methods were applied to patient data at four key points in the cardiac cycle: early LV filling, late LV filling, isovolumetric contraction, and ejection. Row 1 shows the PIV velocity fields derived from ultrasound contrast images. Row 2 presents the vector fields obtained using classic Vector Flow Mapping. Row 3 displays velocity fields from the echo-PIV color-Doppler fusion method. Qualitatively, the fused data set combines the comprehensive spatial coverage of PIV with the high-velocity detail captured by color-Doppler, resulting in a richer and more informative flow field.	71
4.8	PIV–Doppler Fusion Sensitivity Maps (normalized). Row 1 shows the PIV Doppler fused velocity fields. Rows 2 and 3 present local sensitivity estimates in the radial and azimuthal directions. Higher local sensitivity appears near boundaries, due to input segmentation uncertainty, and within the inlet jet, where qualitative discrepancies between PIV and Doppler inputs are more pronounced.	72

4.9	Doppler Multiscale Fusion at four key points in the cardiac cycle: early LV filling, late LV filling, isovolumetric contraction, and ejection. Rows 1 and 2 show color Doppler acquisitions obtained with high (0.66 m/s) and low (0.33 m/s) velocity encodings, respectively. Rows 3 and 4 present the corresponding radial and azimuthal velocity fields reconstructed using the Multiscale Fusion VFM method.	73
5.1	PINN Parameters	75
5.2	A) Synthetic Doppler slices from CFD Data. B) Ultrasound Triplane Clinical Acquisitions C) Synthetic Triplane Slices for PINNs training.	76
5.3	Correlation Plots	78
5.4	Planes trained with Doppler: Early Filling	79
5.5	Planes trained with Doppler: Late Filling	80
5.6	Planes trained with Doppler: Ejection	81
5.7	90th percentile and mean left ventricle residence time for CFD ground truth vs 3D AI-VFM.	82

LIST OF TABLES

Table Number	Page
2.1 Comparison of r^2 values for Vanilla-VFM and Bayesian VFM (B-VFM) under different synthetic scenarios.	24

LIST OF ALGORITHMS

1	MAP-Based Laplace-Approximate Gibbs Sampler	13
2	Metropolis-Hastings Sampling with Log-normal Hyperpriors	15
3	PIV Uncertainty Quantification with Bootstrap Method	64

ACKNOWLEDGMENTS

This PhD has been one of the most challenging and rewarding experiences of my life. I have learned more than I ever imagined possible, pushed myself to take on difficult problems, and persevered through obstacles that at times felt insurmountable. None of this would have been possible without the support, mentorship, and encouragement of those around me, to whom I owe immense gratitude.

First and foremost, I would like to thank my advisor, Dr. Juan Carlos del Álamo, for his guidance, patience, and unwavering mentorship throughout this journey. From the moment I began my PhD, I knew I wanted to work at the interface of engineering and medicine, and I knew he was the mentor who could help me pursue this path. His knowledge, vision, and judgment have profoundly shaped the way I think about problems; always balancing rigorous scientific inquiry with the ultimate goal of improving patient care. I have deep respect for the way he continually challenged me to grow as both an engineer and a researcher. The skills and perspective I have gained under his mentorship will guide me well beyond this dissertation.

I am also grateful to my committee members for their invaluable feedback and expertise. I had the privilege of learning from a diverse group of engineers, cardiologists, and scientists, each of whom broadened my perspective. I want to give a special thanks to Dr. Tiffany Chen for joining my committee in the later stages of my PhD and offering fresh insights that helped me reflect on the broader impact of my work. Their feedback during and after my defense motivated me to think not only about the technical contributions but also about the long term translational implications of this research.

I extend my gratitude to Dr. Juan Lasheras and Dr. James Kirkpatrick for their men-

torship and guidance throughout my doctoral studies. It has been a privilege to learn from them, and I deeply value the opportunities I had to work with both. I would also like to acknowledge the contributions of my co-authors and collaborators whose work is included in the dissertation chapters. Their insights, expertise, and dedication helped shape this research into what it is today. Thank you Dr. Darrin Wong, Dr. Sachiyo Igata, Dr. Manuel Guerrero-Hurtado, Dr. Pablo Martinez-Legazpi, Christian Chazo Paz, Dr. Javier Bermejo, Dr. Andrew Kahn, and Dr. Anthony DeMaria, for their hard work and for being such generous teachers along the way.

I am thankful to the members of the del Álamo lab and my peers in the Mechanical Engineering Department at UCSD for their encouragement and camaraderie. From San Diego to Seattle, I have been fortunate to be surrounded by such talented colleagues and friends. Thank you to Yi-Ting Yeh, Alejandro Gonzalo Grande, Adithan Kandasamy, Bahetihazi Maidu, Clarissa Bargellini, Yunpeng Tu, María Guadalupe Borja, Amy Schwartz, Antoni de Herreros, Stephanie Sincomb, Ernesto Hidalgo, Ricardo Fernandez-Serrano, Lorenzo Rossini, Shun Zhang, Jennifer Fromm, Hossein Firouznia, and Jack Pechac. I will always treasure our lunches, conferences, late night studies, and especially innertube water polo games that kept the PhD both balanced and memorable.

To my friends and family, I cannot thank you enough. To those who stayed close when I moved to Washington, and to the new friends I made along the way; you brought joy, adventure, and balance into my life, reminding me to keep living fully outside of my research. I could not have imagined being surrounded by so much love and support, and I am endlessly grateful.

I want to express my deepest gratitude to my partner, Kyle Lee. Having you by my side has made this journey immeasurably more meaningful. Your support, encouragement, and love have been constant throughout. I am so thankful for the balance, joy, and perspective you brought to my life during these years. Thank you for always pushing me to be my best,

for believing in me when I doubted myself, and for encouraging me to keep reaching higher. I feel incredibly fortunate to have shared this chapter with you, and I look forward to all the chapters to come.

To my sisters, Grace and Carolyn, we've come such a long way. I'm so proud of each of our paths and how we've pushed through challenges to reach our goals. You both give me so much strength and remind me that no one but us can define our future or our success. I love you both so much, and I'm so excited what the future has in store for the Nguyen sisters.

VITA

- 2025 Ph.D. in Mechanical Engineering
 University of Washington, Seattle
- 2024 PhD Research Intern
 META Reality Labs, Redmond
- 2019 M.S. in Mechanical Engineering
 University of California, San Diego
- 2018 B.S. in Bioengineering
 University of California, San Diego
- 2017 Product Lifecycle Management Intern
 Johnson and Johnson Medical Devices, Irvine

PUBLICATIONS

Nguyen, C.M., Martinez-Legazpi, P., Guerrero-Hurtado M., Kandasamy, A., Maidu, B., Gonzalo, A., Kahn, A.M., Bermejo, J. and del Alamo, J. C. Bayesian vector flow mapping (B-VFM). *Ready to Submit*.

Nguyen, C.M., del Alamo, J. C., A Theoretical Analysis of Ultrasound Uncertainty and Propagation. *In preparation*.

Bahetihazi Maidu, Pablo Martinez-Legazpi, Manuel Guerrero-Hurtado, Cathleen M. Nguyen, Alejandro Gonzalo, Andrew M. Kahn, Javier Bermejo, Oscar Flores, Juan C. del Alamo, Super-resolution left ventricular flow and pressure mapping by Navier–Stokes-informed neural networks, *Computers in Biology and Medicine* 2025.

María Guadalupe Borja, Pablo Martinez-Legazpi, Cathleen Nguyen, Oscar Flores, Andrew M. Kahn, Javier Bermejo, Juan C. del Álamo, Deriving phenotype-representative left ventricular flow patterns by reduced-order modeling and classification, *Computers in Biology and Medicine* 2024.

Rossini L, Braun OÖ, Brambatti M, Benito Y, Mizeracki A, Miramontes M, Nguyen C, Martinez-Legazpi P, Almeida S, Kraushaar M, Vu V, May-Newman K, Bermejo J, Adler ED, Kahn AM, Del Alamo JC. Intraventricular Flow Patterns in Patients Treated with Left Ventricular Assist Devices. *ASAIO Journal* 67(1):p 74-83, January 2021.

Manuel Gálvez-Santisteban, Danni Chen, Ruilin Zhang, Ricardo Serrano, Cathleen Nguyen, Long Zhao, Laura Nerb, Evan M Masutani, Julien Vermot, Charles Geoffrey Burns, Caroline E Burns, Juan C del Álamo, Neil C Chi (2019) Hemodynamic-mediated endocardial signaling controls in vivo myocardial reprogramming eLife 8:e44816.

DEDICATION

*To the people who showed up and stood by me through every turn, thank you for your
unwavering support.*

Chapter 1

INTRODUCTION

In the past two decades, advances in medical imaging have allowed for enhanced visualization of the cardiovascular system. As a result, methods to measure intracardiac blood flow and analyze flow patterns have also improved [10, 34, 48]. Among the existing cardiovascular imaging modalities, echocardiography offers non-ionizing, inexpensive, portable imaging with short acquisition times. In particular, color-Doppler echocardiography remains the workhorse of clinical left ventricular (LV) flow evaluation. However, while it can be measured in linear, planar, or volumetric sectors, color-Doppler only senses the flow velocity in the direction of the ultrasound beam (1.1). Therefore, significant efforts have been devoted to inferring velocities in the cross-beam directions [37]. One of such techniques that is gaining widespread acceptance is vector flow mapping (VFM) [5, 6, 23, 25, 36, 49, 72, 80].

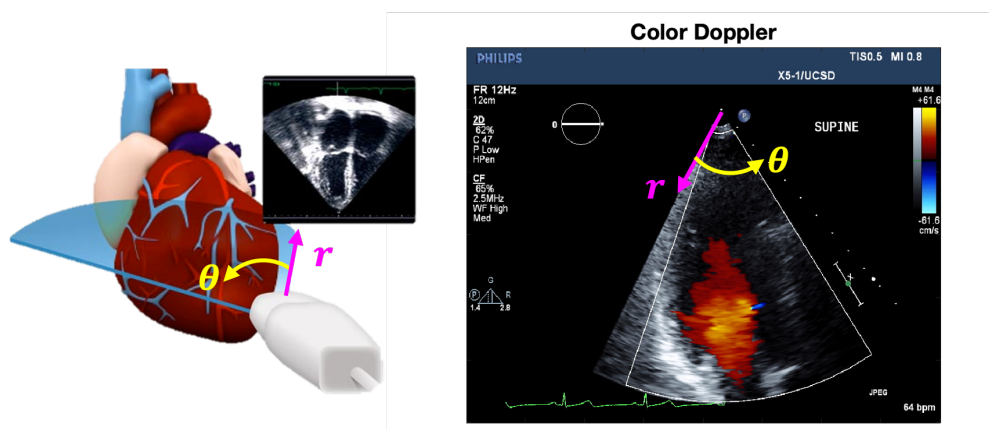


Figure 1.1: Cylindrical Coordinate System

VFM processes color-Doppler data and a time-dependent segmentation of the LV wall in the apical long-axis view, also known as the three-chamber view. This echocardiographic view displays the mitral annulus, lateral LV wall, the apex, the LV outflow tract, and the aortic valve, making it ideal for observing the whole transit of blood through the LV. Each heartbeat, the dominant LV flow pattern alternates between a diastolic inflow jet flanked by a vortex ring, a telediastolic prograde swirling cell that redirects incoming blood towards the LV outflow tract, and a systolic jet [10]. Realizing that these features can be observed in the apical long-axis view [58], and that their dynamics were approximately planar [76] paved the way for VFM. Investigators have validated VFM on this plane using ground-truth data from CFD [5, 6, 36, 49, 77, 80] and *in vitro* models [3, 25]. Head-to-head validation of echocardiographic VFM measurements vs. phase-contrast MRI measurements has also shown favorable agreement in flow velocities [9, 50]. Moreover, numerical descriptors of LV transport computed on the 3-chamber view correspond well with equivalent metrics computed from 4D flow MRI [54]. In addition to the LV, VFM has been evaluated in other vascular geometries, showing encouraging agreement with reference data from particle image velocimetry [4].

Early VFM implementations estimated the cross-beam flow velocity by naively enforcing mass conservation along circular arcs of the color-Doppler sector and no-penetration boundary conditions at the LV walls [25, 36]. Subsequent improvements combined these constraints and Tikhonov regularization by formulating VFM as a maximum-likelihood estimator [5] (Figure 1.2). This formulation assumes that each constraint’s residual follows a normal statistical distribution (i.e., a prior) with zero mean and uniform standard deviation. The inverse of each prior’s standard deviation is a hyperparameter that acts as weight for the associated constraint. The number of hyperparameters can be reduced by introducing Lagrange multipliers to strictly enforce the physical constraints [80]. Several groups have proposed VFM reconstructions based on the streamfunction vorticity formulation, which is applicable in planar, incompressible flow [49, 52]. These methods have the convenience of implicitly introducing regularization by numerically solving a Poisson equation to determine the

streamfunction. Secondary analyses of VFM velocity fields can quantify LV flow transport

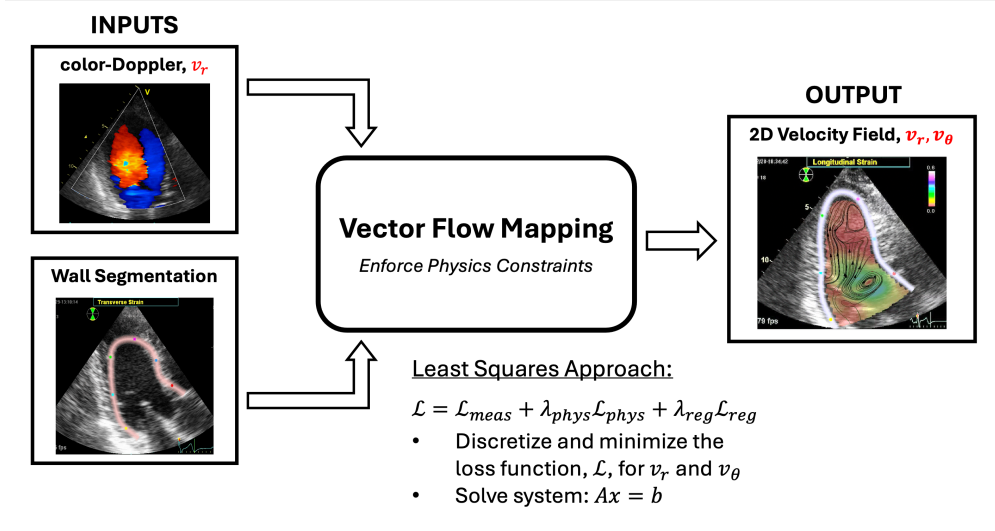


Figure 1.2: Vanilla VFM Pipeline

topology [32], blood stasis [66], and fluctuating pressure fields [74]. Among other applications, VFM has been used to quantify the changes in LV flow caused by age [7], adverse LV remodeling [9, 12, 47, 55], pacemaker or left ventricular assistance device implantation and programming [62, 64, 65], and cardioembolic stroke risk [46, 59, 60]. Apart from the LV, several studies have applied VFM to quantify flow in the right ventricle [1, 16, 50], the aortic root [31, 42] and other regions of the cardiovascular system [33].

While VFM is an effective modality for quantifying LV flow patterns, most implementations enforce in-plane mass conservation and boundary conditions strictly, neglecting out-of-plane fluxes and errors associated with LV wall tracking. Maximum-likelihood implementations offer the possibility of balancing the residual from physical constraints and regularization penalties [5] but this balance depends on the hyperparameter values. While heuristic approaches such as the L-curve criterion can inform the choice of hyperparameters [30], they add significant computational cost, lack robustness, and introduce arbitrariness in the solu-

tion. Adding to the lack of schemes to balance the different error sources, VFM uncertainty propagation is yet to be studied in detail. Naive VFM techniques considered error accumulation when integrating the mass conservation equation along circular arcs to choose weight functions for the integration [25] or to estimate posterior cross-beam velocity errors [73]. However, these methods neglected the errors associated with wall tracking and regularization, and lumped all other errors together without distinguishing their origin.

Bayesian methods have emerged as a powerful tool for solving inverse problems while incorporating uncertainty quantification. For example, Sun et al. applied a Bayesian framework to optical flow estimation, providing spatially resolved confidence in motion fields [71]. Similarly, Borggaard et al. used Bayesian inference to estimate background flows from passive scalar data, highlighting the flexibility of probabilistic approaches in fluid estimation tasks [11]. More recently, Kontogiannis et al. [40] developed a Bayesian inverse Navier–Stokes framework to regularize 3D velocity fields and learn boundary position and parameters such as the viscosity coefficient.

This work introduces B-VFM, a vector flow mapping framework rooted in hierarchical Bayesian inference that uses as inputs 1) color-Doppler velocity images and their variances and 2) left ventricle mask segmentation images and their variances. Its outputs are a corrected LV segmentation and color-Doppler velocity field, the missing cross-beam velocity field, and estimates of the errors of these quantities at each spatial location. B-VFM is based on Gaussian priors for conservation of mass, boundary conditions, and spatial smoothness. Of note, B-VFM uses a heteroskedastic prior for mass conservation that significantly reduces errors in the recovered cross-beam velocity. Our B-VFM approach is validated on cardiovascular relevant synthetic flows and tested on a patient case. To further support the Bayesian formulation, we conduct a theoretical analysis of ultrasound acquisition to characterize sources of uncertainty and study their propagation through the reconstruction process. The first two chapters are organized as follows: 1) Bayesian Vector Flow Mapping B-VFM

and 2) Theoretical Analysis. Building on these foundations, the final chapter broadens the scope by examining alternative approaches to vector flow mapping: Multi-modality fusion and AI-VFM.

Chapter 2

BAYESIAN VECTOR FLOW MAPPING

2.1 Methods

Bayesian VFM (B-VFM) combines physics-informed priors and input data uncertainty to infer LV flow velocity fields and correct the LV mask. B-VFM extends recent least-squares approaches that yield a single best-fit solution by minimizing a cost function [5, 6, 80] by treating the velocity (\vec{v}) and LV mask (m) as random variables. B-VFM relies on a probability model that encodes our beliefs about the likelihood of different solutions, and sample the posterior distribution conditioned on the observed data. The final estimates correspond to the mean of the posterior distribution of \vec{v} and m their uncertainty estimates correspond to the posterior's standard deviation. We introduce the physical models encoding prior beliefs in §2.1.1 below, while §2.1.2 describes the hierarchical methods used to approximate the posterior distribution. We study B-VFM using ground-truth velocity fields from two synthetic flows that resemble the LV diastolic vortex rings: the Lamb-Chaplygin (LC) dipole and Hill's vortex. These flows and the validation methods are described respectively in §2.2. To demonstrate the clinical application of B-VFM, we apply this method to patient-specific echocardiographic acquisitions §2.2.4.

2.1.1 Bayesian Formulation of Vector Flow Mapping

Adopting a cylindrical coordinate system with its origin at the ultrasound transducer (Figure 1.1) B-VFM can be formulated as the inversion of a linear transformation relating the measured color-Doppler velocity (V) and LV mask (M) to the true radial and transverse

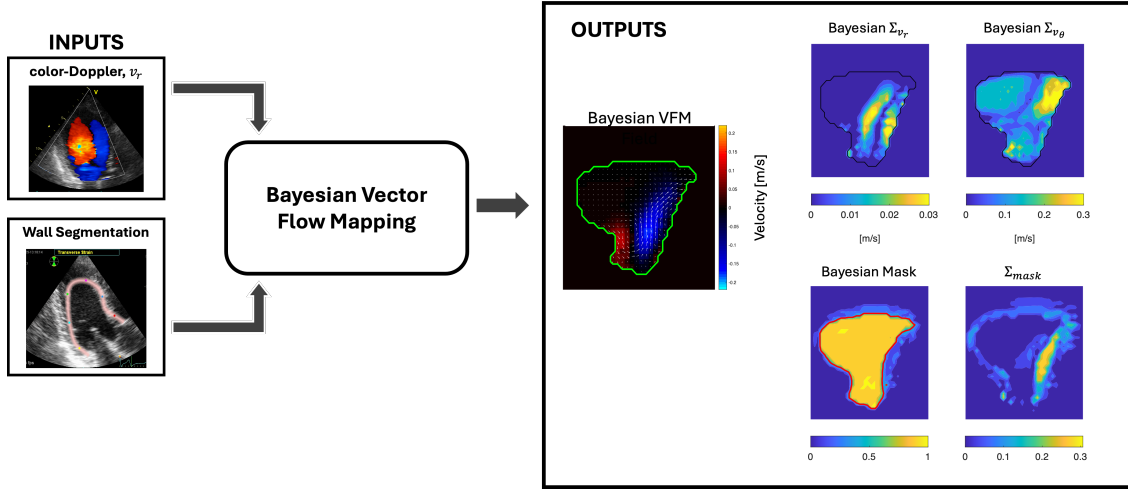


Figure 2.1: Bayesian VFM Pipeline

velocities (v_r and v_θ), and the LV mask, i.e.,

$$\vec{X} = \begin{bmatrix} V \\ M \end{bmatrix} = \begin{bmatrix} I & 0 & 0 \\ 0 & 0 & I \end{bmatrix} \cdot \begin{bmatrix} v_r \\ v_\theta \\ m \end{bmatrix} + \begin{bmatrix} \varepsilon_V \\ \varepsilon_M \end{bmatrix} = A\vec{x} + \vec{\varepsilon}. \quad (2.1)$$

Discretizing this transformation on a regular mesh with N_r and N_θ points in the radial and angular directions, the $3N_r N_\theta \times 1$ parameter vector, $\vec{x} = (v_r, v_\theta, m)^T = (\vec{v}, m)^T$, contains the values of the velocity components and LV mask indicator function at each pixel of the image. The $2N_r N_\theta \times 1$ measurement vector is $\vec{X} = (V, M)$, where V contains color-Doppler velocity values and M indicates whether each pixel belongs to the LV cavity Ω . In a zero-uncertainty scenario, M is a binary mask containing ones and zeros whose boundary, $\partial\Omega$, delineates the LV endocardium. More broadly, M is a probabilistic atlas obtained from a segmentation algorithm [18] or by averaging several binary masks delineated by different users or methods. The measurement errors vectors ε_V and ε_M are assumed to be white noise, so that V and M follow normal distributions with zero mean and diagonal covariance matrices Σ_V and Σ_M ,

$$\varepsilon_V \longrightarrow \mathcal{N}(0, \Sigma_V), \quad \varepsilon_M \longrightarrow \mathcal{N}(0, \Sigma_M). \quad (2.2)$$

These two matrices can be considered as known if the Doppler and segmentation errors are available from the imaging pipeline. Alternatively, they can be considered hyperparameters in the hierarchical Bayesian approach (see §2.1.2). In the latter case, it is convenient to model $\Sigma_{V,M} = \varepsilon_{V,M}I$ assuming spatially homogeneous noise, but more involved spatial models considering, e.g., the depth-dependence of Doppler measurements, are also possible. In particular, it is reasonable to assume that Σ_M is low in regions well inside and well outside the LV mask regardless of the segmentation method.

The linear transformation in equation 2.1 is represented by a rectangular matrix, A , whose I and 0 blocks are respectively identity and square null matrices of size $N_r N_\theta$. Because A is rectangular and $A^T A$ is singular, the transformation cannot be inverted even in least squares sense without injecting prior information about v_θ and the other two variables. Thus, we consider that the divergence of the velocity field in Ω is follows a normal distribution of zero mean and diagonal covariance matrix

$$\nabla \cdot \vec{v}|_\Omega \longrightarrow \mathcal{N}(0, \Sigma_{CM}). \quad (2.3)$$

In discrete form, this prior can be written as

$$p(\vec{x} | \alpha_{CM}, D) \sim |\Sigma_{CM}|^{-1/2} |\delta(m)| \exp \left[-\frac{1}{2} \vec{v}^T D^T \delta(m)^T \Sigma_{CM}^{-1} \delta(m) D \vec{v} \right], \quad (2.4)$$

where $\vec{v}^T = (v_r, v_\theta)$, D is a $2N_r N_\theta \times N_r N_\theta$ matrix representing the discrete divergence operator in cylindrical coordinates, and $\delta(m)$ is a diagonal matrix of size $N_r N_\theta$ that indicates the pixels belonging to the blood pool, Ω . The divergence covariance matrix for this prior is modeled as $\Sigma_{CM} = \alpha_{CM}^{-1} I$ where α_{CM} is a precision hyperparameter. This model assumes that departures from zero divergence are spatially homogeneous and uncorrelated. Empirical knowledge of the spatial structure of $\nabla \cdot \vec{v}$ from, e.g., phase contrast MRI or computational fluid dynamics, could be used to further inform this model, an improvement beyond the scope of this study. Here, the complete representation of the divergence operator including the r^{-1} metric, i.e., $\nabla \cdot \vec{v} = r^{-1}[\partial_r(rv_r) + \partial_\theta(v_\theta)]$, should be used to prevent the appearance of r -dependent errors. Mathematically, it is straightforward to see that not including the

metric term is equivalent to assuming that $\alpha_{CM} \sim r^2$. Previous VFM methods often neglect the r^{-1} metric by multiplying both sides of the $\nabla \cdot \vec{v} = 0$ equation by r , but this simplification lacks justification unless zero divergence is strictly enforced.

Additional priors include wall boundary conditions for mass conservation and smoothness constraints on the velocity field and the LV wall. Wall boundary conditions for mass conservation imply that blood velocity must be parallel to the segmented LV wall, i.e., $\vec{v} \cdot \vec{n} = 0$ where \vec{n} is the vector normal to $\partial\Omega$. Thus, we consider

$$\nabla m \cdot \vec{v} \longrightarrow \mathcal{N}(0, \Sigma_{BC}) \quad (2.5)$$

$$\nabla^2 v_r, \nabla^2 v_\theta \longrightarrow \mathcal{N}(0, \Sigma_{SV}), \quad (2.6)$$

$$\nabla^2 m \longrightarrow \mathcal{N}(0, \Sigma_{SM}), \quad (2.7)$$

where the covariance matrices $\Sigma_{BC} = \alpha_{BC}I$, $\Sigma_{SV} = \alpha_{SV}I$, and $\Sigma_{SM} = \alpha_{SM}I$ are determined by their respective precision hyperparameters, α_{BC} , α_{SV} , and α_{SM} . The corresponding discretized probability functions are:

$$p(\vec{x} | \alpha_{BC}) \sim |\Sigma_{BC}|^{-1/2} |\delta_{BC}| \exp \left[-\frac{1}{2} \vec{v}^T (Gm) \delta_{BC}^T \Sigma_{BC}^{-1} \delta_{BC} (Gm)^T \vec{v} \right], \quad (2.8)$$

$$p(\vec{x} | \alpha_{SM}, \alpha_{SV}) \sim |\Sigma_{SV}|^{-1} |\Sigma_{SM}|^{-1/2} \exp \left[-\frac{1}{2} (v_r^T L^T \Sigma_{SV}^{-1} L v_r + v_\theta^T L^T \Sigma_{SV}^{-1} L v_\theta + m^T L^T \Sigma_{SM}^{-1} L m) \right], \quad (2.9)$$

where $L = DG$ is the discretized Laplacian operator, and δ_{BC} is a diagonal indicator matrix that zeroes out pixels near open cardiac valves. This matrix can also be tailored to exclude enforcement of $\nabla m \cdot \vec{v} = 0$ in regions known *a priori* to lie far from the LV endocardium, septal defects, LVAD inlet cannulas, etc.

Therefore, we model the final prior distribution including all physical and smoothness constraints as

$$p(\vec{x} | \vec{\alpha}) \sim p(\vec{x} | \alpha_{CM}) p(\vec{x} | \alpha_{BC}) p(\vec{x} | \alpha_{SM}, \alpha_{SV}), \quad (2.10)$$

and the likelihood of our observations as

$$p(\vec{X} | \vec{x}) \sim |\Sigma_V|^{-1/2} |\Sigma_M|^{-1/2} |\delta(m)| \exp \left[-\frac{1}{2} (v_r - V)^T \delta(m)^T \Sigma_V^{-1} \delta(m) (v_r - V) - \frac{1}{2} (m - M)^T \Sigma_M^{-1} (m - M) \right], \quad (2.11)$$

where notation has been compacted using the hyperparameter vector $\vec{\alpha} = (\alpha_{CM}, \alpha_{BC}, \alpha_{SV}, \alpha_{SM})$, measurement noise vector $\vec{\varepsilon} = (\varepsilon_V, \varepsilon_M)$, and operator vector $\mathcal{O} = \{D, G, \delta_{BC}\}$. In section § 2.1.2, these models are used to find maximum posteriori (MAP) estimates of \vec{x} that maximize the conditional posterior $p(\vec{X}|\vec{x})p(\vec{x}|\vec{\alpha})$ which, together with marginalization over $\vec{\alpha}$, allows us to infer expected values and variances of \vec{x} . For simplicity and generality, we assume the Doppler map V has already been unwrapped. Appendix §2.4.3 provides probability models for phase unwrapping, showing phase unwrapping can be performed as a standalone preprocessing step prior to running B-VFM while retaining Bayesian rigor.

2.1.2 Hierarchical Bayesian Framework

We adopt a hierarchical Bayesian formulation to infer the unknown field \vec{x} along with hyperparameters $\vec{\alpha}$ that modulate the strength of the prior constraints. In this framework, the hyperparameters are themselves assigned hyperpriors $p(\vec{\alpha})$ to reflect uncertainty, and our goal is to estimate the posterior marginals $p(\vec{x}|\vec{X})$ and $p(\vec{\alpha}, \vec{X})$, or at least compute their expectation and variance by integrating their joint posterior distribution,

$$p(\vec{x}, \vec{\alpha} | \vec{X}) = \frac{p(\vec{x}, \vec{X} | \vec{\alpha})p(\vec{\alpha})}{p(\vec{X})} \sim p(\vec{X}|\vec{x})p(\vec{x}|\vec{\alpha})p(\vec{\alpha}). \quad (2.12)$$

The joint posterior cannot be evaluated directly due to the coupling between \vec{x} and $\vec{\alpha}$, and the intractability of the marginal likelihood $p(\vec{X})$. However, the models developed in §2.1.1 enable the construction of Markov Chain Monte Carlo methods that iteratively sample from known conditional distributions, creating chains that converge to the joint posterior [26]. Each MCMC chain is initialized using plausible hyperparameter values derived from a velocity field estimated by Vanilla VFM. To ensure robustness, an initial segment of each chain (burn-in) is discarded prior to inferring expectation and variance. We also run multiple MCMC chains with different hyperparameter initializations and evaluate whether algorithm consistently converges to similar posterior distributions.

MAP-Based Laplace-Approximate Gibbs Sampler

The Gibbs sampler is a well-known method to create MCMC chains that converge to the joint posterior [13, 57]. In iteration it of this method, one would sample \vec{x}^{it} by drawing from $p(\vec{x}|\vec{X}, \vec{\alpha}^{it-1}) \sim p(\vec{X}|\vec{x})p(\vec{x}|\vec{\alpha}^{it-1})$, and $\vec{\alpha}^{it}$ from $p(\vec{\alpha}|\vec{x}^{it})$. If the hyperprior $p(\vec{\alpha})$ is chosen carefully (e.g., a gamma distribution), then the hyperparameter conditional posterior $p(\vec{\alpha}|\vec{x}) \sim p(\vec{x}|\vec{\alpha})p(\vec{\alpha})$ retains the same functional form as $p(\vec{\alpha})$ and can be obtained in closed form. This property is widely known as conjugacy [26].

However, in B-VFM it is not possible to sample $p(\vec{x}|\vec{\alpha}^{it-1}, \vec{X})$ directly and conjugacy is not strictly observed because the precision matrices Σ^{-1} of some priors, most notably, the boundary conditions prior (eq. 2.8), depend on \vec{x} . Therefore, we use a modification of the traditional Gibbs sampler (see Algorithm 1) for details). We first maximize $p(\vec{x}|\vec{X}, \vec{\alpha}^{it-1})$ to obtain a maximum a posteriori (MAP) estimate \vec{x}^* and estimate the posterior's local curvature by the Hessian $H(\vec{x}^*)$ of the negative log-probability, from which we obtain the covariance as $\Sigma = 2[H(\vec{x}^*)]^{-1}$. Then, a posterior sample \vec{x}_{samp} is drawn from the Gaussian distribution $\mathcal{N}(\vec{x}^*, \Sigma)$, representing a local Laplace approximation to the full conditional.

Hyperparameter values are subsequently sampled using gamma distributions for hyperparameters, i.e., $p(\alpha) = \Gamma(\alpha|\beta, \phi) = \alpha^{\beta-1}e^{-\alpha\phi}$, is convenient due to conjugacy with the Laplace-approximated priors, where the dependence on \vec{x} is treated as fixed during the Gibbs update, preserving the required exponential form, and leading conditional posteriors that are also gamma distributions:

$$p(\alpha_i | \beta_i, \phi_i, \vec{x}) \sim \alpha_i^{r_i/2+\beta_i-1} \exp \left[-\alpha_i \left(\phi_i + \frac{1}{2}\vec{x}^T H_i \vec{x} \right) \right] \sim \Gamma \left(\beta_i + \frac{r_i}{2}, \phi_i + \frac{1}{2}\vec{x}^T H_i \vec{x} \right). \quad (2.13)$$

In this expression, the i subindex denotes the prior, H_i is the associated Hessian block, $r_i = \text{rank}(H_i)$, and the determinant of H_i has been expanded as $|H_i| = \alpha_i^{r_i}$. The tractability of this posterior allows for direct sampling in MCMC.

A closed-form expression for H is provided in §2.1.3. While, in principle, this expression should permit obtaining Σ , the number of degrees of freedom $N_d = 3(N_r N_\theta)$ makes direct

matrix inversion impractical. Thus, we compute the Hessian’s Cholesky decomposition, $H = LL^T$, draw a sample from $N(0, 1)$, and rescale this sample solving linear system with matrix L to obtain \vec{x}_{sample} (see Algorithms 1-2). To ensure the feasibility of the Cholesky decomposition, we perform it on the regularized symmetrized matrix $H_{\text{reg}} = (H + H^T)/2 + \varepsilon I$, where ε is initialized at 10^{-8} and iteratively doubled until positive definiteness is achieved. In addition to enforcing positivity, the gamma distribution offers a flexible means of encoding prior beliefs through its shape and scale parameters, β_i and ϕ_i . For simplicity, we adopt non-informative hyperpriors with $\beta_i = 1$ and $\phi_i = 10^{-6}$. In the non-informative limit—i.e., when $\beta_i \ll \vec{x}^T \Sigma_i^{-1} \vec{x}$ and $\phi_i \ll r_i$ —the resulting posterior distribution for α_i has mean and variance given by:

$$\mathbb{E}(\alpha_i) \approx \frac{2r_i}{\vec{x}^T \Sigma_i^{-1} \vec{x}}, \quad \text{and} \quad \text{Var}(\alpha_i) \approx \frac{\mathbb{E}(\alpha_i)}{r_i}. \quad (2.14)$$

Because the rank $r_i \sim N_r N_\theta$ is large in B-VFM applications, the resulting hyperpriors are sharply peaked and lead to stable, data-adaptive regularization. This effect, commonly referred to as *shrinkage*, ensures that each prior term contributes comparably to the total loss and accelerates convergence of the MCMC chains by reducing burn-in.

The Empirical Bayes Limit of B-VFM with Gamma Hyperpriors: EB-VFM

A simplified version of the hierarchical Bayesian method emerges in the limit $r_i \gg 1$, which could be a good approximation in B-VFM when working with highly resolved data or in potential 3D extensions of the algorithm. In this regime, the variance of the hyperparameters becomes negligible, $\text{Var}(\alpha_i) \approx 0$, and Algorithm 1 converges rapidly toward deterministic values of α_i . This formulation is analogous to an Empirical Bayes method [56], where the hyperparameters are treated as unknown but fixed quantities, estimated directly from the data rather than fully marginalized over.

The resulting Empirical Bayes Vector Flow Mapping (EB-VFM) method is grounded in a more rigid semi-Bayesian framework but it avoids the computational cost of running multiple MCMC chains, enabling substantially faster inference while retaining some level of adaptivity

Algorithm 1 MAP-Based Laplace-Approximate Gibbs Sampler

Data: Degrees of freedom: $N_d = 3N_r N_\theta$; measured field $\vec{X} \in \mathbb{R}^{N_d}$; measurement uncertainty $\Sigma \in \mathbb{R}^{N_d \times N_d}$; initial hyperparameters $\alpha_i^0 > 0$;

Burn-in iterations $it^b = 10$; total iterations $it^m = 50$.

Result: Posterior samples $\{\vec{x}_{\text{samp}}\} \subset \mathbb{R}^{N_d}$; hyperparameter samples $\{\alpha_i\} \subset \mathbb{R}_{>0}$

1: **Initialization:**

$\beta_i \leftarrow 1$; $\phi_i \leftarrow 10^{-6}$;

$\vec{\alpha}_i \leftarrow \alpha_i^0$;

$it \leftarrow 1$;

2: **while** $it \leq it^b + it^m$ **do**

Define loss: $\mathcal{J}(\vec{x}) = -\log p(\vec{x} \mid \vec{X}, \vec{\alpha})$

MAP estimate (Quasi-Newton): $\vec{x}^* \leftarrow \arg \min \mathcal{J}(\vec{x})$

Compute Hessian: $H \leftarrow \nabla^2 \mathcal{J}(\vec{x}^*)$

Regularize Hessian: $H_{\text{reg}} \leftarrow H'(\varepsilon)$, $\varepsilon = \min \{ \delta > 0 \mid H'(\delta) = \frac{1}{2}(H + H^T) + \delta I \succ 0 \}$

Cholesky decomposition: $H_{\text{reg}} = LL^T$, with L lower triangular

Sample \vec{x} : $\vec{y} \sim \mathcal{N}(0, I)$, $\delta \vec{x} \leftarrow L^{-1} \vec{y}$, $\vec{x}_{\text{samp}}^{(it)} \leftarrow \vec{x}^* + \delta \vec{x}$

For each prior term: compute $H_i \leftarrow \nabla^2 \mathcal{J}_i(\vec{x}^*)$, $r_i \leftarrow \text{rank}(H_i)$

Sample hyperparameters: $\alpha_i^{(it)} \leftarrow \Gamma \left(\beta_i + \frac{r_i}{2}, \phi_i + \vec{x}_{\text{samp}}^{(it)\top} H_i \vec{x}_{\text{samp}}^{(it)} \right)$

3: **if** $it > it^b$ **then**

Store sample: $\vec{x}_{\text{samp}}^{(it)}$, $\vec{\alpha}^{(it)}$

4: **end if**

$it \leftarrow it + 1$

5: **end while**

and uncertainty quantification. The posterior uncertainty in \vec{x} can be estimated directly from the diagonal entries of the covariance matrix evaluated at the MAP estimate \vec{x}^* , i.e., $\Sigma = 2[H(\vec{x}^*, \vec{\alpha})]^{-1}$. While the full Hessian may be too large to invert explicitly, efficient stochastic methods such as Hutchinson’s trace estimator [35] or probing techniques [75] can be used to approximate its diagonal.

Metropolis Hastings Sampling with Log-normal Hyperpriors

A more flexible sampling approach combines Metropolis-Hastings updates of the hyperparameters with the Laplace approximation of the conditional posterior described in § 2.1.2 (see Algorithm 2 for details). At each MCMC iteration, a new set of hyperparameters $\vec{\alpha}^*$ is proposed by drawing from log-normal distributions centered at the current values, ensuring positivity. The variance of these proposals is tuned adaptively as described below.

The conditional log-posterior is then optimized to obtain the MAP estimate \vec{x}^* , which is used to compute the acceptance probability based on the ratio of posterior probabilities:

$$\gamma = \min \left(1, \frac{p(\vec{x}^* | \vec{\alpha}^*)}{p(\vec{x}^{(it)} | \vec{\alpha}^{(it)})} \right).$$

The proposal $\vec{\alpha}^*$ is accepted with probability γ ; otherwise, the current state is retained. The acceptance rate is monitored every 10 iterations to adapt the standard deviation σ of the log-normal proposals: σ is halved if the rate falls below 1/3, and doubled if it exceeds 2/3. Crucially, because the prior covariance depends on \vec{x} , the normalization term $|\Sigma|^{-1/2}$ in the conditional posterior $p(\vec{x} | \vec{\alpha})$ may vary significantly across proposals. Therefore, it must be retained in the log-posterior calculation. The corresponding log-determinant is computed via Cholesky decomposition of the symmetrized, regularized H_{reg} described above.

2.1.3 The Log-Posterior Hessian Matrix of Bayesian Vector Flow Mapping

Near the inferred parameters the covariance matrix of the posterior pdf is approximately given by the inverse of its logarithm’s Hessian, i.e.,

$$p(\vec{x} | \vec{X}, \vec{\varepsilon}, \vec{\alpha}) \sim \exp \left[-J(\vec{x}; \vec{X}, \vec{\varepsilon}, \vec{\alpha}) \right] \approx \exp \left(-\vec{x}^T \Sigma^{-1} \vec{x} \right) \longrightarrow H = \frac{\partial^2 J}{\partial \vec{x}^2} = 2\Sigma^{-1}. \quad (2.15)$$

Algorithm 2 Metropolis-Hastings Sampling with Log-normal Hyperpriors

Data: Degrees of freedom: $N_d = 3N_r N_\theta$; measured field $\vec{X} \in \mathbb{R}^{N_d}$; measurement uncertainty $\Sigma \in \mathbb{R}^{N_d \times N_d}$; initial hyperparameters $\alpha_i^0 > 0$; Burn-in iterations $it^b = 10$; total iterations $it^m = 50$

Result: Posterior samples $\{\vec{x}_{\text{samp}}\} \in \mathbb{R}^{N_d}$; $\{\alpha_i\} \in \mathbb{R}_{>0}$

1: **Initialization:**

$\beta_i \leftarrow 1, \phi_i \leftarrow 10^{-6},$

$\alpha_i \leftarrow \alpha_i^0, \mu_i \leftarrow \log \alpha_i^0, \sigma \leftarrow 0.09,$

$\log P \leftarrow -1$

$it \leftarrow 0, itot \leftarrow 1$

2: **while** $it < it^b + it^m$ **do**

Define loss: $\mathcal{J}(\vec{x}) = -\log p(\vec{x} \mid \vec{X}, \vec{\alpha})$

MAP estimate: $\vec{x}^* \leftarrow \arg \min \mathcal{J}(\vec{x}), J_{\min} \leftarrow \mathcal{J}(\vec{x}^*)$

Compute Hessian: $H \leftarrow \nabla^2 \mathcal{J}(\vec{x}^*)$

Regularize Hessian: $H_{\text{reg}} \leftarrow H'(\varepsilon), \varepsilon = \min \{ \delta > 0 \mid H'(\delta) = \frac{1}{2}(H + H^T) + \delta I \succ 0 \}$

Cholesky: $H_{\text{reg}} = LL^T$, with L lower triangular

Compute log-determinant: $\log D \leftarrow \sum_{i=1}^n \log L_{ii}$

Compute log-posterior: $\log P^* \leftarrow -J_{\min} - \log D$

Acceptance probability: $\gamma \leftarrow \min(1, \exp(\log P^* - \log P))$

Draw uniform random number: $u \sim \mathcal{U}(0, 1)$

3: **if** $u < \gamma$ **then**

Accept: $it \leftarrow it + 1$

Store sample \vec{x} : $\vec{y} \sim \mathcal{N}(0, I), \delta \vec{x} \leftarrow L^{-1} \vec{y}, \vec{x}_{\text{samp}}^{(it)} \leftarrow \vec{x}^{(it)} + \delta \vec{x}$

Update log-posterior: $\log P \leftarrow \log P^*$

4: **end if**

Update acceptance rate: $AR \leftarrow it/itot$

5: **if** $\text{mod}(it, 10) = 0$ **then**

6: **if** $AR < 1/3$ **then**

$\sigma \leftarrow \sigma/2$

7: **end if**

8: **if** $AR > 2/3$ **then**

$\sigma \leftarrow 2\sigma$

9: **end if**

10: **end if**

Sample $\vec{\alpha}$: $\mu_i \leftarrow \log \alpha_i^{(it)} - \frac{\sigma^2}{2}, \alpha_i^* \sim \text{LogNormal}(\mu_i, \sigma)$

$itot \leftarrow itot + 1$

11: **end while**

Because the LV mask is an estimated parameter in B-VFM, i.e., m is the third block in \vec{x} , and priors like mass conservation and boundary conditions depend on m , the log posterior probability is not a bilinear mapping. Therefore, the calculation of its Hessian H is more involved than in other VFM methods. This section presents an analytical derivation of H . We proceed by expressing J as a sum of different contributions and taking the Hessian of each part:

$$\begin{aligned}
J &= J_{Dop} + J_{Mask} + J_{MC} + J_{BC} + J_{Reg} \\
J_{Dop} &\sim \frac{1}{2} \left[\log(|\Sigma_V|) + (v_r - V)^T \delta(m)^T \Sigma_V^{-1} \delta(m) (v_r - V) \right] \\
J_{Mask} &\sim \frac{1}{2} \left[\log(|\Sigma_M|) + (m - M)^T \Sigma_M^{-1} (m - M) \right] \\
J_{MC} &\sim \frac{1}{2} \left[-\log(|Q_{MC}|) + \alpha_{MC} \vec{v}^T D^T \delta(m)^2 D \vec{v} \right] \\
J_{BC} &\sim \frac{1}{2} \left[-\log(|Q_{BC}|) + \alpha_{BC} \vec{v}^T (Gm) \delta(\nabla m)^2 (m^T G^T) \vec{v} \right] \\
J_{Reg} &\sim \frac{1}{2} \left[-\log(|Q_{Reg}|) + \alpha_{Reg} (v_r^T L^T L v_r + v_\theta^T L^T L v_\theta + m^T L^T L m) \right]
\end{aligned} \tag{2.16}$$

Computing the Hessian of the mask and regularization terms, J_{Mask} and J_{reg} is simple since these terms are still bilinear mappings in B-VFM. Thus, we have

$$H_{Mask} = \frac{\partial^2 J_{Mask}}{\partial \vec{x}^2} = 2 \Sigma_M^{-1} \otimes I_{mm}^{3 \times 3}, \tag{2.17}$$

$$H_{reg} = \frac{\partial^2 J_{reg}}{\partial \vec{x}^2} = 2 L^T L \otimes (\alpha_{sv} I_{v_r, v_r}^{3 \times 3} + \alpha_{sv} I_{v_\theta, v_\theta}^{3 \times 3} + \alpha_{sm} I_{m, m}^{3 \times 3}). \tag{2.18}$$

The Doppler Hessian has the usual v_r diagonal term but also m -terms,

$$\begin{aligned}
H_{Dop} &= \frac{\partial^2 J_{Dop}}{\partial \vec{x}^2} = 2 \delta(m \circ m) \circ \Sigma_V^{-1} \otimes I_{v_r, v_r}^{3 \times 3} + 2 \delta[(v_r - V) \circ (v_r - V)] \circ \Sigma_V^{-1} \otimes I_{m, m}^{3 \times 3} + \\
&\quad + 2 \delta[(v_r - V) \circ m] \circ \Sigma_V^{-1} \otimes (I_{v_r, m}^{3 \times 3} + I_{m, v_r}^{3 \times 3}).
\end{aligned} \tag{2.19}$$

The $I_{m, m}$ term penalizes the mask in pixels where v_r departs from the measured V , whereas the off-diagonal terms reflect the covariance between v_r and m .

The mass conservation Hessian is

$$\begin{aligned}
H_{MC} &= \frac{\partial^2 J_{MC}}{\partial \vec{x}^2} = 2 D^T \delta(m \circ m) D + 2 \delta(D \vec{v} \circ D \vec{v}) I_{m, m}^{3 \times 3} + \\
&\quad + 2 \delta(D \vec{v} \circ m) \left[D_r \otimes (I_{\vec{v}_r, m}^{3 \times 3} + I_{m, \vec{v}_r}^{3 \times 3}) + D_\theta \otimes (I_{v_\theta, m}^{3 \times 3} + I_{m, v_\theta}^{3 \times 3}) \right]
\end{aligned} \tag{2.20}$$

The first term in this Hessian comes from the bilinear dependence on (v_r, v_θ) and has off-diagonal terms that represent the covariance between the two velocity components arising from mass conservation. The second term penalizes the mask in pixels with strong departures from mass conservation. The last terms reflect that enforcing mass conservation inside the LV mask introduces covariance between \vec{v} and m .

The formal derivation of the boundary condition Hessian is straightforward but the block structure of the parent equation 2.1 complicates its matrix arrangement. Therefore, we pay more attention to its derivation. We begin by defining the matrix

$$N = D_r m \otimes I_{v_r, v_r}^{3 \times 3} + D_\theta m \otimes I_{v_r, v_\theta}^{3 \times 3}, \quad (2.21)$$

which allows us to rewrite the J_{BC} as $J_{BC} = \vec{x}^T N^T N \vec{x}$. Since N only depends on m , this transformation yields the part of H_{BC} corresponding to the velocities.

$$H_{BC}|_{\vec{v}, \vec{v}} = \frac{\partial^2 J_{BC}}{\partial \vec{v}^2} = 2N^T N. \quad (2.22)$$

The remaining terms are obtained from the m -derivatives of N as follows,

$$H_{BC}|_{m, m} = \frac{\partial^2 J_{BC}}{\partial m^2} = 2 [D_r D_r^T \delta(v_r \circ v_r) + D_\theta D_\theta^T \delta(v_\theta \circ v_\theta) + (D_r D_\theta^T + D_\theta D_r^T) \delta(v_r \circ v_\theta)] \otimes I_{mm}^{3 \times 3}, \quad (2.23)$$

$$H_{BC}|_{v_r, m} = H_{BC}|_{v_r, m}^T = \frac{\partial^2 J_{BC}}{\partial v_r \partial m} = 2 \{ [2\delta [(D_r m) \circ v_r] + \delta [(D_\theta m) \circ v_\theta]] D_r + \delta [(D_r m) \circ v_\theta] D_\theta \} I_{v_r, m}^{3 \times 3}, \quad (2.24)$$

and

$$H_{BC}|_{v_\theta, m} = H_{BC}|_{v_\theta, m}^T = \frac{\partial^2 J_{BC}}{\partial v_\theta \partial m} = 2 \{ [2\delta [(D_\theta m) \circ v_\theta] + \delta [(D_r m) \circ v_r]] D_\theta + \delta [(D_\theta m) \circ v_r] D_r \} I_{v_\theta, m}^{3 \times 3}, \quad (2.25)$$

These analytical expressions were validated term by term versus numerical calculations of the Hessian from the fully non-linear J for small meshes of size $N_r, N_\theta \lesssim 32$. The numerical calculation of the Hessian was performed within the quasi-Newton iterative procedure using

values from consecutive iterations and a central finite difference scheme. Differences between the analytical and numerical Hessians were negligible for all terms. For larger mesh sizes, computing the full Hessian matrix of size $3N_r N_\theta \times 3N_r N_\theta$ becomes prohibitively expensive.

$$I_{v_r, v_r}^{3 \times 3} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad I_{v_\theta, v_\theta}^{3 \times 3} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad I_{mm}^{3 \times 3} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (2.26)$$

$$I_{v_r, v_\theta}^{3 \times 3} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad I_{v_r, m}^{3 \times 3} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad I_{v_\theta, m}^{3 \times 3} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}, \quad (2.27)$$

and $I_{v_\theta, v_r}^{3 \times 3} = (I_{v_r, v_\theta}^{3 \times 3})^T$, $I_{m, v_r}^{3 \times 3} = (I_{v_r, m}^{3 \times 3})^T$, $I_{m, v_\theta}^{3 \times 3} = (I_{v_\theta, m}^{3 \times 3})^T$.

2.1.4 Validation with Synthetic Flow Fields

Because the diastolic vortex is both physiologically important and characterized by a relatively simple geometric structure, it provides an ideal benchmark for testing flow reconstruction methods. To this end, we evaluated Vanilla VFM and B-VFM using two synthetic flow fields that approximate this pattern: the Lamb–Chaplygin (LC) dipole and Hill’s vortex. These fields were generated with the MATLAB Ultrasound Toolbox (MUST), which allowed us to simulate controlled, physiologically inspired flow conditions. From each synthetic field, the radial velocity component was extracted and treated as the ‘color-Doppler’ input to both Vanilla VFM and B-VFM. This setup enabled us to directly assess the ability of each method to reconstruct the full velocity field from Doppler-like inputs. Further details on the construction of the synthetic fields, the measurement assumptions, and the reconstruction algorithms are provided in Chapter 3: Theoretical Analysis.

2.2 Results

2.2.1 Lamb Chaplygin Dipole

To validate our method, we introduce an 2D ideal synthetic field as a substitute for the Doppler image, where the radial velocity component of the Lamb-Chaplygin Dipole serve as the color-Doppler input. We look at three cases: 1) An ideal noiseless case to compare BVFM directly with Vanilla VFM, and 2) clinically relevant cases 2) An input color-Doppler data set with untrustworthy velocity regions due to data voids and 3) an untrustworthy boundary conditions due to poor segmentation.

Ideal Case: A direct comparison between Vanilla VFM and BVFM

The purpose of the ideal case is to illustrate that, when no additional information on input data uncertainty is available, BVFM yields results equivalent to those of Vanilla VFM. Results from three chains at different initial conditions were averaged for the final BVFM fields. Due to the well behaved nature of the idealized, noiseless case, Algorithm 1, the Gibbs approach was used. Point wise comparison of reconstructed velocity fields and ground truth fields (Figure 2.3) show that for the ideal case, BVFM produced r^2 values of $u_R = 0.9969$ and $u_\theta = 0.9822$ and Vanilla VFM produced $r_r^2 = 0.9997$ and $r_\theta^2 = 0.9725$. Due to 2D idealized nature of the LC Dipole, Vanilla VFM and BVFM reconstructs the 2D flow field with high accuracy, as expected.




Figure 2.2: Ideal: Hyperprior Markov Monte Carlo Chains (iterations log scale)

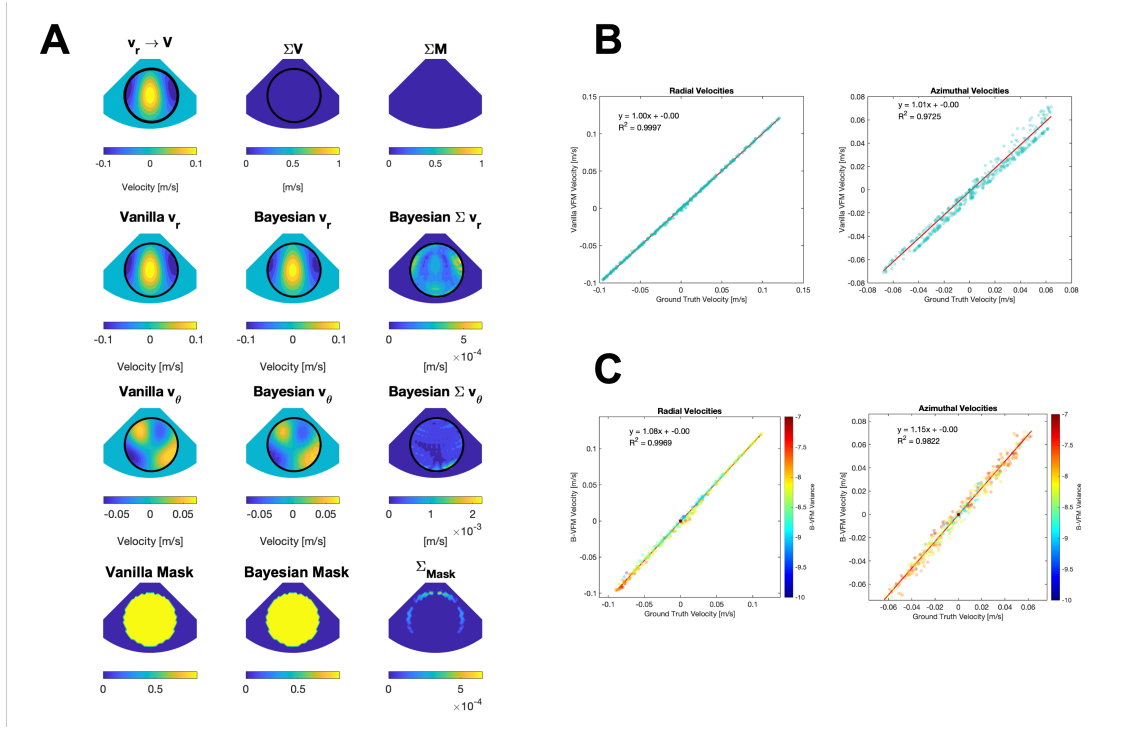


Figure 2.3: Ideal: A) Vanilla-VFM and Bayesian VFM Inputs and Outputs, B) Vanilla-VFM vs Ground Truth Correlation Plots, C) Bayesian-VFM vs Ground Truth Correlation Plots

Doppler Data Void Scenario

To simulate the cases of Doppler data voids, we introduce a hole into our synthetic radial velocity field and process it using both the Vanilla VFM and BVFM pipelines. We run seven independent MCMC chains with different initial conditions with the Gibbs approach (Algorithm 1). Not only did each chain converge, but they all converged to similar posterior estimates, indicating robust and consistent inference across runs. Due to the well behaved nature of the boundary conditions, the Gibbs approach showed to be sufficient in reconstructed the ground truth flow fields. To be thorough, future work should test the Metropolis Hastings approach (Algorithm 2) to compare convergence between the two algorithms. In Figure 2.5, we compare the velocity fields from the Vanilla VFM method to the

averaged results from seven BVFM chains. The BVFM approach not only reconstructs the velocity field more accurately in a qualitative sense but also provides additional information by quantifying uncertainty through posterior variance. Pointwise correlation with the ground truth reveals that the BVFM method achieves higher accuracy, with correlation coefficients of $r^2 = 0.975$ for the u_R field and $r^2 = 0.927$ for the u_θ field (Figure 2.5C). In contrast, the Vanilla iVFM method yielded lower correlation values of $r_r^2 = 0.940$ and $r_\theta^2 = 0.739$ (Figure 2.5B).




Figure 2.4: Doppler Hole: Hyperprior Markov Monte Carlo Chains (iterations log scale)

Poor Segmentation Scenario

Vanilla-VFM heavily depends on properly defined boundary conditions due to the inherent singularity of the system matrix. In cases where boundary conditions are incomplete or ill-posed, the method must rely on regularization to produce a non-singular and stable solution. The Lamb Chaplygin (LC) dipole case is particularly well suited for Vanilla-VFM, as the boundary conditions are naturally satisfied within the synthetic dataset. However, when we intentionally alter the boundary conditions to challenge the reconstruction, the method's sensitivity to hyperparameters, particularly the boundary condition weight, becomes evident. The vanilla approach treats the segmentation as deterministic, assuming no error and independence from the velocity field, while searching for alpha terms that minimize the overall loss. With low boundary condition enforcement, the resulting velocity field is physically

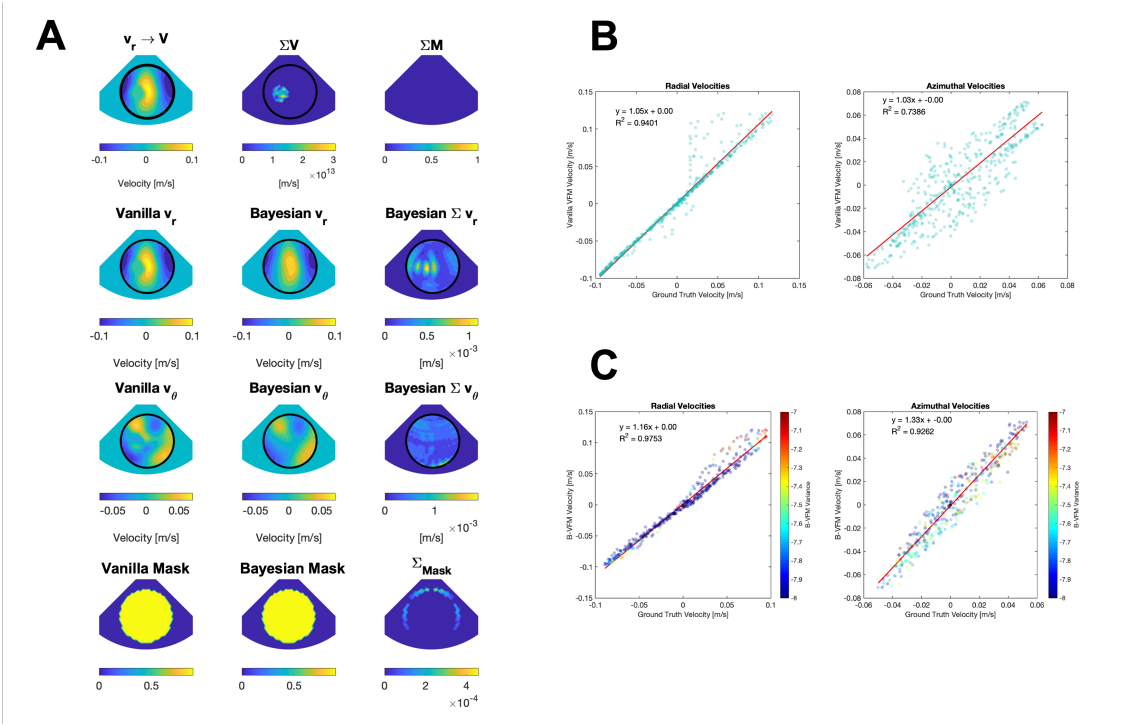


Figure 2.5: Doppler Hole: A) Vanilla-VFM and Bayesian VFM Inputs and Outputs, B) Vanilla-VFM vs Ground Truth Correlation Plots, C) Bayesian-VFM vs Ground Truth Correlation Plots

inconsistent: velocities are not aligned with the boundary, and the reconstruction overprescribes smoothing. In contrast, high boundary condition enforcement attempts to satisfy boundaries in regions that are incorrectly segmented, leading to difficulties in the azimuthal reconstruction, a field that is farther from the ground truth, and incomplete correction of the boundary conditions (Figure 2.6). These observations highlight how the choice of hyperparameters critically affects the reconstructed field.

This sensitivity motivates the use of a Bayesian framework, where uncertainty in the input data is explicitly modeled and hyperparameters are not fixed a priori. Our method allows the mask to update dynamically as new information becomes available, creating bidirectional interactions: the mask is influenced by the velocity field, while the velocity reconstruction is shaped by the evolving mask. Although this introduces nonlinearity and increases compu-

tational complexity, it captures more realistic interactions between the two fields. In cases with more complex boundaries, we rely on the Metropolis-Hastings algorithm, which is non-conjugate and more flexible, unlike the gamma-based conjugate approach (Algorithm 1) that converges quickly but can produce strong shrinkage, narrow posteriors, local minima, and unresolved flow (Figures 2.7, 2.8). For scenarios with poor wall segmentation (Algorithm 2), B-VFM reconstructed a velocity field with r^2 values of 0.9903 and 0.8666 in the radial and azimuthal directions, showing strong agreement with ground truth (Figures 2.9, 2.10). Additionally, B-VFM reconstructs a non-binary mask that better represents uncertainties inherent to real-world data: we are more confident in the LV location at central regions, while near the edges and wall boundaries the mask exhibits higher variance, reflecting increased uncertainty in these areas.

2.2.2 Synthetic Scenarios Summary

In the ideal, noiseless case, B-VFM performs comparably to Vanilla VFM, achieving similar accuracy in both v_r and v_θ components. However, in scenarios with data voids or poor segmentation, B-VFM outperforms Vanilla VFM by more effectively reconstructing the velocity field, rather than simply oversmoothing. This improved performance arises from B-VFM’s ability to balance physical constraints with regularization, while Vanilla VFM remains highly dependent on the choice of hyperparameters (*). Unlike Vanilla VFM, B-VFM is not fixed to a single hyperparameter value and additionally provides uncertainty maps for the reconstructed outputs, offering insight into regions where the reconstruction is less certain. Overall, B-VFM performs better or at least comparable to Vanilla VFM across a range of conditions, while providing the added benefit of principled uncertainty quantification.

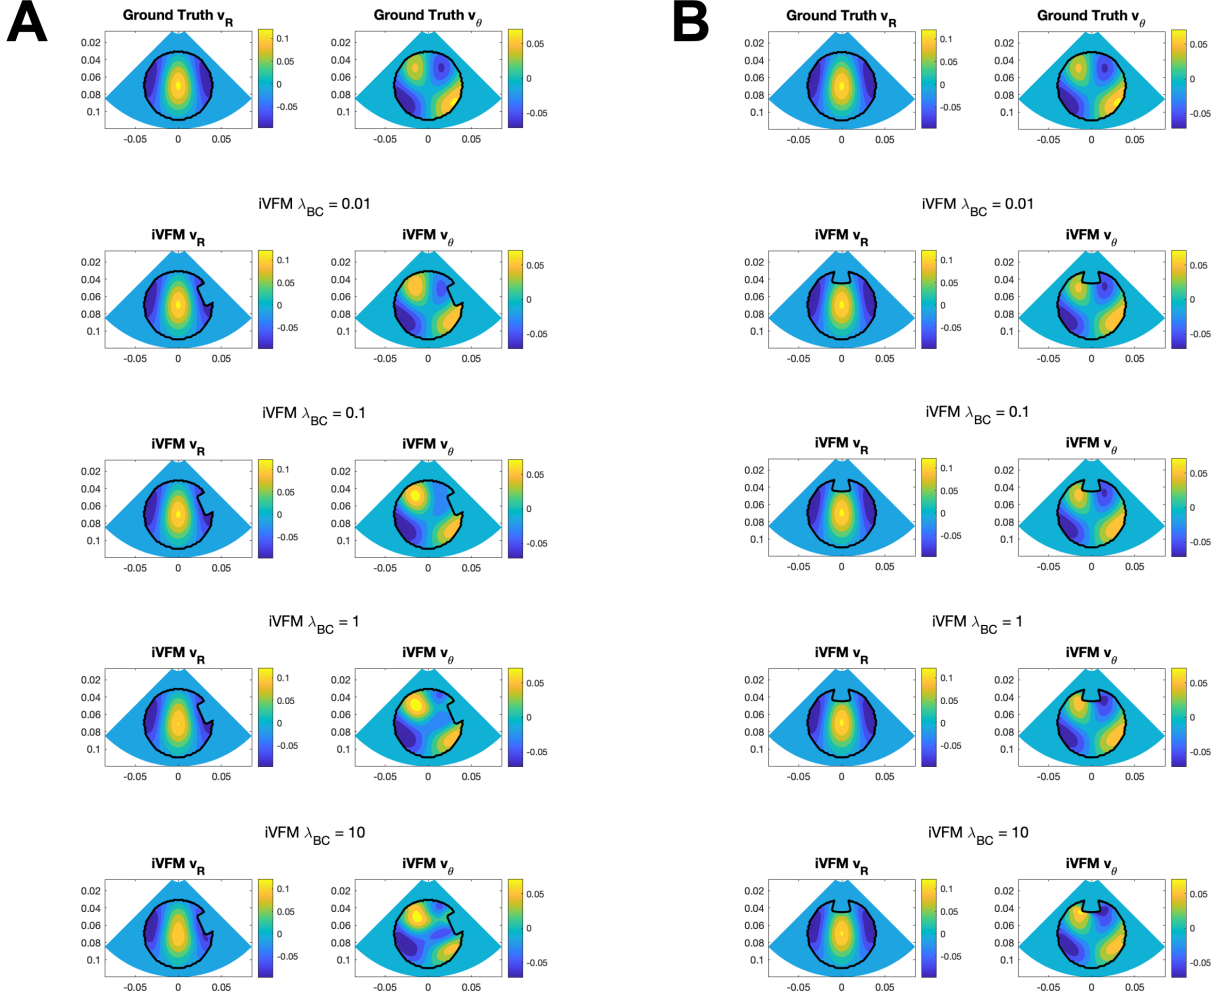


Figure 2.6: Vanilla-VFM: Impact of boundary condition enforcement with poor wall segmentation. A) Poor side wall segmentation, B) Poor "apex" wall segmentation.

Table 2.1: Comparison of r^2 values for Vanilla-VFM and Bayesian VFM (B-VFM) under different synthetic scenarios.

	Ideal		Data Void		Poor Segmentation	
	Vanilla VFM	B-VFM	Vanilla VFM	B-VFM	Vanilla VFM	B-VFM
v_r	$r^2 = 0.9997$	$r^2 = 0.9969$	$r^2 = 0.9401$	$r^2 = 0.9753$	*	$r^2 = 0.9903$
v_θ	$r^2 = 0.9725$	$r^2 = 0.9822$	$r^2 = 0.7386$	$r^2 = 0.9262$	*	$r^2 = 0.8666$




Figure 2.7: Poor Segmentation Scenario with Gamma Approach: Hyperprior Markov Monte Carlo Chains (iterations in log scale) Algorithm 1

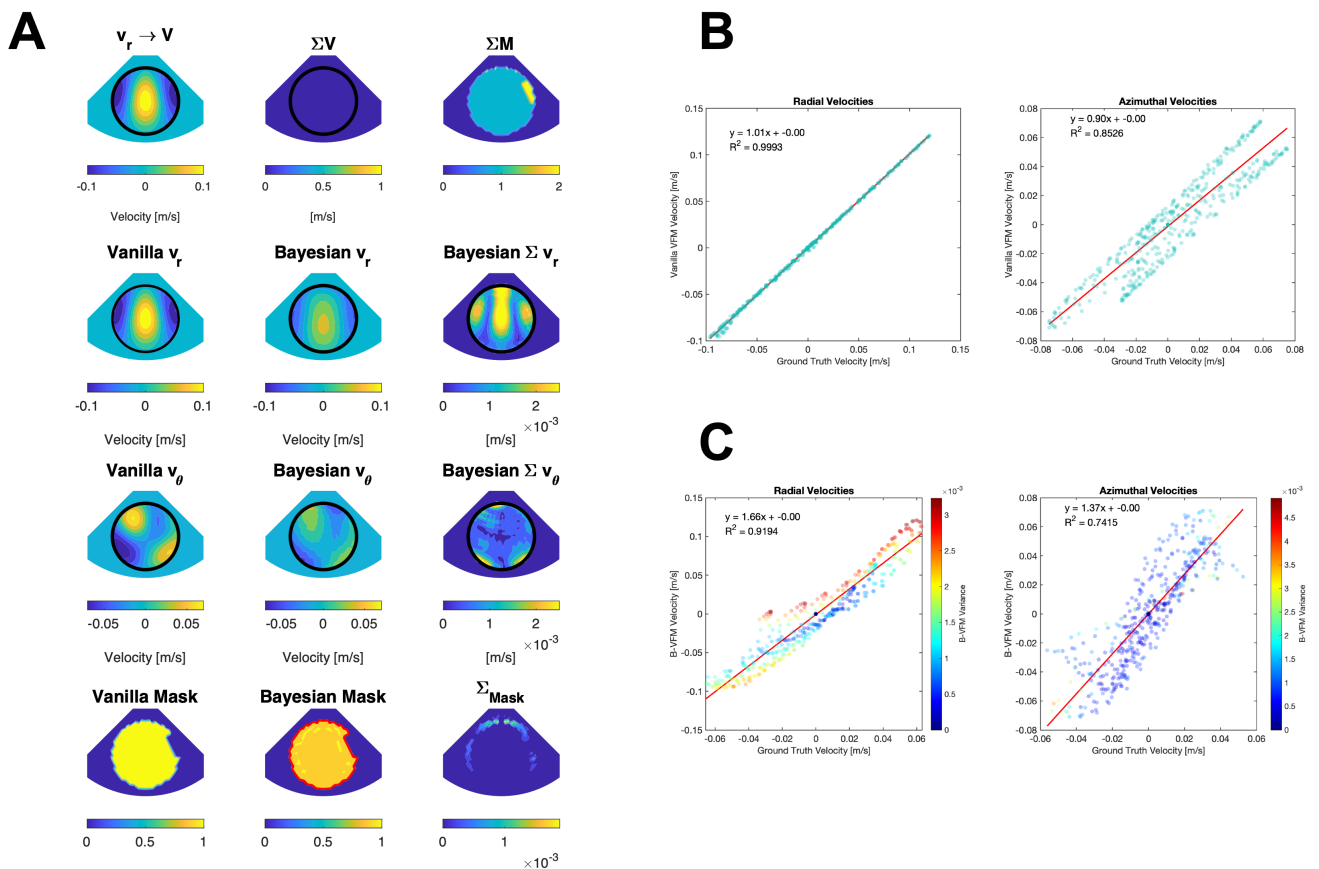


Figure 2.8: Poor Segmentation Scenario with Gamma Approach: A) Vanilla-VFM and Bayesian VFM Inputs and Outputs, B) Vanilla-VFM vs Ground Truth Correlation Plots, C) Bayesian-VFM vs Ground Truth Correlation Plots (Algorithm 1)




Figure 2.9: Poor Segmentation Scenario with Metropolis Hastings Approach: Hyperprior Markov Monte Carlo Chains (iterations in log scale) Algorithm 2

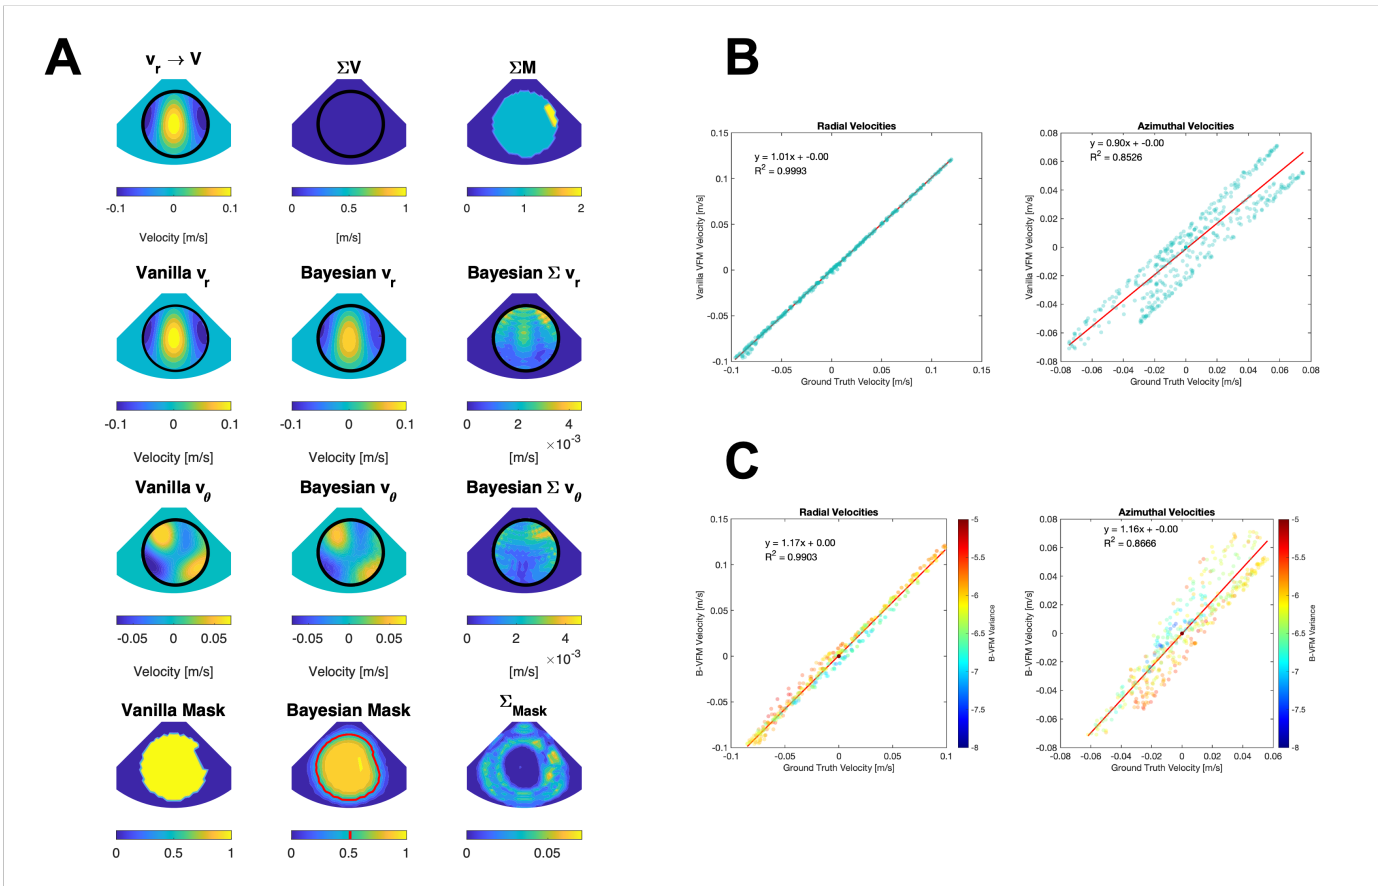


Figure 2.10: Poor Segmentation Scenario with Metropolis Hastings Approach: A) Vanilla-VFM and Bayesian VFM Inputs and Outputs, B) Vanilla-VFM vs Ground Truth Correlation Plots, C) Bayesian-VFM vs Ground Truth Correlation Plots (Algorithm 2)

2.2.3 3D Hill's Vortex

So far, our evaluations have focused on 2D idealized cases, such as the Lamb-Chaplygin dipole, which satisfy 2D continuity and are naturally compatible with the planar flow assumptions used in vector flow mapping. Here, we extend our analysis to the 3D Hill's spherical vortex, which directly challenges these 2D assumptions. The Hill's spherical vortex is an axisymmetric 3D incompressible flow where the velocity components are given by:

$$\psi = \begin{cases} -\frac{3U}{4}\left(1 - \frac{r^2}{a^2}\right)r^2 \sin^2 \theta, & \text{in } r \leq a. \\ \frac{U}{2}\left(1 - \frac{a^3}{r^3}\right)r^2 \sin^2 \theta, & \text{in } r \geq a. \end{cases} \quad (2.28)$$

$$v_r = \frac{1}{r^2 \sin \theta} \frac{\partial \psi}{\partial \theta} \quad v_\theta = -\frac{1}{r \sin \theta} \frac{\partial \psi}{\partial r} \quad v_\phi = 0$$

While the flow field only depends on (r, θ) , the flow still exists in 3D with symmetry around the ϕ -direction. Thus, the 2D planar flow assumption in vector flow mapping is not valid for this field. This makes the Hill's spherical vortex a good choice for testing the limitations of the 2D planar flow assumption. Analytically we know that the 2D planar assumption lacks both a radial and polar factor that exists in axisymmetric incompressibility. We calculate the error as follows:

2D planar incompressibility (D_{2D}):

$$\frac{\partial v_r}{\partial r} + \frac{1}{r} \frac{\partial v_\theta}{\partial \theta} = 0$$

Axisymmetric incompressibility (D_{axisym}):

$$\frac{1}{r^2} \frac{\partial}{\partial r} (r^2 v_r) + \frac{1}{r \cos \theta} \frac{\partial}{\partial \theta} (v_\theta \cos \theta) = 0$$

The error in the 2D planar flow assumption is given by

$$Error = D_{axisym} - D_{2D} \quad (2.29)$$

$$= \left(\frac{1}{r^2} \frac{\partial}{\partial r} (r^2 v_r) + \frac{1}{r \cos \theta} \frac{\partial}{\partial \theta} (v_\theta \cos \theta) \right) - \left(\frac{\partial v_r}{\partial r} + \frac{1}{r} \frac{\partial v_\theta}{\partial \theta} \right) \quad (2.30)$$

$$= \frac{2v_r}{r} + v_\theta \tan \theta \quad (2.31)$$

Figure 2.11 shows the error in 2D planar flow assumption (log scale) solved from Equ. 2.31. The blue histogram shows the divergence error from Vanilla VFM, while the red histogram corresponds to Bayesian VFM divergence error. The longer tails in the Vanilla VFM distribution indicate a higher frequency of large error outliers compared to BVFM. Although BVFM does not fully capture the 3D nature of the flow (i.e., achieving zero error), it generally performs better than VFM because it models priors such as mass conservation probabilistically rather than enforcing them strictly. This softer constraint provides greater flexibility to adapt solutions locally, helping to reduce extreme errors.

In summary, by modeling physical priors probabilistically, rather than enforcing them strictly, our Bayesian approach provides the flexibility to adapt locally to complex flow features. This softer constraint allows the method to reduce extreme errors without aggressively imposing conditions that may conflict with the underlying flow physics.

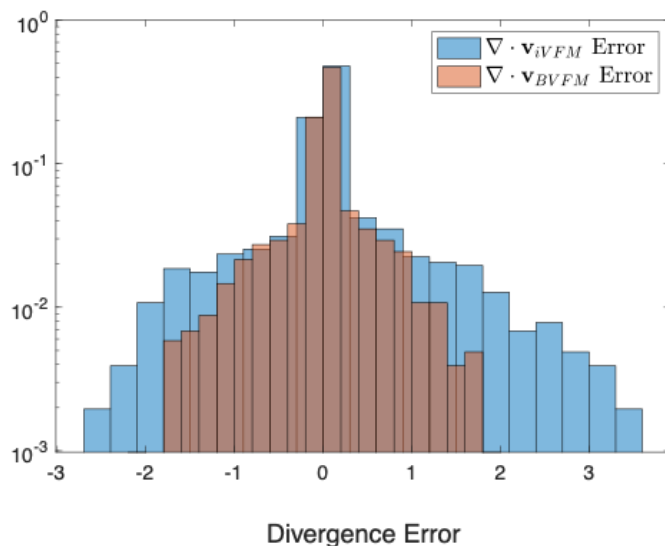


Figure 2.11: Error in 2D planar flow assumption (log scale) Equ. 2.31. The blue histogram shows the divergence error from Vanilla VFM, while the red histogram corresponds to Bayesian VFM divergence error.

2.2.4 Clinical Scenario: Ejection

2D echocardiographic studies were performed using a Vivid 7 scanner with 2- to 4-MHz transducers (GE Healthcare). Conventional B-mode and Doppler data were measured following current recommendations [41]. The LV myocardial wall was segmented from LV from B-mode series using speckle-tracking software (EchoPac 214, GE Healthcare). Flow field in the LV was calculated, combining color Doppler and the endocardial boundary using VFM [25]. Data was acquired three times in each patient.

We applied both the Vanilla VFM and B-VFM methods to a representative time frame during the ejection phase of the cardiac cycle for qualitative comparison. Input variances were derived from interobserver variability, calculated across three independent color-Doppler acquisitions and segmentations (Figure 2.12A,B). While interobserver variability provides a reasonable first choice for specifying input variance in the Bayesian framework, it represents only one option; in practice, the selection of measurement error models should be carefully tailored to the clinical context.

Both methods produced similar reconstructed flow fields, capturing the primary left ventricular velocity patterns. However, the Bayesian approach additionally resolved flow structures in regions of the ventricle where no color-Doppler data were available, highlighting its ability to infer plausible velocities in areas of missing information. The resulting variance maps reflected the chosen input variances: for example, high variability in the right side of the ventricle, arising from segmentation differences, propagated into both the mask variance and the reconstructed velocity field due to the coupling between boundary conditions and velocity estimates in our formulation. This was clearly visible in the velocity variance maps, where the same region exhibited elevated uncertainty (Figure 2.12D).

Although this clinical example was limited to a single frame and used input error models that were not necessarily optimal or clinically rigorous, it nonetheless serves as an important proof of concept for future applications of B-VFM. A key advantage of the Bayesian approach is that it provides not only reconstructed velocity components (v_r , v_θ) and the mask (M),

but also associated variance maps. These maps quantify uncertainty in both the velocity field and the mask, offering a probabilistic representation of regions where the reconstruction is less reliable. In contrast, the Vanilla VFM relies on a deterministic mask and cannot account for such uncertainty. The ability of B-VFM to capture and propagate measurement variability therefore yields a richer, more clinically meaningful representation of flow, with potential downstream use in secondary analyses such as residence time estimation or risk assessment of abnormal flow patterns.

Future work should extend this analysis across the full cardiac cycle to capture the temporal evolution of both reconstructed flow and associated uncertainties. Quantifying the impact of uncertainty modeling on derived hemodynamic metrics, such as residence time, will further clarify its clinical value. Finally, validation against other modalities, including 4D MRI or CT, will be essential to benchmark the accuracy and establish the relevance of B-VFM in clinical practice.

2.3 Ill-posedness and Regularization of VFM

Vector Flow Mapping (VFM) poses an ill-posed inference problem due to the absence of direct measurements for the azimuthal velocity component v_θ , rendering the observation matrix A in equation 2.1 severely rank-deficient. This challenge is analogous to optical flow estimation in directions perpendicular to image gradients, but more severe than in applications such as phase-contrast MRI, where all flow components are measured and A is typically full-rank. The divergence operator—VFM’s principal physical constraint—is itself singular and insufficient for regularization, often leading to spurious v_θ artifacts unless supplemented with smoothing and boundary conditions.

This section presents a theoretical analysis of the negative log-posterior Hessian in Bayesian VFM, which is proportional to the precision matrix Σ^{-1} and equivalent to the normal matrix $Q^\top Q$ in least-squares formulations [5] of the form $Q^\top Q \vec{x} = Q^\top \vec{b}$. We examine the structure of its inverse, proportional to the covariance matrix Σ , to formalize the ill-posedness of VFM and assess its impact on inference and uncertainty.

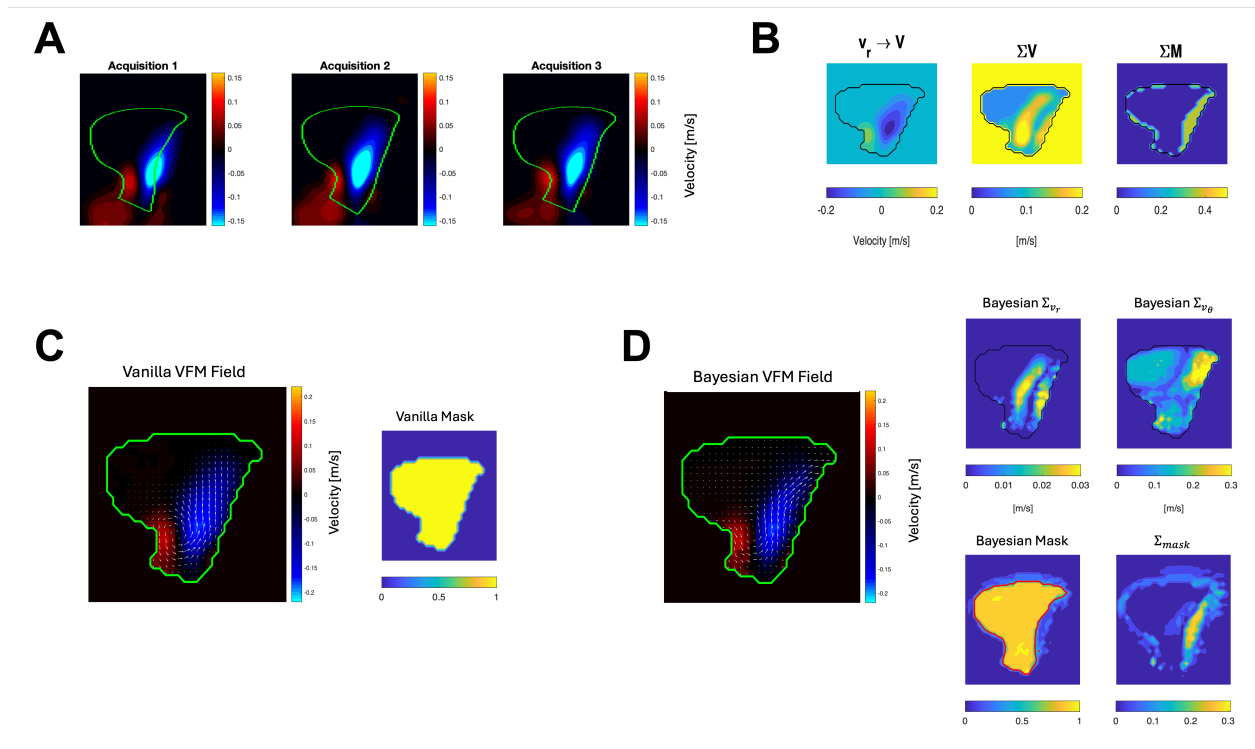


Figure 2.12: Clinical Case (Ejection): A) Ejection color-Doppler acquisitions and segmentations, B) B-VFM Input color-Doppler and data errors, C) Vanilla VFM reconstructed field and input deterministic mask, D) B-VFM reconstructed velocity field and mask with variance maps.

We derive the null space of the mass conservation constraint to characterize the spatial patterns of VFM artifacts and analyze how boundary conditions modify this null space. This yields precise conditions on the LV mask geometry required to ensure regularity without the need for smoothing. Finally, we demonstrate that although smoothing lacks a physical basis, it robustly mitigates ill-posedness in a mask-independent manner. Readers primarily interested in the practical use of VFM may prefer to skip the detailed derivations and refer instead to the subsection summaries, which are written to be more accessible and interpretation-focused.

Structure of the VFM Covariance Matrix

Here, we investigate the structure of the log-conditional posterior Hessian inverse for VFM, which is proportional to the covariance matrix at the MAP estimate, $H = 2\Sigma^{-1}(\bar{x}^*)$. To simplify our analysis, we start with the block decomposition of H

$$H = \begin{bmatrix} H_{rr} & H_{r\theta} & H_{rm} \\ H_{r\theta} & H_{\theta\theta} & H_{\theta m} \\ H_{rm} & H_{\theta m} & H_{mm} \end{bmatrix}, \quad \text{with } H_{ij} = H_{ji}^\top.$$

See §2.1.3 for an explicit derivation of this Hessian matrix. To keep the analysis concise, we further assume we are working with a good-quality image with low Doppler and mask observation errors, i.e., $\Sigma_V, \Sigma_M \ll 1$, so that

$$H = \begin{bmatrix} \sigma^{-1}I & \alpha\tilde{H}_{r\theta} & \alpha\tilde{H}_{rm} \\ \alpha\tilde{H}_{r\theta} & \alpha\tilde{H}_{\theta\theta} & \alpha\tilde{H}_{\theta m} \\ \alpha\tilde{H}_{rm} & \alpha\tilde{H}_{\theta m} & \sigma^{-1}I \end{bmatrix},$$

where $\sigma \sim \Sigma_V \sim \Sigma_M$, α represents a generic strength of the VFM priors, i.e., the hyper-parameters, and the tilde is used to indicate that the Hessian blocks have been normalized with α . Using a Neumann expansion, the inverse of H is approximated as

$$H^{-1} \approx \begin{bmatrix} \sigma I & -\alpha\sigma\tilde{H}_{r\theta}\tilde{H}_{\theta\theta}^{-1} & -\alpha\sigma^2\tilde{H}_{rm} \\ -\alpha\sigma\tilde{H}_{\theta\theta}^{-1}\tilde{H}_{r\theta}^\top & \alpha^{-1}\tilde{H}_{\theta\theta}^{-1} & -\sigma\tilde{H}_{\theta\theta}^{-1}\tilde{H}_{\theta m} \\ -\alpha\sigma^2\tilde{H}_{rm}^\top & -\sigma\tilde{H}_{\theta m}^\top\tilde{H}_{\theta\theta}^{-1} & \sigma I \end{bmatrix}. \quad (2.32)$$

In contrast, a data assimilation problem with observation in both v_r and v_θ would have a precision matrix of the form

$$H_{\text{observed v.}}^{-1} \approx \sigma I - \sigma^2\alpha \begin{bmatrix} \tilde{H}_{rr} & \tilde{H}_{r\theta} & \tilde{H}_{rm} \\ \tilde{H}_{r\theta}^\top & \tilde{H}_{\theta\theta} & \tilde{H}_{\theta m} \\ \tilde{H}_{rm}^\top & \tilde{H}_{\theta m}^\top & \tilde{H}_{mm} \end{bmatrix}.$$

This difference has an important implication. The Neumann approximation for H^{-1} breaks down and H becomes singular in VFM when the weights associated to mass conservation,

boundary conditions, and smoothing are very small, i.e., $\alpha \ll 1$, a reminder of the impossibility of inferring v_θ without observing this velocity component or imposing physical constraints. This issue does not arise in data assimilation with observed v_θ .

Focusing on the VFM Hessian (equation 2.32), we note that only the (r, m) blocks in H^{-1} are quadratic in σ , indicating that Doppler and mask are more weakly coupled than other variables. The coupling block \tilde{H}_{rm} arises from the interdependence between mask and Doppler velocity in the boundary conditions (equation 2.24). Detailed analysis indicates that this coupling is activated at pixels where the mask gradient and radial velocity fail to satisfy the free-slip boundary condition prior, and that its magnitude scales as $\mathcal{O}(\Sigma_V \Sigma_M \alpha_{BC})$, reflecting weaker coupling when either the Doppler velocity or LV mask are well observed (i.e., low error), and stronger coupling when the boundary condition prior is more strongly enforced. In traditional VFM methods, there is no possible coupling since the mask is not a latent variable, which effectively sets $\Sigma_M = 0$. Therefore, these methods cannot correct segmentation errors using velocity data.

Also, the diagonal of H^{-1} provides an estimate of the posterior uncertainty in the latent variables (at least when the hyperparameter variance is low as expected for $\sigma \ll 1$). In the limit of small measurement noise ($\sigma \ll 1$), our analysis indicates that the uncertainty in v_r and m is uniquely dictated by their measurement noise. In contrast, the uncertainty in v_θ depends entirely on the hyperparameters and the inverse of the central block, $\tilde{H}_{\theta\theta}$, which encodes the θ -component of the mass conservation, boundary conditions, and smoothing priors. This finding underscores the importance of rigorous Bayesian inference of hyperparameters to accurately quantify uncertainty in the inferred v_θ .

Effect of Priors on VFM Regularization and Regularity Criteria

Equation 2.32 shows that, at leading order, the regularity of H depends solely on the regularity of the central block $\tilde{H}_{\theta\theta}$. Here, we analyze the regularity of this matrix by finding the singular eigenvectors of the θ -component of the discretized divergence operator (i.e., spurious spatial patterns that could be arbitrarily added to v_θ without breaking mass con-

servation) and computing how the boundary conditions and smoothing priors modify the singular eigenvalue (i.e., how those priors enforce that the strength of the spurious patterns should be zero).

To simplify our derivations, we will work first in 1D and will later extend key results to the 2D case relevant to VFM. Let $D_\theta \in \mathbb{R}^{N_\theta \times N_\theta}$ be the centered finite difference approximation of the derivative operator $\partial/\partial\theta$ with periodic boundary conditions. This operator is circulant and its nullspace is defined by the constant $N_\theta \times 1$ vector $\vec{1} = (1, \dots, 1)$, i.e., $D_\theta \vec{1} = \vec{0}$, so that $D_\theta^\top D_\theta \vec{1} = 0$. This vector would generally belong to the nullspace of any 1D smoothing operator as well, so boundary conditions would be the only regularizer in 1D. Bayesian enforcement of a free-slip boundary condition $\vec{n}^\top \vec{v} = 0$ modifies leads to the normal equation to

$$H_{\theta\theta} \vec{v}_\theta = (\alpha_{MC} D_\theta^\top D_\theta + \alpha_{BC} B_\theta^\top B_\theta) \vec{v}_\theta = \text{rhs},$$

where $B_\theta = \text{diag}(n_\theta) = \text{diag}(D_\theta \vec{m})$ is a diagonal matrix encoding the θ -component of the normal vector to the mask, used to construct a least-squares penalty on the flux across the mask, while α_{MC} and α_{BC} are the hyperparameters controlling each penalty's strength. Thus, imposing boundary conditions can be seen as a rank update to the precision matrix. This update affects the zero-valued eigenvalue of D_θ , corresponding to the constant vector $\vec{1}$, by computing

$$\tilde{\lambda}_{\vec{1}} = \vec{1}^\top H_{\theta\theta} \vec{1} = \vec{1}^\top D_\theta^\top D_\theta \vec{1} + \alpha_{BC} \cdot \vec{1}^\top B_\theta^\top B_\theta \vec{1} = 0 + \alpha_{BC} \cdot \sum n_\theta^2.$$

Extending this analysis to two dimensions is straightforward because, while the nullspace of D_θ adopts dimension N_r , the associated eigenvectors $\vec{1}_{r_0}$ – constant in θ at $r = r_0$ and zero elsewhere – are orthogonal to each other. Therefore, we can study each eigenvector separately, similar to the 1D case. However, the 2D singular eigenvalues get an additional update from the smoothing prior because the eigenvectors $\vec{1}_{r_0}$ do not belong to the nullspace of the discretized Laplacian operator, i.e., $L \vec{1}_{r_0} \neq 0$. Assuming second-order centered finite differences, the regularized 2D eigenvalues become:

$$\tilde{\lambda}_{\vec{1}_{r_0}}^{2D} = \vec{1}_{r_0}^\top D_\theta^\top D_\theta \vec{1}_{r_0} + \alpha_{BC} \cdot \vec{1}_{r_0}^\top B_\theta^\top B_\theta \vec{1}_{r_0} + \alpha_{SV} \cdot \vec{1}_{r_0}^\top L^\top L \vec{1}_{r_0} \approx \alpha_{BC} \sum_\theta n_{\theta(r_0, \theta)}^2 + \alpha_{SV} \frac{6N_\theta}{\Delta r^4}, \quad (2.33)$$

where α_{SV} is the smoothing penalty hyperparameter, Δr is the polar mesh spacing in the radial direction and we have made the approximation $r_0 \gg \Delta r$.

To complete the 2D analysis, we consider the intersection of the nullspaces of the mass conservation and smoothing operators, where regularization can only come from boundary conditions. This space has dimension one and is spanned by a $N_r \times N_\theta$ vector full of ones, representing a constant function in both the r and θ directions, $\vec{1}_{r,\theta}$. For this vector, it is straightforward to obtain that $\tilde{\lambda}_{\vec{1},r,\theta}^{2D} = \alpha_{BC} \sum_{r,\theta} n_{\theta(r,\theta)}^2$, guaranteed to be positive as long as there is one pixel in the entire domain with $n_\theta \neq 0$. Therefore, this degeneracy should not be concerning in VFM.

Finally, it is worth studying how the singular eigenvalues of mass conservation would be modified when no-slip instead of free-slip boundary conditions were enforced. In that case, n_θ should be replaced by an indicator function $\delta_{LV}(r, \theta)$ that labels endocardial pixels where flow velocity matches endocardial velocity. The modified eigenvalues would then be

$$\tilde{\lambda}_{\vec{1},r_0}^{2D,\text{no-slip}} \approx \alpha_{BC} \sum_{\theta} \delta_{LV}(r_0, \theta)^2 + \alpha_{SV} \frac{6N_\theta}{\Delta r^4} = \alpha_{BC} N_{LV,\theta} + \alpha_{SV} \frac{6N_\theta}{\Delta r^4}, \quad (2.34)$$

where $N_{LV,\theta}$ is the number of pixels segmented as endocardial at $r = r_0$.

In summary, interpreting boundary conditions and smoothing priors as rank updates to the mass conservation constraint provides insight into how these terms regularize the inference of v_θ in VFM. When smoothing is omitted (i.e., $\alpha_{SV} = 0$), the degeneracy of the mass conservation operator can still be resolved by boundary conditions, provided that the segmented LV mask satisfies certain geometric criteria. Under free-slip conditions, regularity is ensured when the LV mask boundaries are not aligned with the θ -direction at each constant- r arc. In this case, the azimuthal normal component n_θ is nonzero for at least one pixel along the arc, preventing the term multiplying α_{BC} in equation 2.33 from vanishing. For no-slip conditions, regularity requires that there are no holes in the LV mask at any constant- r arc, so that the number of segmented pixels $N_{LV,\theta}$ is nonzero and the term multiplying α_{BC} in equation 2.34 remains active. When these geometric criteria are not met, v_θ becomes undetermined up to an additive constant at all θ -positions along the affected arc Figure

2.13.

This insight enables a priori assessment of segmentation quality for VFM: singular arcs can be identified and flagged for refinement, or α_{SV} can be locally increased to interpolate from adjacent pixels in the r -direction. Alternatively, a thin line with n_θ or δ_{LV} equal to ε can be added outside the LV mask spanning $r = r_{\text{base}}$ to $r = r_{\text{apex}}$, enabling Tikhonov-like regularization of v_θ with strength $\varepsilon\alpha_{BC}$ via boundary conditions. This approach would stabilize inference without significantly altering the reconstructed flow within the LV mask.

Our analysis suggests that free-slip and no-slip regularize VFM in a similar manner. Free-slip is more consistent with mass conservation physically. However, no-slip could result in more stable inference because the indicator function $\delta_{LV} = 1$ at endocardial pixels whereas n_θ could reach near-zero values depending on the orientation of the LV mask edge. We also note that inference stability is independent of the particular values of velocity or flux imposed at the LV wall. For example, similar regularization properties would be obtained imposing zero velocity or the measured LV wall velocity. This is an important detail because measuring LV wall velocity via mask registration can be involved and lead to errors that would translate into smaller α_{BC} in practice, potentially hindering stability.

Finally, our analysis suggests that boundary conditions become non-essential for regularity when the smoothing weight α_{SV} is sufficiently high. In fact, least-squares implementations of VFM, such as iVFM [5], can be run stably using very low weights for the boundary condition penalty. However, it is important to recognize that smoothing is not a physical constraint—it alters not only the singular modes of the mass conservation operator but also a broad range of its eigenvalues.

Moreover, in VFM formulations where smoothing is not used as a penalty on v_θ , the parameter α_{SV} should be interpreted as zero, even if smoothing is applied to v_θ or other fields outside the inference step. As a result, such algorithms remain vulnerable to the intrinsic ill-posedness of the problem and are prone to producing spurious banded structures in v_θ , gaining little benefit from the unphysical smoothing imposed on the data.

These results underscore the need for rigorous analyses of VFM robustness, and data-

adaptive Bayesian estimation of hyperparameters.

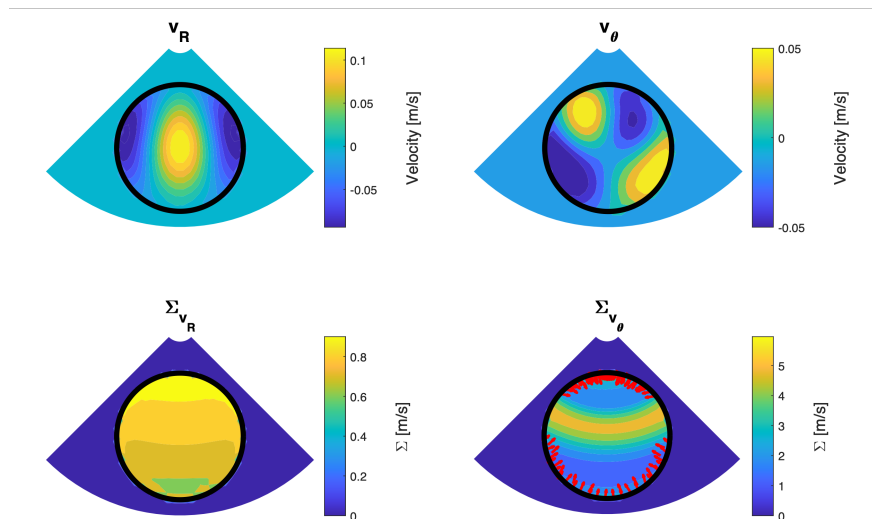


Figure 2.13: Vanilla-VFM challenged with unmet geometric constraints, producing reasonably accurate velocity field reconstructions (v_r, v_θ but exhibiting high uncertainty along the θ arc (Σ_{v_θ} , with N_r, N_θ in red).

2.4 Discussion

Vector flow mapping poses an involved inference problem because the ultrasound cross-beam velocity component, v_θ , lacks measurement data. This situation is analogous to the ill-posedness of optical flow in the direction perpendicular to the local image gradient [81], and contrasts with other scenarios, such as regularization of phase-contrast MRI data, where there is observational support for all flow components.

The ill-posed nature of VFM makes it highly sensitive to implementation details. Significant progress has been made since the earliest formulations, in which v_r and m were fixed to their observed values, V and M , and mass conservation was used to infer v_θ along each radial arc of the Doppler sector. In these early or “naive” implementations, phase unwrapping and spatio-temporal filtering of V were performed prior to inference, while v_θ was regularized

post hoc using spatial filtering. It is now clear that phase unwrapping can be performed *a priori* without affecting performance (see § 2.4.3), but that decoupling smoothing from the inference step degrades the balance between regularity and accuracy. Instead, incorporating smoothing directly into the inference of v_θ leads to more stable and accurate reconstructions (see § 2.3). This evolution in understanding highlights the importance of carefully evaluating VFM’s stability and error propagation—particularly as the method is increasingly applied to complex secondary analyses and to echocardiographic views beyond the relatively well-characterized apical long-axis view of the left ventricle. Motivated by this need, the present work provides a brief theoretical analysis of VFM stability and introduces a fully Bayesian implementation of the method. This new formulation enables uncertainty quantification and offers a principled foundation for the future development of VFM algorithms.

The formal formulation of VFM as a regularization problem constituted a significant advance that enabled subsequent refinements of the method. Furthermore, the primary regularizer in VFM, the divergence operator, is also singular. For instance, in a naive formulation, respectively, the problem admits degenerate solutions in which v_θ takes on arbitrary values along constant-radius arcs, in addition to the values required to balance the radial mass flux. These spurious banded flow structures appear in VFM reconstructions [29, 61] unless additional constraints are imposed.

Banded v_θ artifacts are typically mitigated by incorporating spatial smoothing and boundary condition operators for v_θ . Although these operators are also singular individually, their combined action ensures regularization, since the latent state \vec{x} cannot generally lie within the null spaces of the divergence, smoothing, and boundary-condition operators simultaneously. It is important to recognize that smoothing is not a physical constraint; ideally, VFM would rely as much as possible on mass conservation and boundary conditions alone. But the boundary-condition operator depends on the observed LV mask M , which complicates standardization of the associated hyperparameters and contributes to VFM’s sensitivity to both smoothing and boundary assumptions. Despite this sensitivity, the impact of uncertainty propagation in VFM has not yet been formally investigated.

2.4.1 Hierarchical Bayesian VFM

We adopted a hierarchical Bayesian framework in which the flow velocity \vec{v} , segmentation mask m , and hyperparameters $\vec{\alpha}$ governing the precision of each physical constraint were all treated as random variables. Due to the interdependence between \vec{v} , m , and $\vec{\alpha}$, the joint posterior distribution could not be sampled directly. Instead, we explored it using Markov Chain Monte Carlo (MCMC), where each variable was updated by sampling from its conditional distribution given the others [57], [13].

MCMC enables efficient exploration of the high-dimensional, coupled posterior while accounting for uncertainty in both the physical model and the data. However, the conditional posteriors in B-VFM are also complex, as the segmentation mask m is itself a latent variable, introducing nonlinearity in the model and making priors pseudo-Gaussian via dependence of the precision matrix Σ^{-1} on the latent state $\vec{x} = (\vec{v}, m)$. This complexity is particularly pronounced in the boundary condition prior $\nabla m \cdot \vec{v} = 0$, which strongly couples the mask and the flow. To address this complexity, we employ Laplace approximations to the conditional posteriors, representing them as multivariate Gaussians with precision matrix $\Sigma^{-1}(\vec{x}^*)$, where \vec{x}^* is the MAP estimate obtained by minimizing the conditional log-posterior. We explored two strategies for hyperparameter sampling: Gibbs sampling with gamma hyperpriors and Metropolis–Hastings (MH) with log-normal proposals. Below, we discuss the trade-offs between these two methods in terms of Bayesian rigor and computational efficiency.

Gamma hyperpriors are attractive due to their conjugacy, which allows for efficient sampling from the conditional hyperposteriors, themselves gamma distributed. The Gibbs sampler exhibited rapid convergence and low hyperparameter variance across MCMC chains, yielding stable results, driven by the shrinkage properties of gamma distributions in high-dimensional latent spaces \vec{x} . The shrinkage experienced by conditional hyperpriors could be exploited to formulate a semi-Bayesian version of VFM where hyperparameter values are deterministically determined based on input data, and a rough estimate of error propagation can be directly computed from the diagonal of $\Sigma^{-1}(\vec{x}^*)$.

It should be noted, however, that while fast convergence is computationally appealing, it does not guarantee reliable posterior inference—particularly for uncertainty quantification. A key limitation is that conjugacy holds only approximately for pseudo-Gaussian priors, after fixing the precision matrix to its MAP estimate during each Gibbs update. This approximation may restrict the sampler’s ability to explore the joint posterior adequately. Consistent with this concern, the Gibbs sampler failed to correct clear inconsistencies in the input data, such as segmentation artifacts that induced flow across the endocardial boundary (Figures: 2.7, 2.8). Bootstrapping from iVFM [5] or similar velocity estimation methods may further bias the Gibbs updates, potentially accelerating convergence at the expense of Bayesian rigor. When working with high-quality images with low noise (Σ_M, Σ_V), complete wall segmentation, and minimal Doppler voids — the v_r and m blocks become approximately constant and diagonally dominant, making conditional conjugacy a more accurate approximation.

Metropolis–Hastings (MH) hyperparameter sampling does not rely on conjugacy and is better suited to explore the full pseudo-Gaussian structure of the posterior, albeit at the cost of slower convergence due to proposal rejections. Notably, MH can leverage both the Doppler velocity field and the physical priors to correct segmentation errors—an ability that traditional VFM methods lack. Moreover, in our implementation, the acceptance probability of MH proposals is computed using the residual of the MAP solver, thereby embedding solver convergence directly into the sampling criterion. This design ensures that poorly converged MAP estimates do not bias the posterior approximation and reflects a Bayesian internal consistency when approximating the conditional posteriors via Laplace’s method.

Traditional VFM formulations [5, 6, 23, 25, 36, 49, 72, 80] treat the LV mask as fixed, effectively enforcing $m = M$ as a deterministic constraint. If a similar assumption were applied in our model, the precision matrix of the conditional posterior would adopt a significantly simpler, constant structure, independent of \vec{x} . Consequently, the conditional posterior for \vec{x} would be Gaussian, and gamma hyperpriors on $\vec{\alpha}$ would yield exact conjugacy. This simplification would allow for direct sampling from the conditional posterior without requiring MAP estimation, thereby eliminating a significant source of computational cost in the B-

VFM algorithm. Moreover, we speculate that the added complexity of MH sampling would be less advantageous in this setting, as the conditional distributions would be tractable and well-behaved. In contrast to the full model, where the nonlinearity introduced by the latent mask m can lead to poorly explored posterior regions or low acceptance rates, a fixed-mask formulation would allow for more efficient and fully Gibbs-based inference.

However, enforcing $m = M$ sacrifices the ability to use velocity data to refine or correct segmentation errors. This trade-off is particularly relevant in certain pathologies where endocardial border delineation is challenging or unreliable. For example, in dilated cardiomyopathy, the myocardium is thinned offering reduced contrast with the blood pool, and endocardial borders become poorly defined the apical view used for VFM [53]. In LV non-compaction, prominent trabeculations and deep intertrabecular recesses obscure endocardial borders complicate precise segmentation and color-Doppler is intuitively used to improve diagnostic accuracy [70]. In apical hypertrophic cardiomyopathy, the apex may nearly or completely obliterate during systole, and segmentation is further complicated by asymmetric wall thickening, aberrant trabeculation, and displaced papillary muscles [45].

2.4.2 Extensions of B-VFM

The hierarchical Bayesian framework presented in this study is flexible and its computational cost is relatively insensitive to the number of hyperparameters, making it suitable for customization to include additional priors and or observation modalities. For instance, data from echo-partice image velocimetry or multiple color-Doppler acquisitions could be incorporated as an additional observation. We will explore more of extensions of VFM in Chapter 4: Fusion-VFM.

Supporting information

2.4.3 Phase Unwrapping

The color Doppler maps $V(r, \theta)$ used as input to B-VFM may exhibit aliasing artifacts due to phase wrapping of the Doppler phase when the true radial velocity exceeds the encoding limit V_{enc} . These artifacts create ambiguity in the observed phase, $\Phi = \pi V/V_{\text{enc}}$. In pixels affected by aliasing, the true (φ) and observed phase differ by an integer multiple of 2π :

$$\frac{\varphi - \Phi}{2\pi} = \frac{v_r - V}{2V_{\text{enc}}} = n = \text{int}(\phi), \quad (2.35)$$

where $n \in \mathbb{Z}$ is the aliasing index, and $\phi \in \mathbb{R}$ is its real-valued extension.

To account for this ambiguity within the Bayesian framework, we introduce ϕ as an additional latent variable in the forward model $\vec{X} = A\vec{x} + \vec{\varepsilon}$, modifying the observation equation to:

$$\begin{bmatrix} V \\ M \end{bmatrix} = \begin{bmatrix} I & 0 & 0 & -2V_{\text{enc}} \\ 0 & 0 & I & 0 \end{bmatrix} \cdot \begin{bmatrix} v_r \\ v_\theta \\ m \\ \text{int}(\phi) \end{bmatrix} + \begin{bmatrix} \varepsilon_V \\ \varepsilon_M \end{bmatrix}. \quad (2.36)$$

When $n = \phi = 0$, there is no phase wrapping, and the model reduces to the scenario described in the Main Text.

To correct phase-wrapping artifacts, we incorporate prior knowledge on the spatial structure of ϕ following the approach of Loecher et al. [43]:

$$\varepsilon_\phi = \phi - L^{-1}f \sim \mathcal{N}(0, \Sigma_\phi), \quad (2.37)$$

where L is the Laplace operator, and

$$f = \cos \Phi \circ (L \cdot \sin \Phi) - \sin \Phi \circ (L \cdot \cos \Phi)$$

is a nonlinear function of Φ , itself a function of V . The prior covariance $\Sigma_\phi = \alpha_\phi^{-1}I$ is diagonal, with precision hyperparameter α_ϕ . The corresponding probability density function

is:

$$p(\phi | L, V) \sim \exp \left[-\frac{1}{2} (\phi - L^{-1}f)^T \Sigma_\phi^{-1} (\phi - L^{-1}f) \right], \quad (2.38)$$

introducing a multiplicative term in the conditional posterior (eq. 2.10 in the Main Text).

In addition to this prior, the likelihood for V is modified to:

$$p(V | \vec{x}, \vec{\varepsilon}) \sim \exp \left[-\frac{1}{2} (v_r - V - 2V_{\text{enc}} \text{int}(\phi))^T \Sigma_V^{-1} (v_r - V - 2V_{\text{enc}} \text{int}(\phi)) \right]. \quad (2.39)$$

Differentiating the negative log conditional posterior with respect to ϕ yields the following MAP condition:

$$\alpha_\phi (\phi - L^{-1}f) - \frac{2V_{\text{enc}}^2}{\varepsilon_V} \cdot \frac{v_r - V - 2V_{\text{enc}} \text{int}(\phi)}{V_{\text{enc}}} \sum_{i \in \mathbb{Z}} \delta(\phi - i) = 0, \quad (2.40)$$

where the final term is a Dirac comb used to approximate the distributional derivative of $\text{int}(\phi)$ with respect to ϕ . This derivative is zero almost everywhere (for $\phi \notin \mathbb{Z}$), effectively decoupling the MAP estimate of ϕ from other latent variables and yielding $\phi = L^{-1}f$. When $n = \phi \in \mathbb{Z}$, equation 2.40 reduces to the aliasing identity $2V_{\text{enc}}n = v_r - V$. These results hold independently of α_ϕ and shows that phase unwrapping can be performed as a standalone preprocessing step prior to executing the B-VFM algorithm while retaining Bayesian rigor.

Chapter 3

THEORETICAL ANALYSIS

Intracardiac blood flow analysis is an emerging field that allows for assessing flow dynamics, vortex formation, cardiac dysfunction, energetic efficiency, and other quantitative biomarkers of cardiovascular health [34]. These markers are increasingly recognized for their potential to detect early signs of cardiac dysfunction adding to already existing metrics such as ejection fraction and wall motion.

Color Doppler echocardiography is a widely used, noninvasive imaging modality for visualizing blood flow in the heart. Its portability and real-time capabilities make it a valuable tool in clinical cardiology. Contrast-enhanced echocardiography, while minimally invasive due to the requirement of an intravenous contrast agent, also offers powerful insights into intracardiac flow dynamics. In particular, contrast ultrasound has been shown to outperform conventional Doppler in certain diagnostic tasks. For example, Rossi et al. demonstrated that contrast-enhanced sonography is significantly more sensitive and accurate than color Doppler for detecting and characterizing thrombosis [63]. Contrast ultrasound can be especially useful in scenarios where Doppler imaging is limited, such as in regions with low-velocity flow, deep anatomical structures, or when aliasing artifacts hinder interpretation. In such cases, contrast echocardiography provides an effective alternative for capturing detailed flow information.

Despite these methods' strengths, ultrasound-based imaging is inherently prone to multiple sources of uncertainty. These include probe motion such as with hand-held transducers, variability in acquisition angle, depth of region of interest, limited spatial resolution, and aliasing [27]. Additionally, data voids and artifacts often arise during acquisition, particularly when patients have a challenging acoustic window or cardiovascular conditions. When

these noisy and error prone images are used as inputs to computational models, such as reconstructing velocity fields and pressure gradients, uncertainty is further propagated due to additional physics model assumptions, interpolation methods, and discretization steps. Despite the increasing reliance on image-based flow analysis, most existing methods and pipelines do not rigorously quantify or propagate these uncertainties. The absence of this analysis undermines the reliability of model-derived clinical metrics and makes it challenging to translate intracardiac blood flow analysis to guide medical decisions.

In this chapter, we use an echocardiographic simulator, MATLAB Ultrasound Toolbox and SIMUS [24], to perform a theoretical analysis of how uncertainties in ultrasound image acquisition arise and propagate through flow analysis pipelines. While there is no universally accepted framework for quantifying uncertainty in medical imaging, several methods have been developed to address different aspects of the problem. In a survey by Gillman et al., uncertainty in medical imaging is discussed in depth [27]. Berge et al. visualized confidence maps using the equilibrium solution of a random walk problem [8]. Gueziri et al. measured trustworthiness by assessing the positional uncertainty of the transducer, incorporating target registration information into their uncertainty estimation [28]. Building on these prior efforts, our approach focuses specifically on acquisition-related variability and its downstream impact on computational flow modeling pipelines. By simulating a range of imaging conditions and introducing controlled noise and perturbations, we investigate how these uncertainties influence image quality and flow estimation. This sensitivity analysis allows us to evaluate the robustness of flow based biomarkers and assess the broader implications for both model reliability and clinical applicability.

3.1 Understanding Sources of Uncertainty in Medical Imaging

Understanding the sources of uncertainty in medical imaging is essential for both accurate interpretation and reliable downstream modeling. In their comprehensive survey, Gillmann et al. categorize uncertainty according to the major stages of medical imaging: acquisition, transformation, and visualization, and identify multiple distinct sources such as positional

uncertainty due to hand-held probe placement variability, pixel/voxel value uncertainty, incompleteness of data, model inaccuracy or incompleteness, and boundary/parameter uncertainty, among others [27, 82]. Positional uncertainty arises especially in modalities like ultrasound where transducer posture or orientation may vary between acquisitions; voxel intensity uncertainty stems from noise, reconstruction errors, or signal quantization; and incomplete data may result from limited field of view, occlusions, or sampling gaps. Additionally, uncertainties from physics-based modeling, such as model physics parameters, interpolation methods, and discretization, can propagate through pipelines when using images for computational analysis.

Beyond visualization-specific frameworks, the broader field of medical image analysis recognizes two root causes: aleatoric uncertainty, which captures inherent variability or noise in the data (e.g. patient motion, imaging noise), and epistemic uncertainty, arising from the underlying model. Recent reviews show that explicitly quantifying both forms using Bayesian models, ensemble methods, or dropout techniques not only improves model reliability, but also offers interpretable confidence metrics for clinical decision support [17]. This further supports our choice in a Bayesian approach to Vector Flow Mapping as presented in Chapter 2. For this section, we focus on aleatoric uncertainty, examining how a theoretical analysis of acquisition-related noise and image variability can be incorporated into the Bayesian framework to improve flow estimation accuracy and reliability.

3.2 *MUST: MATLAB Ultrasound Toolbox*

3.2.1 Varying Ultrasound Parameters

To study uncertainties in echocardiographic Doppler images, we use MATLAB Ultrasound Toolbox, MUST, to simulate and control the ultrasound imaging environment [68]. Its flexible features closely replicates real-world acquisitions, enabling us to systematically introduce and isolate sources of uncertainty. We generate a diverse set of ultrasound scenarios by varying key ultrasound parameters such as scatterer brightness, scatterer location, attenuation

coefficient, transducer type and angle, and viewing depth. Ultimately, this synthetic pipeline allows us to evaluate how these variables influence reconstructed flow fields downstream. The PFIELD function is a key feature in MUST that models the acoustic pressure field generated by a uniform linear or convex ultrasound transducer array. It computes the radio-frequency (RF) signals received by each transducer element based on wave propagation, scattering, and transducer parameters. The acoustic pressure $P_{Tx}(\omega)$ measured at the q th element is given by:

$$P_{Tx}(\omega) = \sum_{m=1}^{\#of\ particles} \left[BSC_m \left(\sum_{n=1}^N \left[D_b(\theta_{nm}, k) \frac{e^{ikr_{nm}}}{\sqrt{r_{nm}/b}} \right] \right) D_b(\theta_{qm}, k) \frac{e^{ikr_{qm}}}{\sqrt{r_{qm}/b}} \right] e^{-i\omega t}$$

The term BSC_m represents the backscattering coefficient of the m th scatterer. It accounts for how much of the ultrasound energy is reflected back from the scatterer to the transducer. Since each scatter in the tissue has a different backscattering coefficient, we model it as a random variable drawn from a Gaussian distribution, $BSC_m \sim \sigma_m \mathcal{N}(0, 1)$. Attenuation is also represented in $P_{Tx}(\omega)$ through $\frac{1}{\sqrt{r}}$ and e^{ikr} . This describes how the propagated wave weakens or attenuates as it travels through different tissues. The farther the wave travels in r , the weaker the signal becomes due to energy spreading and absorption. Within these parameters, we also vary the distance between the target of interest to the transducer origin. Overall, in addition to the backscattering coefficient, we also adjust the attenuation and distance r traveled in our simulations to analyze its effects on acquired images.

3.2.2 Choice in Synthetic Flow Patterns

Left ventricular (LV) diastolic vortex formation is widely recognized as an important mechanism for efficient blood pumping. During diastole, the formation of a LV vortex plays a critical role in cardiac efficiency by minimizing energy loss and smoothly redirecting inflowing blood toward the outflow tract, thereby optimizing ventricular filling and ejection. In a healthy LV, this manifests as a stable, symmetric vortex ring that forms during early diastolic filling and persists into diastasis, minimizing energy loss and shear stresses. [2]. Abnormalities in this vortex pattern have been linked to LV-related disorders such as diastolic

dysfunction and heart failure [21]. Motivated by both the physiological importance and the relatively simple geometric structure of the diastolic vortex, we evaluate vanilla VFM and B-VFM using two synthetic flow models that approximate this pattern: the Lamb–Chaplygin (LC) dipole and Hill’s vortex (Figure 3.1). The Lamb–Chaplygin (LC) dipole is used as a verification benchmark because it generates an exact, steady 2D incompressible flow that can be fully described by an analytical streamfunction. This streamfunction takes the form:

$$\psi(r, \theta) = \begin{cases} -\frac{2U J_1(kr)}{k J_0(kR)} \sin(\theta), & r < R \\ U \left(\frac{R^2}{r} - r \right) \sin(\theta), & r \geq R \end{cases} \quad (3.1)$$

where U is the dipole’s translational velocity, R is the radius of the vortex core, J_0 and J_1 are the zeroth and first-order Bessel functions of the first kind, and k satisfies the boundary condition $kR = 3.8317\dots$, which is the first non-trivial zero of J_1 . Because the Lamb–Chaplygin Dipole velocity field is entirely in-plane and divergence-free by construction, B-VFM, as well as any other VFM method, should reconstruct the velocity field with perfect accuracy given noise-free inputs. This makes the LC dipole an ideal ”sanity check” case, ensuring that the algorithm can match an exact analytic solution in a controlled 2D setting. In contrast, the Hill’s spherical vortex is an axisymmetric 3D flow characterized by a compact vortex structure embedded in a surrounding irrotational field. Its velocity field can also be described by a streamfunction in spherical coordinates:

$$\psi(r, \theta) = \begin{cases} -\frac{3U}{4} \left(1 - \frac{r^2}{a^2} \right) r^2 \sin^2 \theta, & r \leq a, \\ U \left(1 - \frac{a^3}{r^3} \right) r^2 \sin^2 \theta, & r > a, \end{cases} \quad (3.2)$$

where U is a constant velocity far away from the origin and a is the radius of the spherical vortex core. For $r > a$, the vorticity is zero and the solution described above is simply the potential flow past a sphere of radius a . Unlike the LC Dipole, the Hill’s vortex contains out of plane velocities, ie 3D flow, introducing more realistic challenges for vector flow mapping algorithms. For classic VFM and B-VFM, the Hill’s vortex case tests the methods’ ability

to recover the velocity field in the presence of unmeasured out-of-plane fluxes, making it a more challenging and practically relevant benchmark.

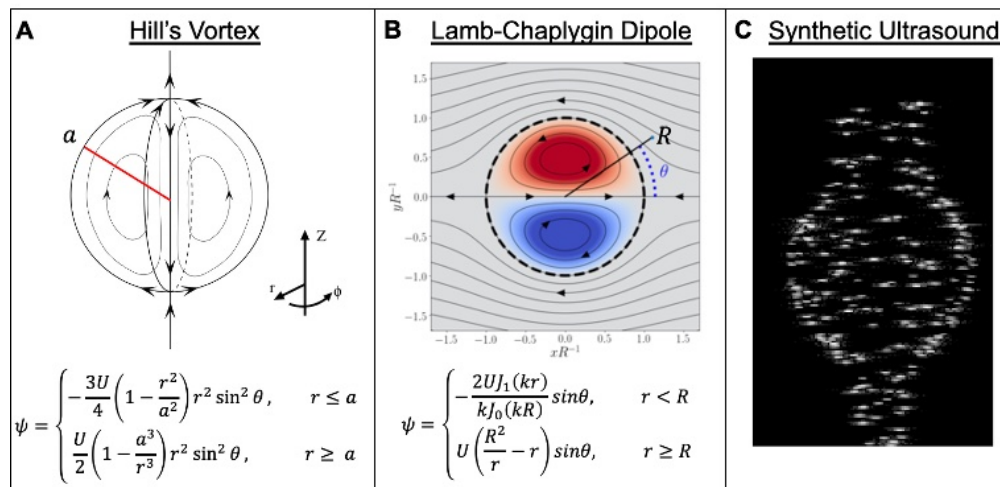


Figure 3.1: A) Streamfunction of Hill's Vortex B) Streamfunction of Lamb-Chaplygin Dipole C) Generated synthetic ultrasound contrast echo using MUST

Particle coordinates were defined to follow the synthetic flow fields. The MATLAB Ultrasound Toolbox (MUST) simulated ultrasound wave propagation and scatterer interactions at these particle locations using standard transthoracic echocardiography parameters: a linear-array probe transmitting pulses with a center frequency of 2.8 MHz and an 86-beam scanning sector with 1° spacing. From the resulting simulated radiofrequency data, MUST generated synthetic B-mode and color Doppler images. Variance fields were then computed by systematically varying the parameters described in Section 3.2.1 (Figure 3.2). MUST offers a versatile simulation environment that allows systematic variation of numerous parameters. Our selections of transducer type, ultrasound environment settings, and other factors represent just a few examples within this synthetic framework. Different parameters can and should be varied to suit specific scenarios and research goals.

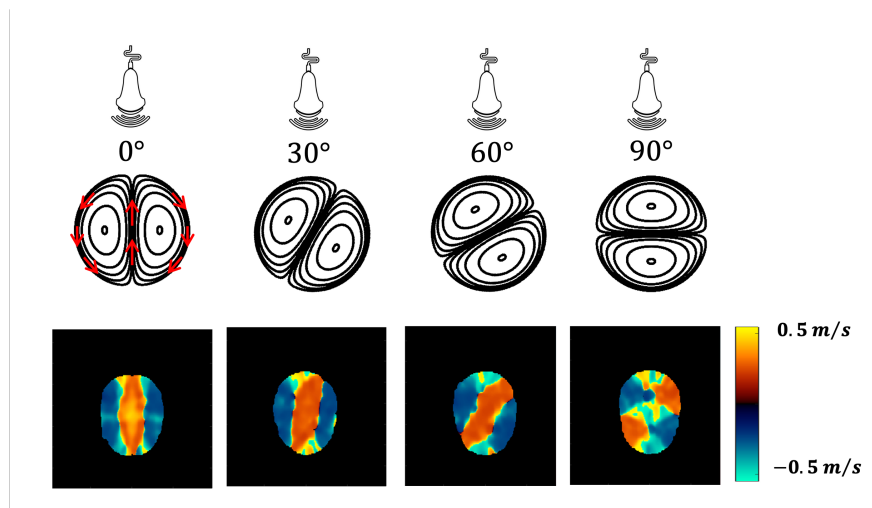


Figure 3.2: Particle coordinates were defined within the MATLAB Ultrasound Toolbox to simulate flow patterns of synthetic velocity fields at different transducer angles. Simulated ultrasound wave propagation and scatterer echoes generate realistic radiofrequency data, which is then processed to create synthetic the corresponding Doppler images.

3.3 Results

We introduced variance levels of 10%, 30%, 50%, 70%, and 90% relative to the mean scatterer brightness value of 1 across a range of transducer angles: 0°, 22.5°, 45°, 67.5°, and 90° (Figure 3.3). As expected, increasing variability in scatterer brightness and position results in a proportional increase in the variance observed in the synthetic Doppler images. This relationship highlights how intrinsic heterogeneity in tissue scattering properties directly impacts the stability and noise characteristics of Doppler signals. Furthermore, we observed that the variance in Doppler images systematically increases with the transducer angle from 0° to 67.5°, indicating that beam orientation influences signal consistency. The variance plateaus at 90°, which may be attributed to the physical limitation that at orthogonal insonification, the Doppler shift is minimized, thus reducing sensitivity to scatterer brightness changes, effectively capping the variance increase.

To further elucidate the effects of tissue properties on image variance, we investigated

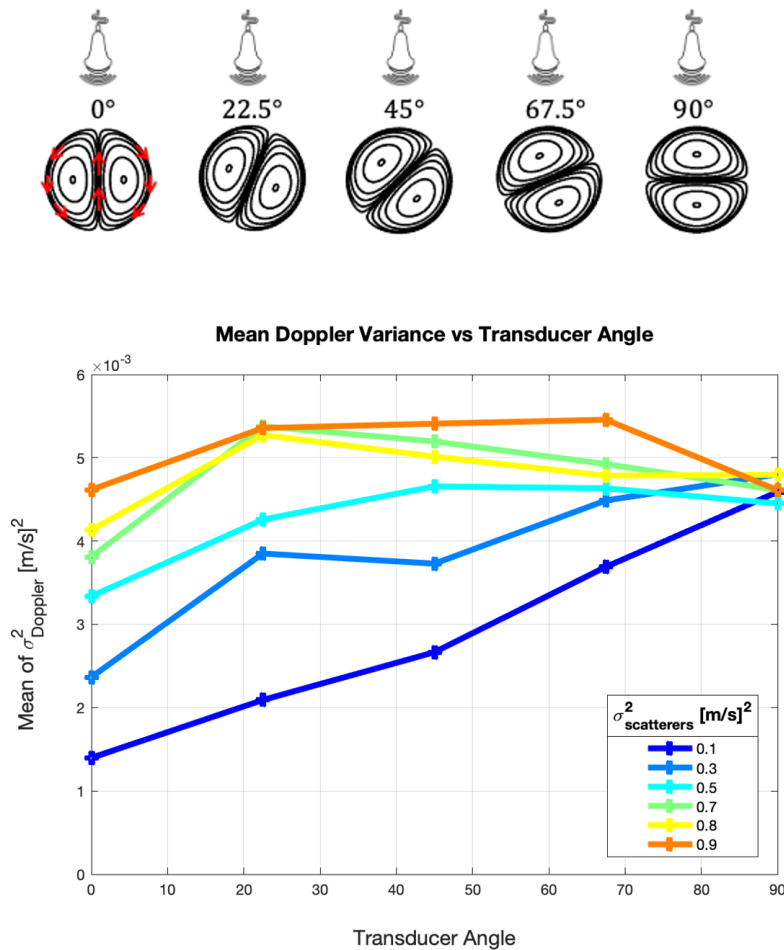


Figure 3.3: Measurement of mean variance in synthetic doppler images when angle between the transducer and the flow field is varied.

acoustic attenuation by simulating two attenuation coefficients, $\alpha = 0$ and $\alpha = 0.5$ dB/cm/MHz, the latter representing typical soft tissue attenuation. These simulations spanned depths ranging from 0.01 m to 0.1 m, all performed with a fixed transducer angle of 0° (Figure 3.4). We found that increasing scatterer brightness variability consistently elevated Doppler variance across all depths. Moreover, as the imaging depth increased, the Doppler variance grew significantly, likely due to cumulative attenuation effects reducing signal-to-noise ratio at greater distances from the transducer. Quantitatively, Doppler variance differed by ap-

proximately 2% to 40% when comparing attenuation coefficients of 0 and 0.5 dB/cm/MHz, underscoring attenuation’s nonlinear impact on signal quality over depth.

Interestingly, our results indicate that variability in the backscatter coefficient exerts a more substantial influence on Doppler image variance than attenuation within the tested parameter range. Physically, this is intuitive because backscatter heterogeneity directly modulates the amplitude and phase of the returned echoes, producing immediate fluctuations in Doppler signal strength. In contrast, attenuation acts as a gradual signal decay mechanism, which predominantly affects deeper targets but may be partially compensated by system gain or post-processing.

By leveraging the flexible parameter control offered by the synthetic ultrasound toolbox, we have established a framework to introduce realistic, clinically relevant noise profiles into simulated Doppler images. This approach enables systematic exploration of how acquisition related variabilities propagate through vector flow mapping reconstruction algorithms, providing crucial insight into the sensitivity and robustness of derived hemodynamic metrics.

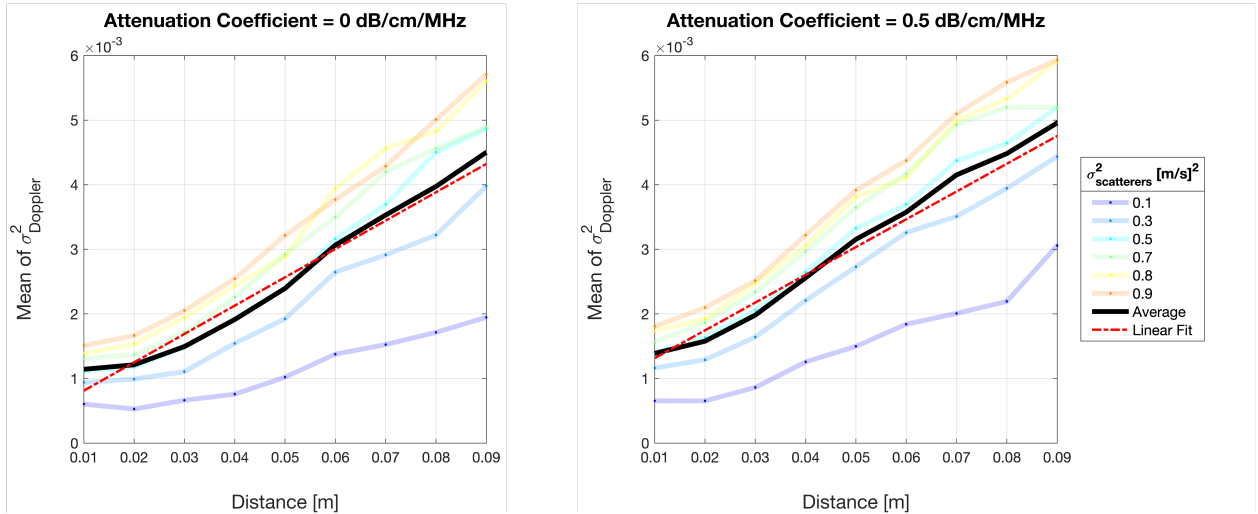


Figure 3.4: Measurement of mean variance in synthetic doppler images r distance between transducer and flow field is varied for $\alpha = 0$ and $\alpha = 0.5$ dB/cm/MHz.

3.4 Discussion

Our findings underscore the significant impact of aleatoric uncertainty, manifested in this theoretical analysis as variability in scatterer brightness, transducer angle, and attenuation, on the quality and reliability of Doppler ultrasound images. These acquisition related noise sources propagate through the image formation process and ultimately affect downstream flow estimation, as evidenced by the systematic increase in Doppler variance under controlled perturbations. This behavior is consistent with the broader medical imaging literature, which differentiates aleatoric uncertainty as inherent data noise that cannot be mitigated simply by acquiring more data, unlike epistemic uncertainty stemming from limitations in the underlying models [17]. By explicitly quantifying how acquisition variability translates into signal variance, our theoretical analysis sheds light on the stochastic nature of ultrasound imaging and the challenges it poses for accurate flow measurement.

Building on this foundation, our study motivates the use of a Bayesian framework for Vector Flow Mapping, as presented in Chapter 2. Unlike the conventional Vanilla-VFM method which assumes a simplistic identity covariance matrix Σ_V and neglects Doppler measurement errors, our Bayesian approach replaces Σ_V with an informed covariance matrix that reflects realistic, spatially varying error characteristics. For instance, our theoretical analysis revealed that Doppler error tends to increase with imaging depth, allowing Σ_V to be adaptively adjusted to account for this spatially varying uncertainty. Additionally, alternative sources of uncertainty, such as cycle-to-cycle variability, can also be incorporated within this flexible Bayesian framework. By embedding these uncertainty estimates directly into the reconstruction process, the model gains enhanced robustness and produces confidence metrics that are clinically interpretable.

It is important to emphasize that no single approach exists for quantifying uncertainty, but regardless of the method, acknowledging and integrating multiple uncertainty sources is essential for deriving reliable, clinically meaningful metrics. Incorporating rigorous uncertainty quantification helps to distinguish genuine physiological flow patterns from artifacts

introduced by imaging noise, ultimately improving the trustworthiness of the flow estimates. Furthermore, while this chapter focuses on aleatoric uncertainty related to image acquisition, future work should also integrate epistemic uncertainty arising from model assumptions to further bolster transparency and clinical confidence.

By bridging detailed synthetic simulation studies with practical clinical challenges, the methods and insights developed here provide a critical step toward more comprehensive uncertainty-aware intracardiac flow analysis. In summary, our work demonstrates the value of explicitly modeling acquisition-level variability within probabilistic flow estimation pipelines, thereby closing the gap between theoretical image simulation and real-world echocardiographic diagnostics.

Chapter 4

MULTI-MODALITY FUSION

Mapping left ventricular (LV) flow from ultrasound can be approached through multiple complementary modalities, each offering distinct advantages and limitations. As we have already seen, left ventricle (LV) mapping by ultrasound through echo color-Doppler velocimetry (echo-CDV) has shown to be a strong imaging modality to quantify flow in the LV, but it is still limited by 1D radial velocities and enforces planar flow assumption in vector flow mapping. Echocardiographic Particle Image Velocimetry (Echo-PIV) is another method that has been utilized in the clinical setting to quantify flow patterns in the heart. Echo-PIV applies particle image velocimetry to contrast ultrasound images, producing a full 2D velocity field without making any assumptions about the flow. Its main drawback is the need for well-controlled contrast injections, which may not always be clinically feasible. Vector Flow Mapping (VFM), by contrast, estimates flow from color Doppler images and wall motion tracking without contrast agents. While simple and widely used, VFM is constrained by modeling assumptions and the need for good selection of velocity encoding value to avoid aliasing while capturing low velocities. Fusion Vector Flow Mapping aims to combine such complementary datasets within a unified reconstruction framework, leveraging the strengths of each modality while mitigating their weaknesses. In this work, we explore two fusion strategies:

1. Echo-PIV and Doppler Fusion: Using PIV to capture 2D velocities from contrast echocardiography and Doppler to resolve regions of higher velocity such as during left ventricular filling, creating a more complete and physically consistent flow field.
2. Doppler Multiscale Fusion: Combining color-Doppler acquisitions acquired at different

velocity encodings to increase the dynamic range, allowing simultaneous recovery of both high and low velocity regions.

Both fusion strategies are implemented within a least-squares optimization framework that enforces physical flow constraints, including conservation of mass, adherence to boundary conditions, and spatial-temporal regularity (Figure 4.1). While the formulation is deterministic, it incorporates explicit uncertainty metrics into the reconstruction, enabling propagation of measurement variability through to the final velocity estimates. However, because hyperparameter selection is fixed rather than inferred from probability distributions, the method is not fully Bayesian.

This distinction is important because although the current fusion results are produced with a deterministic least-squares approach, the underlying structure naturally accommodates a Bayesian formulation. Since the fusion framework already integrates multiple datasets into a common forward model with uncertainty metrics, extending it to a Bayesian setting would simply require probabilistic treatment of the hyperparameters and priors. This would enable principled uncertainty quantification across modalities, enhancing robustness, interpretability, and clinical reliability. As such, the work presented here serves both as a demonstration of multimodal fusion capabilities and as a natural stepping stone toward full Bayesian fusion methods.

4.1 Formulation

4.1.1 Echo-PIV and color-Doppler Fusion

Maximum likelihood inference is a common approach in experimental flow mapping when data are sparse or incomplete. Here, we adopt an algorithm based on the intraventricular flow mapping (iVFM) method by KC et al. [5]. The method discretizes and minimizes a likelihood function, yielding a system of linear equations whose solution provides the reconstructed velocity field. The cost function for the optimization contains five weighted terms: Doppler data fidelity, PIV data fidelity, mass conservation, boundary conditions, and spatial

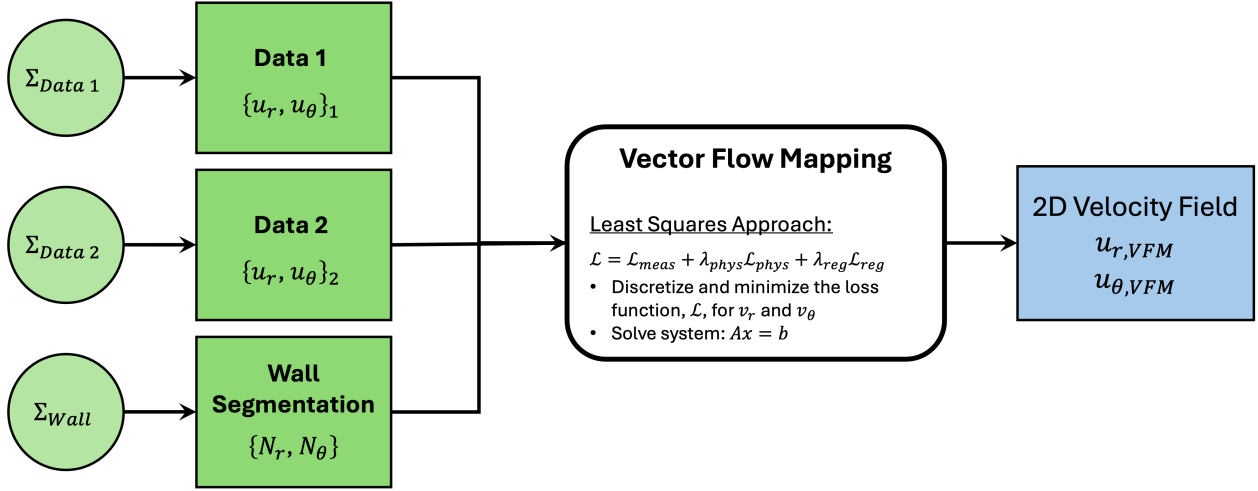


Figure 4.1: Fusion Flowchart: 1) Echo-PIV and Doppler Fusion, Data 1 would contain both u_r and u_θ components provided by PIV's 2D field and Data 2 would contain u_r and $u_\theta = 0$ from color-Doppler acquisition. 2) Doppler Multiscale Fusion has $u_{r,1}, u_{\theta,1} = 0, u_{r,2}, u_{\theta,2} = 0$ as inputs. Both fusion modalities have corresponding error measurements and wall segmentation acquisitions as inputs.

smoothness regularization. The construction of the Doppler data fidelity, priors, and regularization follows KC et al. Our extension incorporates an additional PIV term, which allows for fusion between Echo-PIV and Doppler data while still enforcing the physical constraints.

$$\begin{aligned}
 J(\mathbf{v}) = & \lambda_{DOP} \int_{\Omega_{DOP}} (\mathbf{v}_r - \mathbf{v}_D)^2 + \lambda_1 \int_{\Omega} (r \partial_r \mathbf{v}_r + \mathbf{v}_r + \partial_\theta \mathbf{v}_\theta)^2 \\
 & + \lambda_{BC} \int_{\Omega} (\mathbf{v} \cdot \mathbf{d}_{wall})^2 + \lambda_{SM} \sum_{m \in r, \theta} \int_{\Omega} (r^2 \partial_r^2 \mathbf{v}_m)^2 + 2(r \partial_{r\theta}^2 \mathbf{v}_m)^2 + (\partial_\theta^2 \mathbf{v}_m)^2 \\
 & + \lambda_{PIV} \int_{\Omega} (Q_{PIV}(\mathbf{v} - \mathbf{v}_{r\theta,PIV}))^2
 \end{aligned}$$

Here, Q_{PIV} is a weighting matrix derived from the local signal-to-noise ratio (SNR) of the PIV measurements. This weighting ensures that high-confidence PIV vectors contribute more strongly to the reconstruction, while low-confidence measurements have reduced influence. PIV measurements are obtained by cross-correlating two interrogation windows from frames separated by a time Δt . The SNR is quantified from a bilinear fit to the cross-correlation

peak, which approximates a paraboloid:

$$\text{paraboloid} = \mathbf{b}x^2 + \mathbf{c}y^2 + \mathbf{d}x + \mathbf{e}y + \mathbf{f}$$

A steeper, well-defined paraboloid corresponds to a higher SNR (and thus more trustworthy velocity estimates), whereas a shallow paraboloid indicates greater uncertainty. Figure 4.2 illustrates the geometry of the paraboloid fit. To construct Q_{PIV} , we make two simplifying

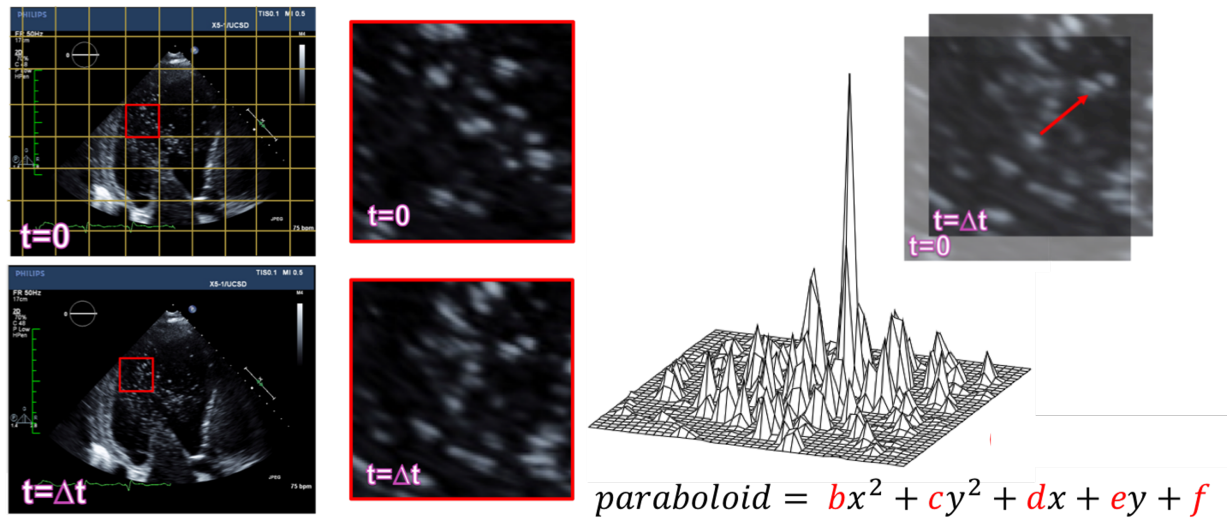


Figure 4.2: PIV compares corresponding interrogation windows between two frames via cross-correlation. The resulting correlation surface is approximated by a paraboloid. The location of its maximum gives the displacement between frames, while the curvature coefficients (b, c) reflect measurement confidence.

assumptions: 1) symmetry and no tilt by neglecting cross terms d, e , 2) the vertical offset is not relevant therefore we neglect f , which shifts the paraboloid along the vertical axis. Under these assumptions, the coefficients b, c fully characterize the paraboloid curvature in the x, y directions. These curvature values directly inform the entries of Q_{PIV} , scaling each PIV velocity vector by its associated measurement confidence.

While classic Particle Image Velocimetry (PIV) offers the advantage of producing a two-dimensional velocity field, thereby providing richer spatial flow information compared to the

one-dimensional velocity component obtained from color Doppler, it faces challenges in accurately capturing regions of low velocity. PIV works by dividing consecutive ultrasound images into smaller interrogation windows and estimating the average particle displacement within each window by calculating the cross-correlation between corresponding regions in sequential frames. Importantly, PIV does not track individual particles but rather measures the collective, averaged motion of groups of particles within these windows. A notable limitation arises when particle displacement between frames is large relative to the window size, i.e. particles can move outside the interrogation region, causing a decrease or loss of cross-correlation and thereby underestimating displacement. This issue becomes particularly important when quantifying left ventricular (LV) flow, which exhibits a wide range of dynamic flow patterns over the cardiac cycle. For example, during early diastolic filling, a high-velocity inflow jet enters through the mitral valve and rapidly decelerates as it propagates toward the apex. PIV often underestimates velocities in such high-speed inflow regions, which can bias downstream flow estimates and reduce the accuracy of derived hemodynamic metrics.

To mitigate this problem, multipass algorithms have been developed [67]. These methods begin with a first pass using relatively large window sizes to capture large scale displacements. The displacement estimates from this initial pass are then used to inform subsequent passes with progressively smaller window sizes, enabling the capture of finer flow details while maintaining accuracy. This iterative approach can be repeated multiple times, gradually refining the velocity field by reducing the window size and improving spatial resolution. Additionally, due to the cyclic nature of cardiac flow, an ensemble averaging step can be incorporated, where PIV results from multiple cardiac cycles are averaged to enhance signal-to-noise ratio and reduce measurement variability. This combination of multiscale, iterative processing and temporal ensemble averaging, significantly improves PIV's robustness and accuracy in complex cardiovascular flow measurements. We must note that even with these improvements to PIV metrics, PIV alone still tends to underestimate regions of very high velocity compared to color Doppler. Capturing regions of high velocities is where color-Doppler generally per-

forms better, making a fusion of the two modalities an effective strategy to recover a more robust and physiologically accurate flow field.

Contrast Ultrasound Acquisition Enrollment. With IRB approval and HIPAA compliance, we prospectively enrolled 40 patients who underwent clinically indicated contrast enhanced 2-D echocardiograms between September 2016 and April 2017. Patients who were 18 years or older were eligible to be enrolled in the study. Those with complex congenital heart disease, pregnant women, known intracardiac shunts, or history of an allergic reaction to Bracco Optison or Lumason were excluded from the study.

Contrast Ultrasound Imaging Acquisition. 2-D echocardiograms was performed on across three different vendor platforms: Siemens Acuson S2000, General Electric Vivid e9, and Philips iE33. In the apical 3-chamber view, the focal depth was set to mid-cavity and the depth reduced to include the mitral valve and aortic valve in order to maximize spatial resolution. Sector width was narrowed in order to maximize temporal resolution ≥ 50 Hz. Harmonic imaging or vendor-specified contrast presets were not used for acquisition. The mechanical index was modulated between 0.5-0.8. The contrast agent, Lumason, was diluted in a 1:2 fashion in sterile normal saline. 0.5 mL of diluted contrast was hand-bolused and then followed by a 2 mL normal saline bolus. Acquisitions were over 5 cardiac cycles. In the apical 3-chamber view, the focal depth, depth, and sector width were duplicated to perform color-Doppler imaging. The color-Doppler field included the entire left ventricular cavity. Acquisitions were acquired over 5 cardiac cycles.

Temporal Alignment of PIV and Doppler data. From ultrasound acquisitions, ECG signals were used to temporally align contrast echocardiography and color-Doppler echocardiography data. For each acquisition, 3 to 4 consecutive cardiac cycles were recorded, averaged, and normalized to reduce noise and improve signal consistency. Temporal alignment was achieved by matching each Doppler frame to the closest corresponding contrast echo frame, leveraging the higher temporal resolution of the contrast echo data.

Depending on the acquisition resolution and the quality of the data, alternative methods such as moving average filters, interpolation techniques, or advanced signal registration

algorithms may be necessary to improve alignment accuracy. These methods can help compensate for discrepancies in frame rates or timing offsets, ensuring that multimodal data are synchronized precisely for subsequent analysis and fusion

4.1.2 Doppler Multiscale Fusion

While the Echo-PIV Doppler fusion described above combines complementary modalities, the same least-squares framework can be extended to fuse multiple Doppler datasets acquired under different imaging configurations. In our Doppler–Doppler multiscale fusion approach, two Doppler velocity fields, $v_{D,1}$ and $v_{D,2}$, are incorporated into the cost function through separate data fidelity terms:

$$\begin{aligned} J(\mathbf{v}) = & \lambda_{D,1} \int_{\Omega_{D,1}} (\mathbf{v}_r - \mathbf{v}_{D,1})^2 + \lambda_{D,2} \int_{\Omega_{D,2}} (\mathbf{v}_r - \mathbf{v}_{D,2})^2 \\ & + \lambda_{CM} \int_{\Omega} (r\partial_r \mathbf{v}_r + \mathbf{v}_r + \partial_\theta \mathbf{v}_\theta)^2 \\ & + \lambda_{BC} \int_{\Omega} (\mathbf{v} \cdot \mathbf{d}_{wall})^2 + \lambda_{SM} \sum_{m \in r, \theta} \int_{\Omega} (r^2 \partial_r^2 \mathbf{v}_m)^2 + 2(r\partial_{r\theta}^2 \mathbf{v}_m)^2 + (\partial_\theta^2 \mathbf{v}_m)^2 \end{aligned}$$

Each Doppler dataset is acquired with different velocity encoding value to target different portions of the flow’s dynamic range. For example, one dataset may prioritize high-velocity jets with minimal aliasing, while the other captures low-velocity recirculating regions with improved sensitivity. By fusing them, the reconstruction leverages the strengths of both acquisitions, producing a velocity field that better resolves the full spectrum of cardiac flow dynamics. As in the Echo-PIV Doppler fusion case, each Doppler term can be weighted by a confidence metric derived from acquisition geometry and estimated measurement uncertainty. While these confidence weights allow uncertainty information to propagate through the reconstruction, the hyperparameter selection (λ_{CM} , λ_{BC} and λ_{SM}) remains deterministic in this least-squares formulation. Therefore, although uncertainty is incorporated at the measurement level, the method is not fully Bayesian, unlike the framework proposed in Chapter 2, which directly models and infers uncertainty in both the measurements and the model parameters.

Color-Doppler Multiscale Acquisition. We reanalyze data from previous animal experiments used to determine the relationships between VFM flow and the risk of stroke after myocardial infarction. Briefly, after median sternotomy, epicardial echocardiographic examinations were performed using a Vivid 7 scanner and phased-array 2-4-MHz transducers (GE Healthcare). Without tilting or displacing the probe, we acquired color-Doppler sequences (8 to 14 beats) in the three-chamber view, while varying the Nyquist limit between 0.33 m/s and 0.66 m/s, followed by B-mode sequences (four to six beats) at a high frame rate.

4.2 Uncertainty Modeling

4.2.1 Uncertainty in PIV Field with Bootstrapping

One approach to quantify uncertainty in PIV measurements is with bootstrapping. Bootstrapping is a classical statistical technique used to estimate uncertainty by resampling data with replacement to create multiple synthetic datasets [20]. In the context of Particle Image Velocimetry (PIV), the algorithm begins with the reference image data, for example at t_o , I_1 and generates a bootstrap sample $I^{(i)}$ by randomly replacing some pixels with values sampled from I_1 . The same steps are repeated for the following image at $t = \Delta t$, I_2 . For each bootstrap samples $I_1^{(i)}$ and $I_2^{(i)}$, the correlation matrix $C^{(i)}$ is computed, capturing particle displacement information. The particle displacement $\mathbf{d}^{(i)}$ is then estimated by locating the peak of $C^{(i)}$. This resampling and displacement estimation is repeated for N bootstrap iterations, resulting in a set of displacement estimates $\{\mathbf{d}^{(1)}, \mathbf{d}^{(2)}, \dots, \mathbf{d}^{(N)}\}$. From this ensemble, statistical quantities such as the mean displacement $\bar{\mathbf{d}}$ and variance $\Sigma_{\mathbf{d}}$ are calculated, providing an empirical measure of uncertainty in the velocity vectors. This process offers a robust, nonparametric approach to quantify measurement variability arising from noise and other errors in PIV imaging, without assuming a specific underlying distribution (Figure 4.3, Algorithm 3).

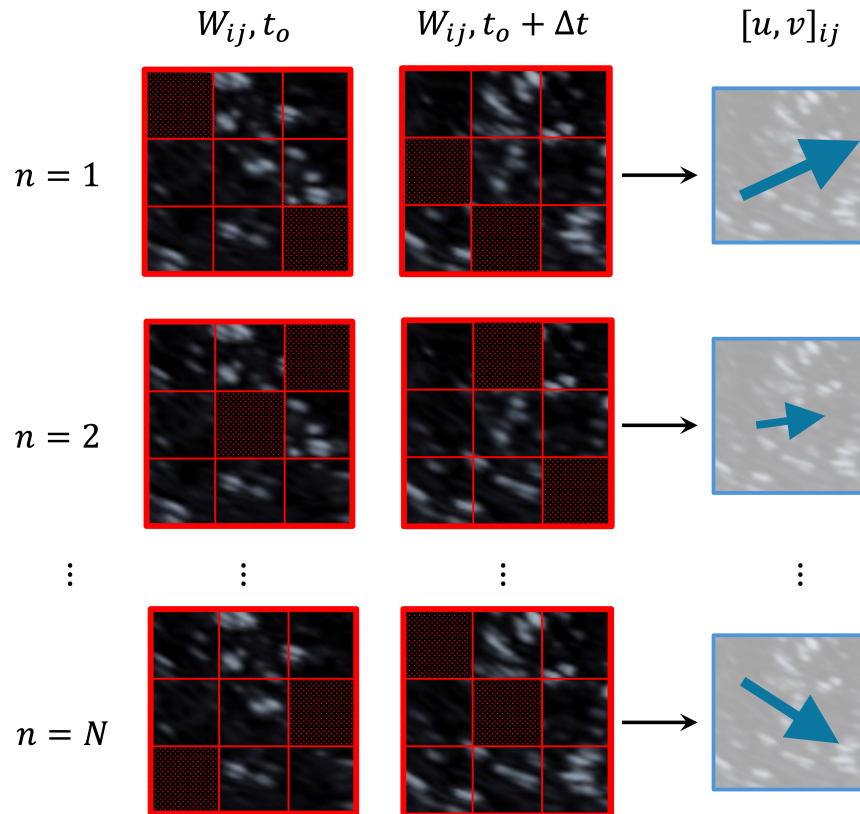


Figure 4.3: Bootstrapping is used to estimate uncertainty in PIV measurements by resampling and replacing random pixels (denoted by red patterned boxes in image) in each input data interrogation window. PIV is rerun to obtain results for each new sample, providing an empirical distribution of the measurements for uncertainty quantification.

4.2.2 Uncertainty in color-Doppler

There is currently no universally accepted method to quantify uncertainty in color Doppler measurements. However, several practical approaches have been proposed to estimate and characterize uncertainty in color-Doppler data:

Cycle-to-cycle variability: One approach we used to assess uncertainty in color Doppler acquisitions is calculating deviations of flow parameters at corresponding points within the cardiac cycle across multiple cardiac cycles. We averaged around 10 cycles acquired at 24 frames per second. This method assumes a stable and regular heart rhythm, attributing

Algorithm 3 PIV Uncertainty Quantification with Bootstrap Method

Input: Original PIV image pair $I_1 = I$ at t_o , $I_2 = I$ at $t_o + \Delta t$, number of bootstrap samples N

Result: Estimated velocity uncertainty for each interrogation window

for $i = 1$ to N **do**

for $W_{ij,1}, W_{ij,2}$ to N **do**

 Generate bootstrap sample: Randomly replace pixels in I_1 with pixels from I_1

 Generate bootstrap sample $I_2^{(i)}$ similarly from I_2

 Compute cross-correlation matrix $C^{(i)}$ between $I_1^{(i)}$ and $I_2^{(i)}$

 Determine displacement vector $\mathbf{d}^{(i)}$ by locating peak in $C^{(i)}$

end for

end for

Compute mean displacement $\bar{\mathbf{d}}$ and uncertainty (e.g., standard deviation) from $\{\mathbf{d}^{(i)}\}_{i=1}^N$

return $\bar{\mathbf{d}}$ and uncertainty estimates

any observed variability in Doppler frames primarily to measurement noise or acquisition artifacts rather than physiological changes. Thus, the cycle-to-cycle variability offers an empirical measure of acquisition uncertainty.

Inter-observer Variability: Another approach involves acquiring color Doppler data and performing segmentations by different sonographers or using different ultrasound probes for the same patient. The differences in measurements caused by operator-dependent factors or equipment variation provide an estimate of uncertainty stemming from user and hardware variability.

Power Doppler: Power Doppler imaging, which maps the total integrated power of the Doppler signal rather than velocity direction or magnitude, can be used as an indirect indicator of signal quality and uncertainty. Areas exhibiting higher power Doppler signal generally correspond to stronger and more reliable flow signals, reflecting higher signal-to-noise ratio (SNR). Conversely, regions with low power Doppler signal often coincide with weak flow

signals or locations outside the scan sector and should be treated with greater caution in interpretation due to increased measurement uncertainty.

Ultrasound Machine Variance Maps: Many ultrasound systems generate variance or spectral width maps by analyzing the Doppler spectrum within each small region. This is accomplished by emitting multiple ultrasound pulses and calculating how the returned signal fluctuates over time. Specifically, the first-lag autocorrelation yields the mean Doppler frequency (velocity estimate), while the second central moment of the Doppler spectrum represents the frequency variance, corresponding to velocity variance. These variance maps are available in clinical practice to better understand complex flow patterns and provide a real-time qualitative measure of flow turbulence and uncertainty [39]. More specifically, these variances maps have shown to be sensitive to flow disturbances, such as those found in stenotic, regurgitant, or shunt lesions, and is often displayed as green or mosaic color overlays in clinical systems. However, variance displays are not purely physiological measures, machine parameters significantly influence them. Thus efforts have been made to investigate variance in Doppler color flow mapping from an instrumentation perspective. Utsunomiya et al. demonstrated that factors such as packet size, pulse repetition frequency, frame rate, system gain, and wall filter settings directly affect measured variance area in controlled flow experiments. For example, larger packet size and higher pulse repetition frequency were inversely correlated with variance area, while higher system gain and wall filter settings increased it. These findings highlight that, to use variance displays quantitatively for uncertainty assessment, machine parameters must be carefully standardized [78].

Together, these approaches offer complementary methods to evaluate uncertainty in color Doppler imaging, informing both clinical interpretation and quantitative flow analyses.

4.3 Hyperparameter Selection and Wall Uncertainty Quantification

In our fusion framework, the regularization terms for boundary conditions, Doppler data, and PIV data are determined using uncertainty metrics, while the spatial smoothness parameter is selected through a hyperparameter search. The PIV regularization weight (λ_{PIV}) is

obtained from the PIV bootstrap method (Figure 4.3), and the Doppler λ_{Dop} is derived from cycle-to-cycle variability. Previous approaches have relied heavily on boundary conditions to resolve the inherent singularity in this optimization problem; without a well-defined wall, the system matrix becomes singular. Our approach builds on this dependence by explicitly incorporating uncertainty in the wall definition itself, rather than assuming perfect knowledge. Specifically, we use the convolutional neural network developed by Dahal et al. [18] to generate a probabilistic atlas that quantifies confidence in the wall location.

Uncertainty in wall segmentations arises from two sources: (1) the ultrasound images themselves, which are subject to the acquisition-related uncertainties as described in Chapter 3, and (2) the segmentation process, which is a typically manual, labor-intensive, and error-prone process (Figure 4.4). Even though sonographers are trained experts, manual segmentations are subjective, with potential variability between operators. Moreover, poor image quality can obscure parts of the wall, forcing the operator to make uncertain judgments about its location. As a result, it is undesirable to assign equal weight to all wall regions and segments that are highly visible should be trusted more than those with ambiguous boundaries.

To address this, the CNN-based segmentation approach uses Monte Carlo dropout as a Bayesian approximation, producing a probabilistic atlas that represents the average of inter-modal variations across dropout iterations. The atlas values range from 0 (no certainty) to 1 (full certainty) for each pixel location of the wall. We trained the network on over 1,000 pairs of synthetic images (developed in MUST) and more than 16,000 pairs of clinical images. Figure 4.5 illustrates a clinical left ventricle contrast image alongside its probabilistic atlas. In this example, the apex region is poorly defined in the ultrasound image, and correspondingly, the atlas assigns lower certainty there, while the lateral walls that are clearly visible in the image receive higher certainty scores.

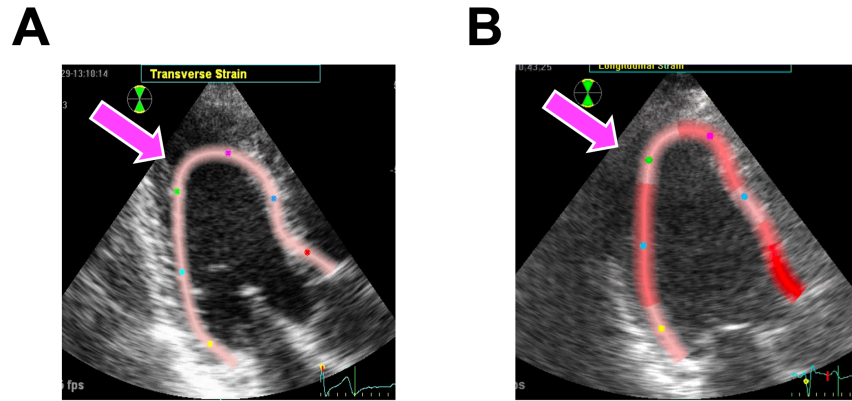


Figure 4.4: **Examples of left ventricle ultrasound acquisitions and their corresponding segmentations.** Figure A shows a higher-quality ultrasound acquisition with clear imaging of the left ventricle walls, enabling more confident and accurate segmentation of the side walls. However, some uncertainty remains at the apex due to intrinsic imaging limitations. Figure B illustrates a noisier ultrasound acquisition, which results in a more challenging segmentation process with greater uncertainty and variability. These examples highlight the subjective nature of manual segmentation, demonstrating how image quality can directly impacts segmentation reliability and downstream analyses.

4.4 Results

4.4.1 Echo-PIV Color-Doppler Fusion

Intracavitary velocity fields using echo-PIV can appreciably underestimate the high velocities in the diastolic jet when compared to color-Doppler data (Figure 4.6). In contrast, the hybrid modality overcame the planar flow assumption and resolved high velocity values, even for moderate acquisition frame rates. Normalized differences in radial velocity over the LV by color Doppler and echo-PIV to the hybrid modality were 0.46 ± 0.12 and 0 ± 0 , respectively. Echo-PIV and the hybrid modality underestimated the color-Doppler peak filling velocities by 0.13 ± 0.08 m/s and 0.02 ± 0.10 m/s, respectively.

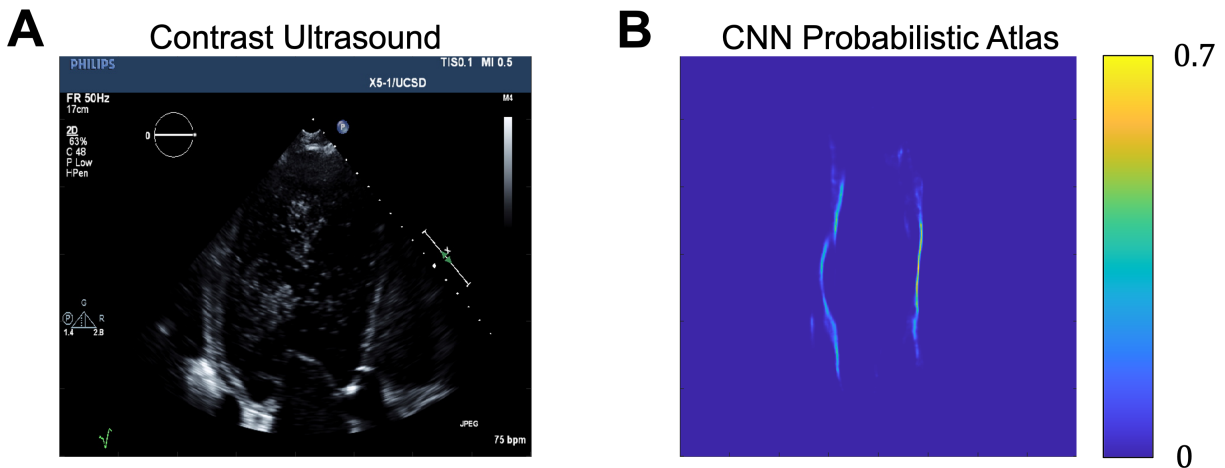


Figure 4.5: **CNN Segmentation.** Panel A shows the input ultrasound image fed into the convolutional neural network. Panel B displays the resulting probabilistic atlas, where values range from 0 to 1 — with 0 indicating no confidence in the wall location and 1 representing full confidence in the detected wall location.

4.4.2 Color-Doppler Multiscale Fusion

Color Doppler imaging faces the well-known challenge of balancing aliasing suppression with accurate measurement of low velocities, a trade-off governed by the choice of velocity encoding (VENC). From our multiscale Color Doppler acquisitions, we observe this trade-off in practice: a high VENC setting enables capture of clinically important high-speed events, such as left ventricular filling, but underestimates low velocities. Conversely, a low VENC captures slow flow more accurately, but introduces aliasing artifacts in high-velocity regions. With our multiscale fusion VFM approach, we combine high and low VENC datasets to produce a velocity field with a greater dynamic range and a more faithful representation of the true flow. Qualitatively, the fused method corrects aliased regions by relying more heavily on the high-VENC data, while the low-VENC field contributes to improved reconstruction in low-velocity areas. In this way, our azimuthal field reconstruction is derived from a more robust radial field, yielding improved accuracy across the entire velocity spectrum (Figure 4.9).

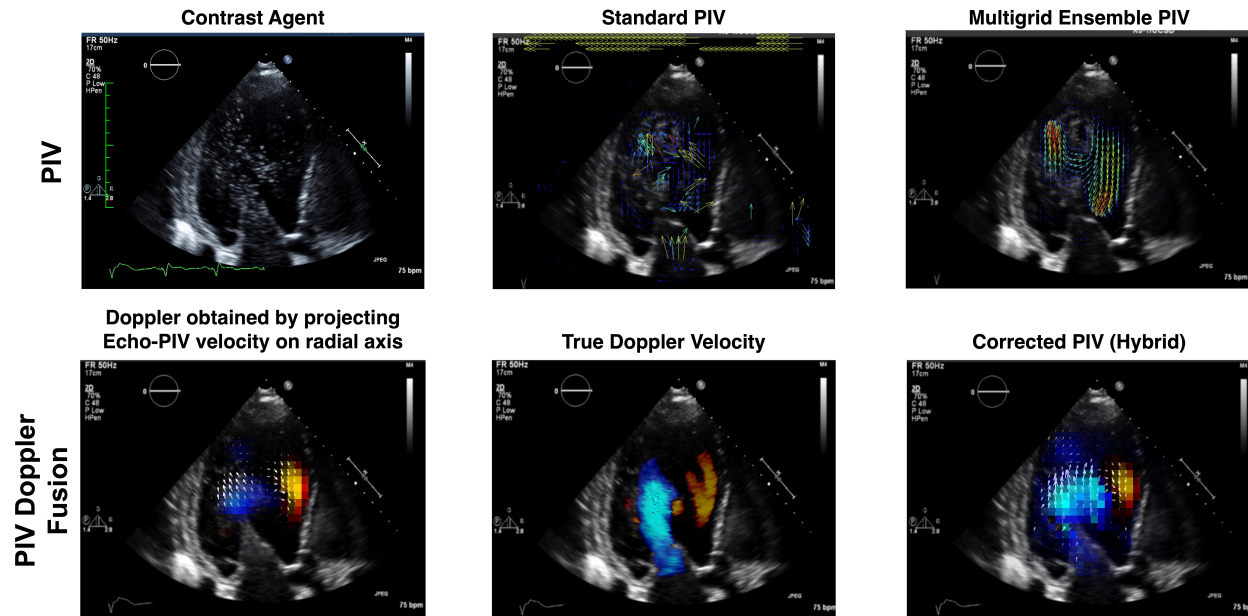


Figure 4.6: Echo-PIV Color-Doppler Fusion: Row 1 illustrates the comparison of obtaining echo-PIV velocity fields from the standard PIV method and the improved Multigrid-Ensemble method. Multigrid-Ensemble significantly improved the PIV result, providing higher dynamic range of the velocity field within the left ventricle. Row 2 illustrates the multimodality fusion results. After obtaining the 2D PIV velocity field, the result is projected into the polar coordinate system (ie the Doppler home coordinate system) to allow for fusion with Doppler acquisition. The final fused velocity field captures richer and more robust flow information by combining the strengths of both modalities.

4.5 Discussion

The results from both fusion strategies demonstrate the potential of integrating complementary echocardiographic modalities within a unified reconstruction framework to overcome individual limitations and enhance LV flow quantification. Echo-PIV and Doppler fusion successfully combines the broad spatial coverage of PIV on contrast ultrasound with the high-velocity resolution of Doppler, producing fields that more accurately capture diastolic jet dynamics while maintaining physical consistency. Similarly, Doppler multiscale fusion resolves the longstanding VENC trade-off in Color Doppler imaging by merging high and low VENC acquisitions to achieve a greater dynamic range, correcting aliased regions while

preserving low velocity fidelity. Although the present implementation is deterministic due to some fixed hyperparameter selection, it explicitly propagates uncertainty from the input data (color Doppler and segmentation variability, and PIV uncertainty) to the reconstructed fields. This feature makes the framework a natural foundation for fully Bayesian multimodality fusion, in which priors and hyperpriors would be treated probabilistically to produce principled uncertainty estimates. Such an extension to Bayesian methods is straightforward. Our algorithms described in Chapter 2 can be directly applied to this fused dataset to incorporate uncertainty quantification and prior modeling. Bayesian multimodality fusion would be particularly valuable in the clinical setting, where understanding the contributions of different error sources is essential for informed interpretation and decision making. By making uncertainty an explicit part of the flow mapping process, these fusion methods not only improve quantitative accuracy but also lay the groundwork for robust, interpretable, and reliable patient-specific diagnostics.

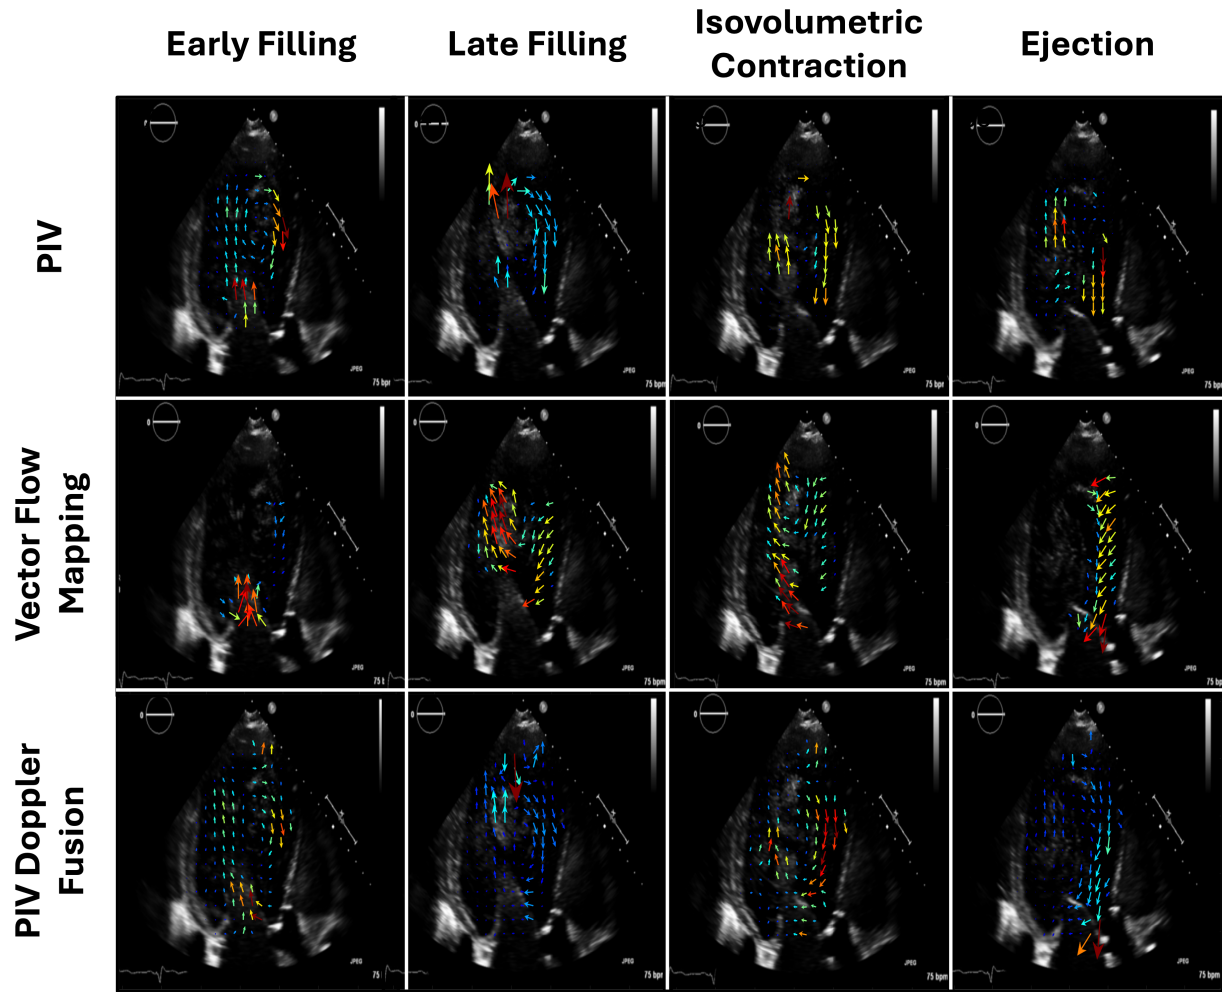


Figure 4.7: Comparison of PIV, Vector Flow Mapping, and Echo-PIV Color-Doppler Fusion: Three left ventricular (LV) flow quantification methods were applied to patient data at four key points in the cardiac cycle: early LV filling, late LV filling, isovolumetric contraction, and ejection. Row 1 shows the PIV velocity fields derived from ultrasound contrast images. Row 2 presents the vector fields obtained using classic Vector Flow Mapping. Row 3 displays velocity fields from the echo-PIV color-Doppler fusion method. Qualitatively, the fused data set combines the comprehensive spatial coverage of PIV with the high-velocity detail captured by color-Doppler, resulting in a richer and more informative flow field.

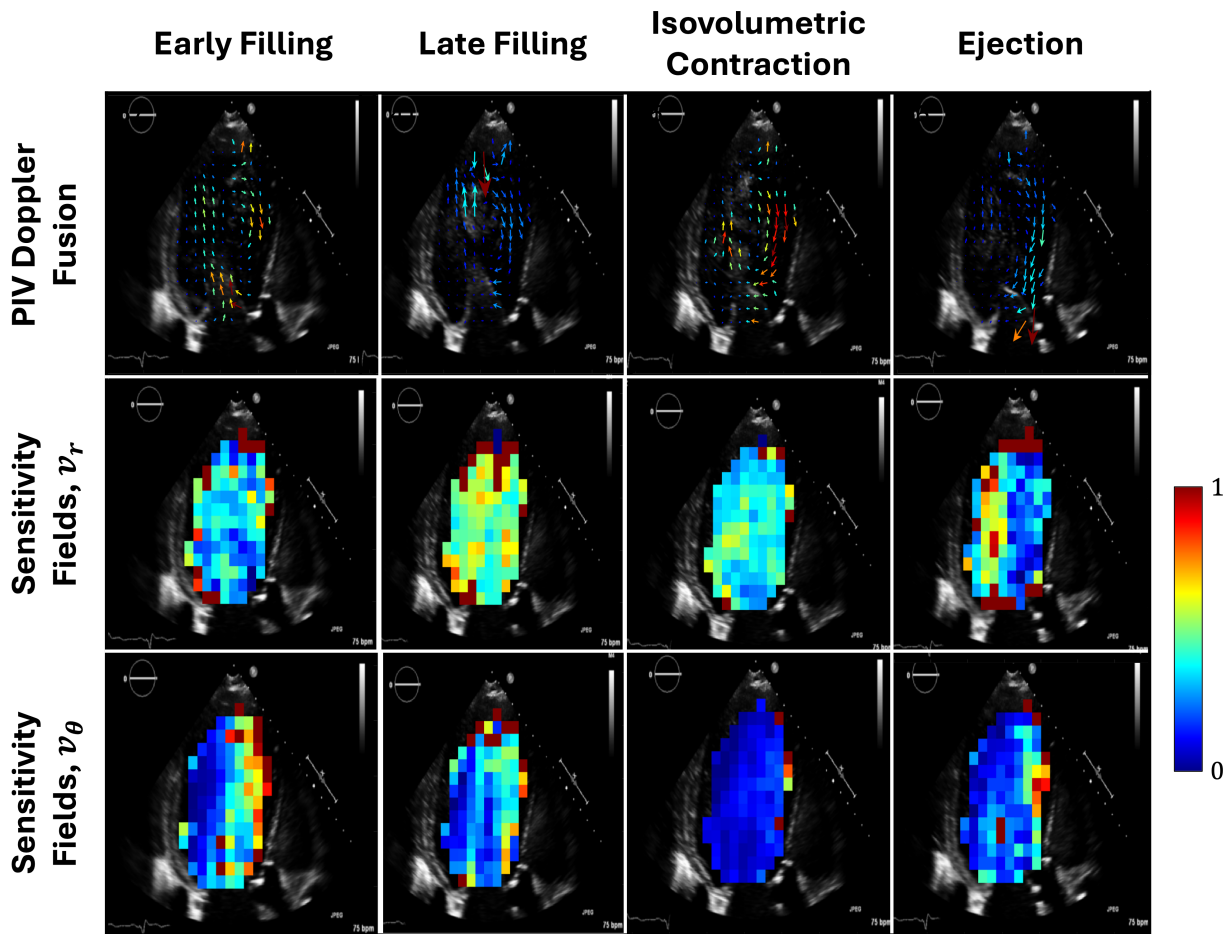


Figure 4.8: PIV–Doppler Fusion Sensitivity Maps (normalized). Row 1 shows the PIV Doppler fused velocity fields. Rows 2 and 3 present local sensitivity estimates in the radial and azimuthal directions. Higher local sensitivity appears near boundaries, due to input segmentation uncertainty, and within the inlet jet, where qualitative discrepancies between PIV and Doppler inputs are more pronounced.

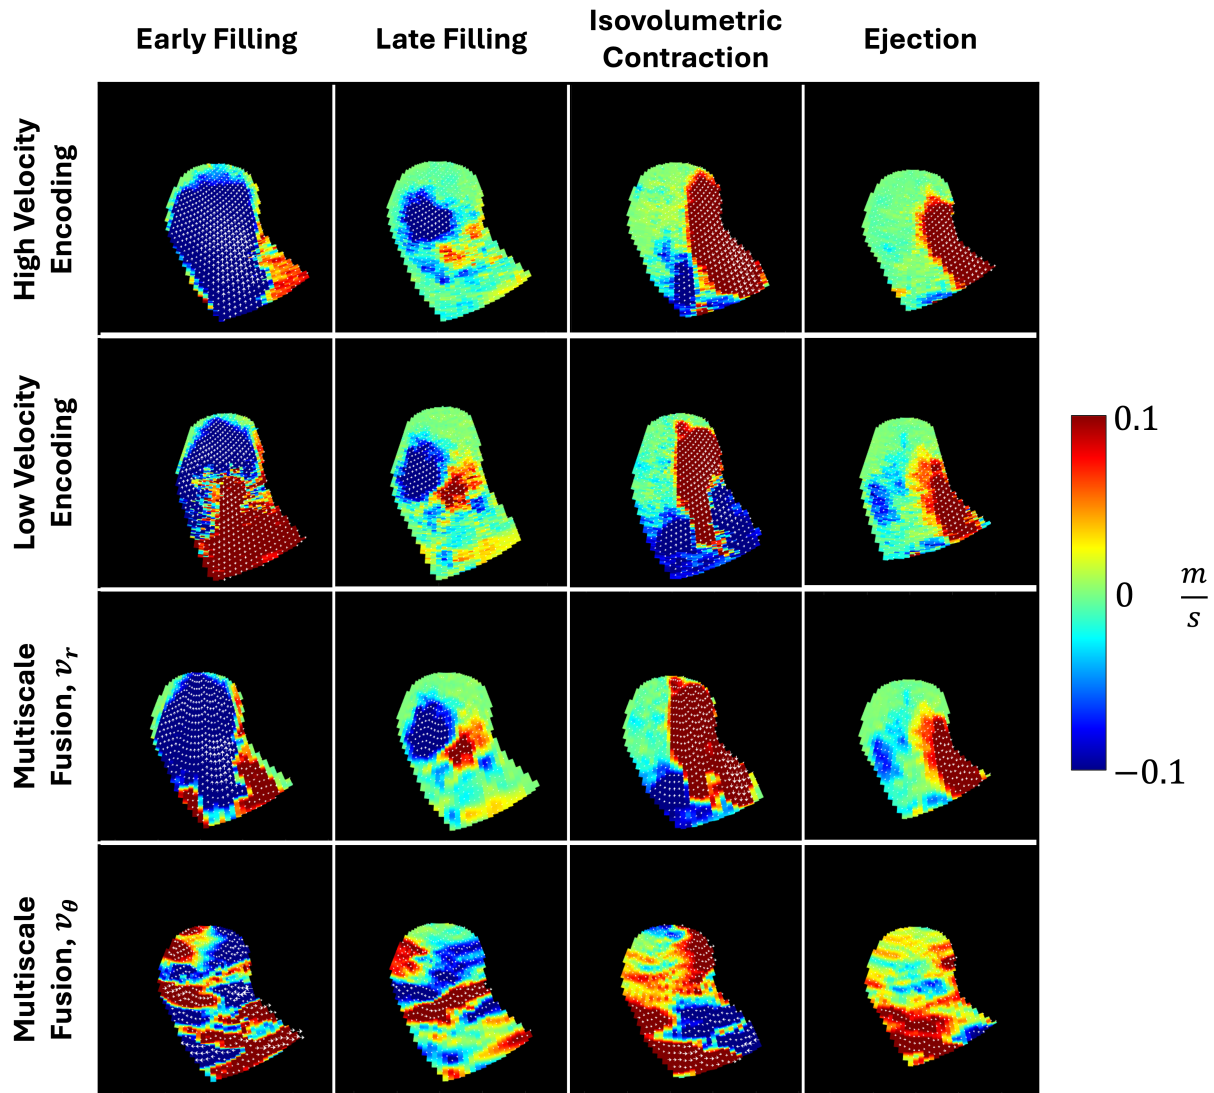


Figure 4.9: Doppler Multiscale Fusion at four key points in the cardiac cycle: early LV filling, late LV filling, isovolumetric contraction, and ejection. Rows 1 and 2 show color Doppler acquisitions obtained with high (0.66 m/s) and low (0.33 m/s) velocity encodings, respectively. Rows 3 and 4 present the corresponding radial and azimuthal velocity fields reconstructed using the Multiscale Fusion VFM method.

Chapter 5

3D AI VECTOR FLOW MAPPING

5.1 *Physics Informed Neural Networks (PINNs)*

Physics-informed neural networks (PINNs) leverage recent advances in deep learning to infer hidden variables from flow measurements while enforcing the governing equations of fluid mechanics [14, 38]. A key advantage of PINNs is their ability to learn from sparse datasets and adapt to patient-specific conditions, making them particularly promising for velocity field mapping (VFM). Unlike conventional machine learning models, PINNs integrate physical laws directly into the loss function, ensuring that predictions respect conservation principles.

However, despite growing interest and promising developments, the use of deep learning models, especially PINNs, for VFM remains relatively unexplored. Maidu et al introduced a novel AI-driven VFM approach using PINNs, incorporating mass conservation, momentum balance, and boundary conditions to reconstruct intraventricular flow and pressure fields from echocardiographic images [44]. AI-VFM performs phase unwrapping and fills in spatial and temporal gaps in color Doppler input, yielding super-resolution flow maps. Crucially, AI-VFM is guided entirely by patient-specific flow physics without relying on explicit smoothness priors or data from other patients or empirical models. Validation against computational fluid dynamics (CFD) ground truth demonstrated that AI-VFM outperforms both traditional VFM techniques and prior PINN-based methods that enforce only mass conservation.

5.2 *3D VFM from Triplane color-Doppler*

Triplane color-Doppler captures three orthogonal planes of the left ventricle, typically the apical four-chamber, two-chamber, and long-axis views at about 60° separation. This provides a more comprehensive view of blood flow in the left ventricle compared to a single

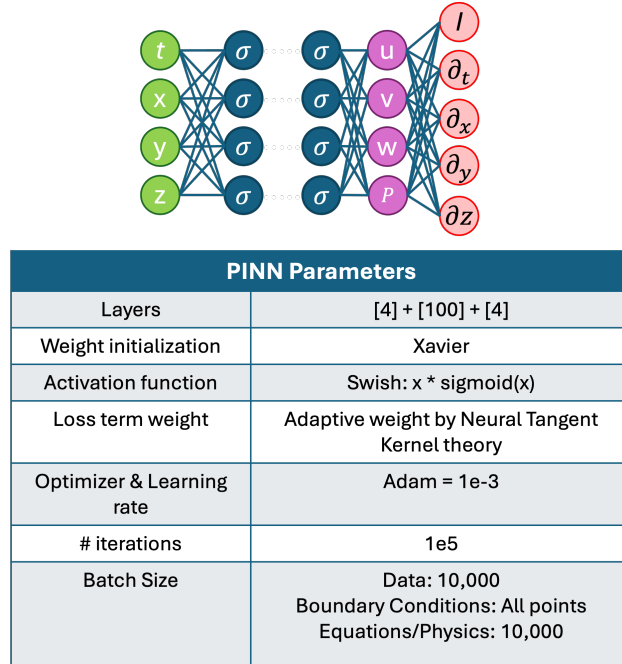


Figure 5.1: PINN Parameters

plane color Doppler. Vixege et al extended their previous 2D iVFM method to a 3D iVFM framework by applying vector flow mapping to triplane color Doppler data [79]. Traditional 2D and 3D VFM techniques typically rely on interpolation to estimate flow in regions lacking Doppler information, making them prone to inaccuracies in areas with data gaps or imaging artifacts. These methods enforce mass conservation and free-slip boundary conditions along the data planes but omit other critical physical constraints, such as momentum balance. Furthermore, they do not directly estimate flow related variables like pressure, which must be recovered through secondary analyses.

5.3 3D AI-VFM

Similarly, we extend our 2D AI-VFM method into 3D by applying AI-VFM on triplane color-Doppler data. As a result, not only do we reconstruct the cross-beam velocity, we also reconstruct fluctuating pressure fields. 3D AI-VFM enforces mass conservation and

momentum balance on a cartesian grid, automatically handling Doppler artifacts and data gaps, outperforming traditional VFM in accuracy. For validation, we use CFD-generated ground truth of an idealized ventricular flow. Three orthogonal slices separated by 60° are extracted from the CFD data and serve as the only training inputs. The underlying physics, however, is applied across the entire volume, enabling the network to reconstruct the full 3D flow field.

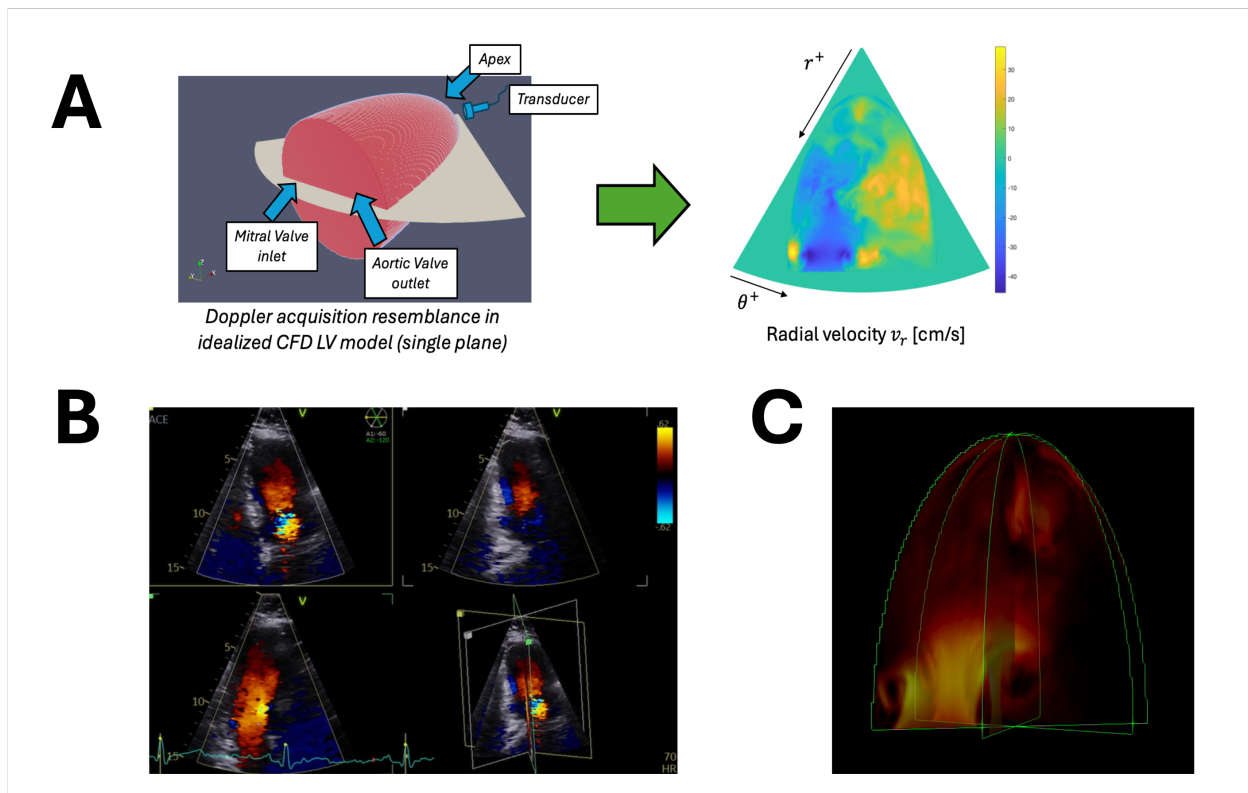


Figure 5.2: A) Synthetic Doppler slices from CFD Data. B) Ultrasound Triplane Clinical Acquisitions C) Synthetic Triplane Slices for PINNs training.

5.4 3D VFM from Triplane color-Doppler

5.5 Results

Qualitative comparisons between 3D AI-VFM and CFD ground truth demonstrate excellent agreement. The primary vortex structures during the filling phase, as well as secondary vortex structures emerging from vortex-wall interactions, are well captured by AI-VFM. These secondary vortices contribute to destabilization of the primary vortex through twisting-core instabilities, a phenomenon accurately reproduced by the model.

During late diastole, vortex breakdown is observed, and while velocity magnitudes are slightly underestimated, the overall flow pattern is consistent with CFD. Ejection-phase flow is similarly well reconstructed, indicating that AI-VFM captures key three-dimensional features of ventricular flow, including pressure distribution and oscillatory vortex dynamics.

Quantitative analysis shows excellent correlation for the U velocity component ($r^2 = 0.91$), moderate correlation for V ($r^2 = 0.65$), and lower correlation for W ($r^2 = 0.37$), consistent with the smaller magnitude of through-plane flow. Pressure estimates correlate strongly with CFD ($r^2 = 0.91$), highlighting the method's reliability for hemodynamic studies.

5.6 Residence Time

Another underlying model used in this method is the transport equation, which enables inference of clotting risk by estimating blood residence time. Neural networks tend to learn the underlying structure of the data, but are biased toward capturing low-frequency (long-wavelength) components first, a phenomenon known as spectral bias, which often results in smoother outputs. This behavior is a well-documented characteristic of physics-informed neural networks (PINNs), where smoothness is introduced not only through the network architecture and activation functions, but also inherently through the spectral bias. As a result, PINNs tend to underestimate the ground truth mean residence time, particularly due to their underestimation of the highest residence time values.




Figure 5.3: Correlation Plots

5.7 Discussion

In this chapter, we demonstrated the extension of AI-VFM to three-dimensional intraventricular flow using triplane color-Doppler data. By enforcing mass conservation and momentum balance within a PINN framework, 3D AI-VFM reconstructs both velocity and pressure fields with high fidelity, even in the presence of sparse or noisy data. The method captures key three-dimensional flow phenomena, including vortex formation, secondary instabilities, and vortex breakdown, which are critical for understanding cardiac hemodynamics. The integration of the transport equation allows for estimation of blood residence times, highlighting

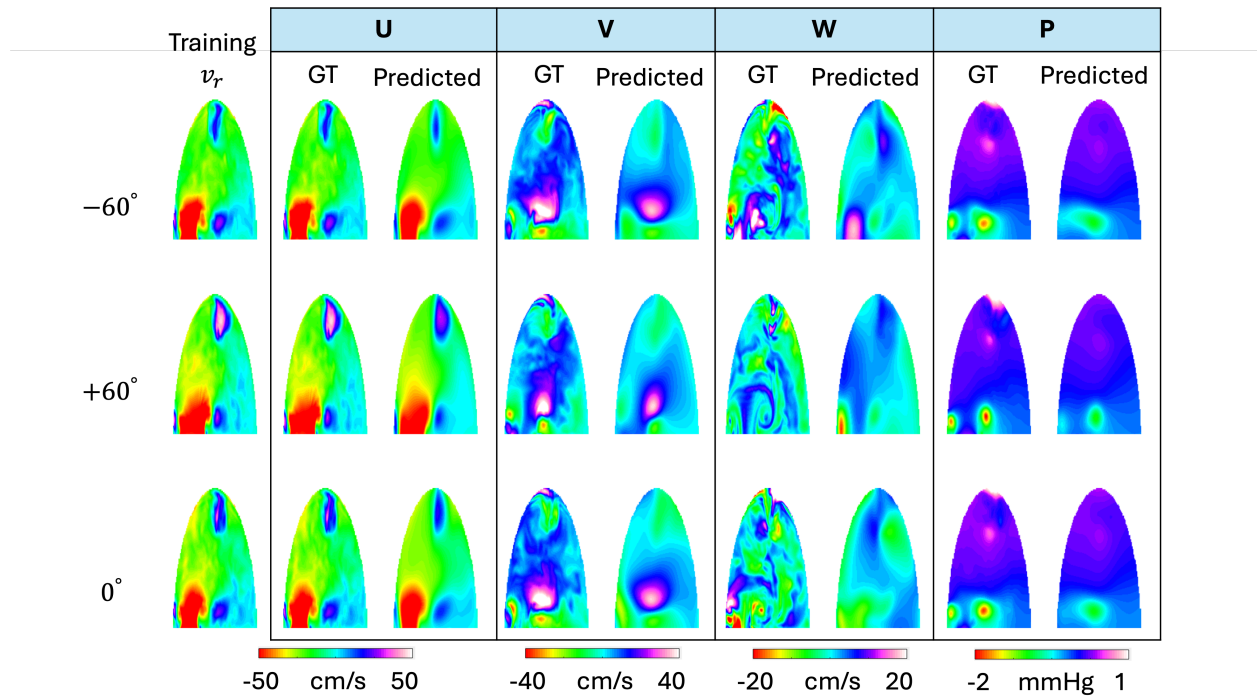


Figure 5.4: Planes trained with Doppler: Early Filling

potential applications in assessing thrombotic risk. While spectral bias leads to smoothing in extreme residence time regions, the overall spatial distribution and trends remain consistent with CFD.

It is important to note that these results are preliminary and based solely on CFD generated data. Further work is needed to validate 3D AI-VFM on patient echocardiograms, including studies on complex geometries, pathological flow conditions, and real clinical imaging artifacts. Despite these limitations, the method demonstrates several advantages over traditional VFM approaches, including physics guided reconstruction, the ability to handle data gaps and artifacts, and simultaneous estimation of velocity and pressure fields in three dimensions. Its adaptability to patient specific conditions will enable personalized hemodynamic assessment, bridging the gap between advanced computational modeling and clinical echocardiography.

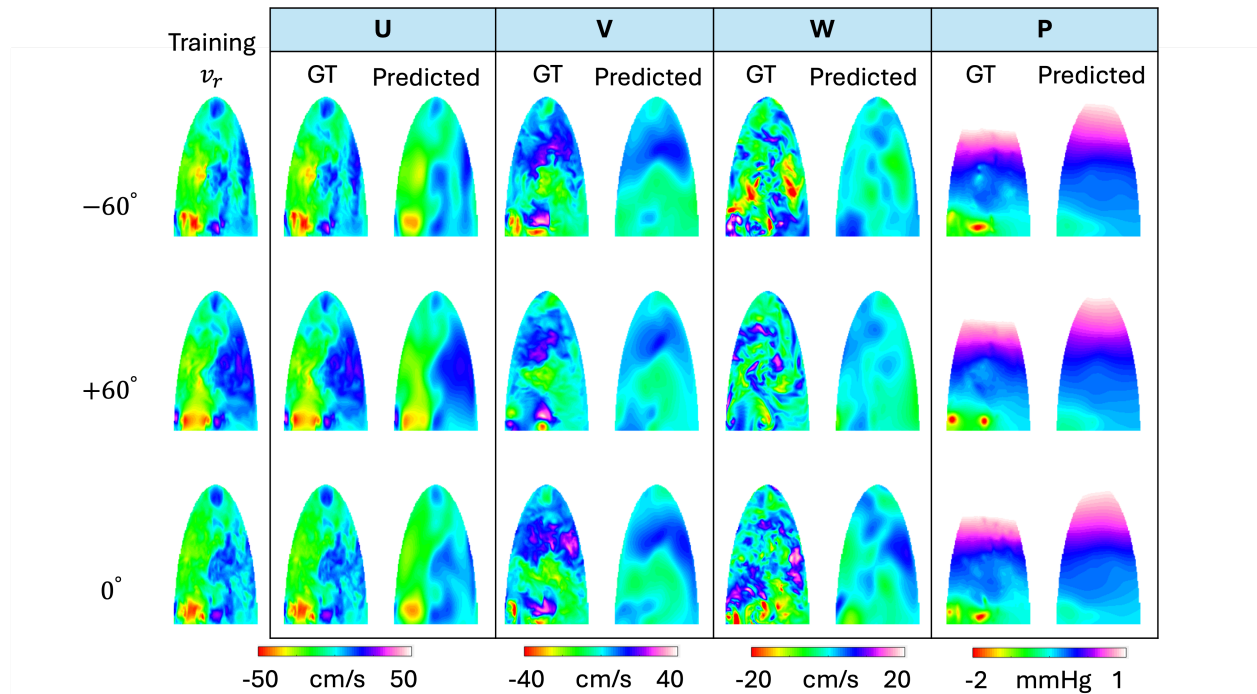


Figure 5.5: Planes trained with Doppler: Late Filling

Future work may also explore a Bayesian extension of AI-VFM, allowing the incorporation of prior knowledge and quantification of uncertainty in the reconstructed velocity and pressure fields. A Bayesian framework could provide confidence intervals for flow estimates, improve robustness to noisy or sparse input data, and facilitate probabilistic risk assessment for clinical decision making. Additionally, efforts to mitigate spectral bias and integrate the method into real time imaging workflows will further enhance the clinical applicability of AI-VFM. Overall, 3D AI-VFM represents a promising tool for high resolution, patient specific assessment of intracardiac flow, with the potential to provide new insights into cardiac function and hemodynamic health.

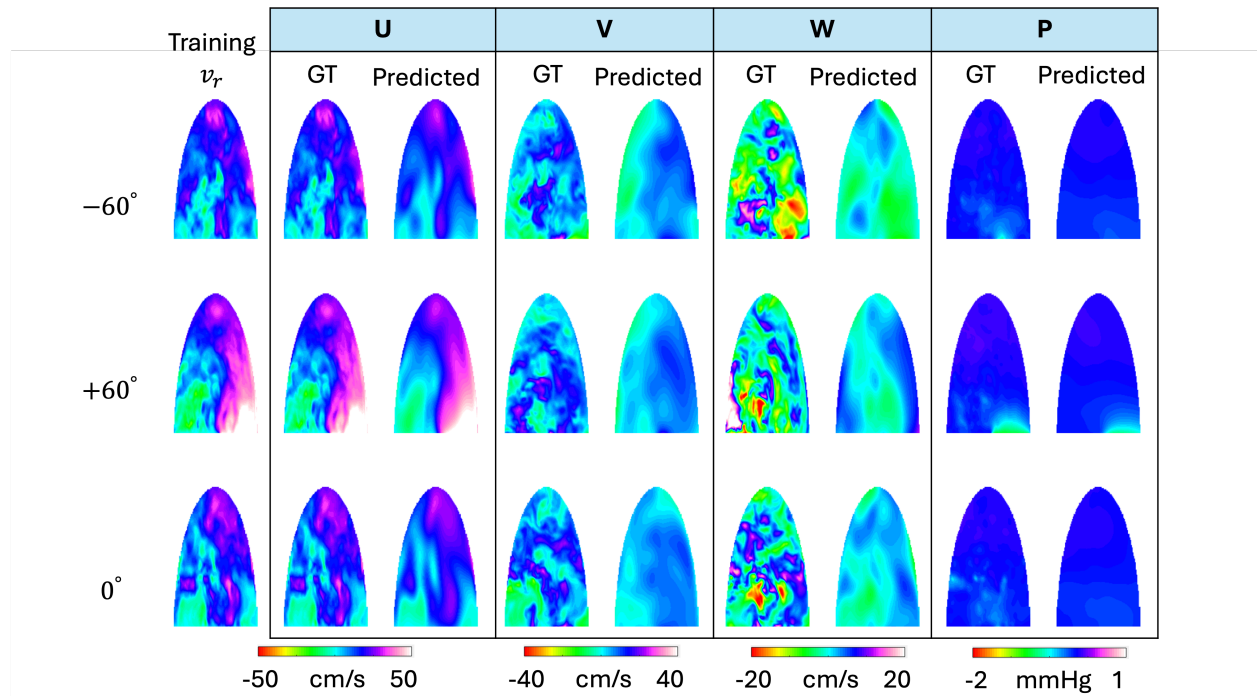


Figure 5.6: Planes trained with Doppler: Ejection

5.8 Acknowledgments

This chapter was carried out in collaboration with my colleague, Bahetihazi Maidu, who developed the underlying AI-VFM methodology on which this work builds.

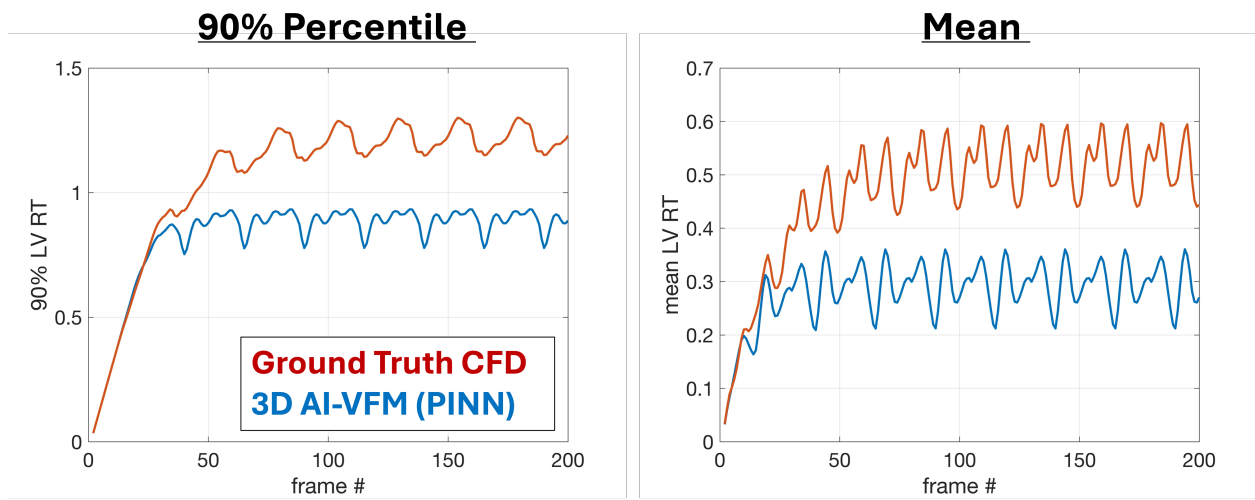


Figure 5.7: 90th percentile and mean left ventricle residence time for CFD ground truth vs 3D AI-VFM.

Chapter 6

CONCLUDING REMARKS

The central aim of this thesis was to develop a vector flow mapping framework that accounts for local measurement uncertainties, with the broader goal of complementing global metrics in cardiovascular clinical evaluation (Chapter 2). While this objective may appear straightforward, it quickly became clear that addressing it required careful consideration of several fundamental questions: How should local measurement uncertainties be defined? How can they be quantified in a manner that is both realistic and clinically relevant? The theoretical analysis presented in Chapter 3 addressed these questions by focusing specifically on the measurement process at the level of image acquisition. Within the Bayesian framework, this analysis contributes directly to the formulation of the error terms (ϵ_V, ϵ_M) , thereby linking data fidelity to realistic models of measurement variability. It also highlights the possibility of incorporating structured error into the inference process. For example, the systematic increase in uncertainty with radial depth, r .

Once this foundation was established, it became clear that uncertainties defined at the acquisition stage cannot be treated in isolation. They naturally propagate into the subsequent steps of reconstruction, influencing modeling choices, secondary analyses, and ultimately clinical interpretation. This recognition broadened the scope of the work beyond a single Bayesian formulation and led to the development of a more comprehensive roadmap for vector flow mapping. The roadmap emphasizes not only the propagation and management of uncertainty at each stage but also the systematic integration of prior knowledge about flow physics, boundary behavior, and measurement fidelity. In this sense, the thesis contributes both a methodological framework and a conceptual perspective for uncertainty-aware cardiovascular flow analysis.

Building on this roadmap, several avenues for future research emerge naturally. Boundary condition selection, for example, warrants careful consideration: while free-slip conditions are often sufficient in low-resolution ultrasound, no-slip constraints can become advantageous when wall motion or tangential flow information is available, enabling more physiologically consistent reconstructions. Another promising direction is the explicit enforcement of fluid dynamics principles, such as the Navier–Stokes equations. Within a Bayesian framework, these constraints can be incorporated as priors, capturing temporal and physical consistency without overly restricting flexibility. Finally, the choice of smoothness constraints, currently a Laplacian operator, can be revisited to better balance regularization with fidelity to complex flow patterns, potentially leveraging more sophisticated differential operators or data-driven priors.

Beyond these methodological refinements, the roadmap also highlights opportunities for broadening the scope of vector flow mapping. As demonstrated in Chapter 4, multimodality integration represents a particularly powerful extension. By fusing complementary imaging data, uncertainty can be more accurately characterized and flow estimates more robustly constrained, ultimately enhancing the clinical relevance of vector flow mapping. In parallel, emerging approaches such as 3D AI-VFM (Chapter 5) have shown remarkable progress by leveraging deep learning to directly reconstruct complex three-dimensional flows. While these methods are not Bayesian in their current form, the framework developed in this thesis suggests a pathway for incorporating explicit uncertainty modeling into AI-VFM, potentially combining the predictive power of machine learning with the interpretability and rigor of probabilistic inference.

Equally important, this framework enables sensitive assessment of flow dynamics, to complement current global metrics in cardiovascular clinical evaluations. Standard measurements such as ejection fraction or stroke volume remain essential, but they can remain unchanged even when local flow dynamics evolve. Vector flow mapping, particularly when uncertainty is quantified, offers a means to assess subtle, spatially heterogeneous changes in flow patterns over time, for example, in a patient imaged three months apart. Clinical scenarios

where such changes are highly relevant include left ventricular remodeling after myocardial infarction, where altered vortex patterns can reflect impaired diastolic filling; valvular disease, where regurgitant or stenotic jets can evolve in severity and spatial distribution; and cardiomyopathies, where abnormal intraventricular flow patterns often precede measurable declines in global systolic function [15, 19, 22, 51, 69]. In these contexts, flow field analysis provides a more granular perspective that complements, rather than replaces, global indices, thereby strengthening monitoring and risk assessment.

In conclusion, this thesis demonstrates that a rigorous, uncertainty-aware approach to vector flow mapping not only improves the fidelity of flow estimation but also provides a flexible, extensible framework for future methodological and clinical innovation. By connecting image acquisition, modeling, and analysis within a principled roadmap, this work lays the foundation for more reliable and interpretable cardiovascular flow assessment, and it opens the door to systematic exploration of advanced priors, multimodality fusion, and physics-informed inference.

BIBLIOGRAPHY

- [1] Koichi Akiyama, Paolo C. Colombo, Eric J. Stöhr, Ruiping Ji, Isaac Y. Wu, Keiichi Itatani, Shohei Miyazaki, Teruyasu Nishino, Naotoshi Nakamura, Yasufumi Nakajima, Barry J McDonnell, Koji Takeda, Melana Yuzefpolskaya, and Hiroo Takayama. Blood flow kinetic energy is a novel marker for right ventricular global systolic function in patients with left ventricular assist device therapy. *Frontiers in Cardiovascular Medicine*, 10, 2023.
- [2] P. M. Arvidsson, M. Ugander, J. Töger, E. Heiberg, M. Carlsson, H. Arheden, and S. J. Kovacs. Vortex ring behavior provides the epigenetic blueprint for the human heart. *Scientific Reports*, 6:22021, 2016.
- [3] Rei Asami, Tomohiko Tanaka, Ken ichi Kawabata, Kunio Hashiba, Takashi Okada, and Tomohide Nishiyama. Accuracy and limitations of vector flow mapping: left ventricular phantom validation using stereo particle image velocimetry. *Journal of Echocardiography*, 15(2):57–66, Jun 2017.
- [4] Rei Asami, Tomohiko Tanaka, Motochika Shimizu, Yoshinori Seki, Tomohide Nishiyama, Hajime Sakashita, and Takashi Okada. Ultrasonic vascular vector flow mapping for 2-d flow estimation. *Ultrasound in medicine & biology*, 45(7):1663–1674, 2019.
- [5] K. C. Assi, E. Gay, C. Chnafa, S. Mendez, F. Nicoud, Jfpj Abascal, P. Lantelme, F. Tournoux, and D. Garcia. Intraventricular vector flow mapping-a doppler-based regularized problem with automatic model selection. *Phys Med Biol*, 62(17):7131–7147, 2017. 1361-6560 Assi, Kondo Claude Gay, Etienne Chnafa, Christophe Mendez, Simon Nicoud, Franck Abascal, Juan F P J Lantelme, Pierre Tournoux, François Garcia, Damien Journal Article England 2017/08/12 Phys Med Biol. 2017 Aug 11;62(17):7131-7147. doi: 10.1088/1361-6560/aa7fe7.
- [6] M. Avesani, B. Degrelle, G. Di Salvo, J. B. Thambo, and X. Iriart. Vector flow mapping: A review from theory to practice. *Echocardiography*, 38(8):1405–1413, 2021. 1540-8175 Avesani, Martina Orcid: 0000-0003-4970-4562 Degrelle, Bastien Di Salvo, Giovanni Thambo, Jean-Benoit Iriart, Xavier Journal Article Review United States 2021/07/15 Echocardiography. 2021 Aug;38(8):1405-1413. doi: 10.1111/echo.15154. Epub 2021 Jul 14.

- [7] Yolanda Benito, Pablo Martinez-Legazpi, Lorenzo Rossini, Candelas Perez del Villar, Raquel Yotti, Daniel Rodríguez-Pérez, Jose Carlos Antoranz, Juan C del Álamo, and Javier Bermejo. Age-dependence of flow homeostasis in the left ventricle. *Frontiers in Physiology*, 10:445666, 2019.
- [8] Christian Schulte Zu Berge, Denis Declara, Christoph Hennersperger, Maximilian Baust, and Nassir Navab. Real-time uncertainty visualization for b-mode ultrasound. In *Proceedings of the 2015 IEEE Scientific Visualization Conference (SciVis)*, SCIVIS '15, page 33–40, USA, 2015. IEEE Computer Society.
- [9] Javier Bermejo, Yolanda Benito, Marta Alhama, Raquel Yotti, Pablo Martínez-Legazpi, Candelas Pérez del Villar, Esther Pérez-David, Ana González-Mansilla, Cristina Santa-Marta, Alicia Barrio, Francisco Fernández-Avilés, and Juan C. del Álamo. Intra-ventricular vortex properties in nonischemic dilated cardiomyopathy. *American Journal of Physiology-Heart and Circulatory Physiology*, 306(5):H718–H729, 2014. PMID: 24414062.
- [10] Javier Bermejo, Pablo Martínez-Legazpi, and Juan C. del Álamo. The clinical assessment of intraventricular flows. *Annual Review of Fluid Mechanics*, 47(1):315–342, 2015.
- [11] Jeff Borggaard, Nathan Glatt-Holtz, and Justin Krometis. A bayesian approach to estimating background flows from a passive scalar, 2020.
- [12] María Guadalupe Borja, Pablo Martinez-Legazpi, Cathleen Nguyen, Oscar Flores, Andrew M Kahn, Javier Bermejo, and Juan C Del Álamo. Deriving phenotype-representative left ventricular flow patterns by reduced-order modeling and classification. *Comput Biol Med*, 179:108760, June 2024.
- [13] Steve Brooks, Andrew Gelman, Galin Jones, and Xiao-Li Meng, editors. *Handbook of Markov Chain Monte Carlo*. Chapman & Hall/CRC, Boca Raton, 2011.
- [14] S. Cai, Z. Mao, Z. Wang, M. Yin, and G.E. Karniadakis. Physics-informed neural networks (PINNs) for fluid mechanics: a review. *Acta Mechanica Sinica*, 37(12):1727–1738, 2021.
- [15] B. T. Chan. Impact of myocardial infarction on intraventricular vortex and flow energetics assessed using computational simulations. *International Journal for Numerical Methods in Biomedical Engineering*, 35(6):e3204, June 2019.
- [16] Zhongxiu Chen, Yajiao Li, Chen Li, Hong Tang, Hui Wang, Yue Zhong, Yuyan Cai, and Li Rao. Right ventricular dissipative energy loss detected by vector flow mapping in children: characteristics of normal values. *Journal of Ultrasound in Medicine*, 38(1):131–140, 2019.

- [17] Mehmet Akif Cifci. A deep learning-based framework for uncertainty quantification in medical imaging using the dropweak technique: An empirical study with baresnet. *Diagnostics*, 13(4), 2023.
- [18] Lavsén Dahal, Aayush Kaffé, and Bishesh Khanal. Uncertainty estimation in deep 2d echocardiography segmentation, 2020.
- [19] Ahmet Demirkiran, Mariëlla E. C. J. Hassell, Pankaj Garg, Mohammed S. M. Elbaz, Ronak Delewi, John P. Greenwood, Jan J. Piek, Rob J. van der Geest, Robin Nijveldt, and Sven Plein. Left ventricular four-dimensional blood flow distribution, energetics, and vorticity in chronic myocardial infarction patients with/without left ventricular thrombus. *European Journal of Radiology*, 150:110233, May 2022.
- [20] Bradley Efron. Bootstrap methods: Another look at the jackknife. In Samuel Kotz and Norman L. Johnson, editors, *Breakthroughs in Statistics: Methodology and Distribution*, pages 569–593. Springer, New York, 1992. Reprinted from the *Annals of Statistics*, 7(1): 1–26, 1979.
- [21] Mohammed S. M. Elbaz, Emmeline E. Calkoen, Jos J. M. Westenberg, Boudewijn P. F. Lelieveldt, Arno A. W. Roest, and Rob J. van der Geest. Vortex flow during early and late left ventricular filling in normal subjects: quantitative characterization using retrospectively-gated 4d flow cardiovascular magnetic resonance and three-dimensional vortex core analysis. *Journal of Cardiovascular Magnetic Resonance*, 16(1):78, Sep 2014.
- [22] Nobuaki Fukuda, Keiichi Itatani, Koichi Kimura, Aya Ebihara, Kazuaki Negishi, Kansei Uno, Kagami Miyaji, Masahiko Kurabayashi, and Katsu Takenaka. Prolonged vortex formation during the ejection period in the left ventricle with low ejection fraction: a study by vector flow mapping. *Journal of Medical Ultrasonics*, 41(3):301–310, July 2014.
- [23] Giorgio Galanti, Gianni Pedrizzetti, Stefano Pedri, Laura Stefani, and Loira Toncelli. In vivo assesment of a novel technique for rest and stress intra-ventricular flow patterns in athletes, 12 2014.
- [24] Damien Garcia. Make the most of must, an open-source matlab ultrasound toolbox. In *2021 IEEE International Ultrasonics Symposium (IUS)*, pages 1–4, 2021.
- [25] Damien Garcia, Juan C. del Álamo, David Tanné, Raquel Yotti, Cristina Cortina, É Bertrand, José Carlos Antoranz, Esther Pérez-David, Régis Rieu, Francisco Fernández-Avilés, and Javier Bermejo. Two-dimensional intraventricular flow mapping by digital processing conventional color-doppler echocardiography images. *IEEE Transactions on Medical Imaging*, 29(10):1701–1713, 2010.

- [26] Andrew Gelman, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin. *Bayesian Data Analysis*. CRC Press, Boca Raton, FL, 3rd edition, 2013.
- [27] Christina Gillmann, Dorothee Saur, Thomas Wischgoll, and Gerik Scheuermann. Uncertainty-aware visualization in medical imaging - a survey. *Computer Graphics Forum*, 40:665–689, 06 2021.
- [28] Houssein-Eddine Gueziri, Michael J. McGuffin, and Catherine Laporte. Visualizing positional uncertainty in freehand 3D ultrasound. In Ziv R. Yaniv and David R. Holmes III, editors, *Medical Imaging 2014: Image-Guided Procedures, Robotic Interventions, and Modeling*, volume 9036, page 90361H. International Society for Optics and Photonics, SPIE, 2014.
- [29] Y. Han, L. Huang, Z. Li, X. Li, Y. Guo, J. Liu, C. Luo, H. Gao, L. He, H. Wang, and C. Xu. Relationship between left ventricular isovolumic relaxation flow patterns and mitral inflow patterns studied by using vector flow mapping. *Scientific Reports*, 9:16264, 2019.
- [30] Per Christian Hansen and Dianne Prost O’Leary. The use of the l-curve in the regularization of discrete ill-posed problems. *SIAM journal on scientific computing*, 14(6):1487–1503, 1993.
- [31] Hideyuki Hayashi, Koichi Akiyama, Keiichi Itatani, Scott DeRoo, Joseph Sanchez, Giovanni Ferrari, Paolo C Colombo, Koji Takeda, Isaac Y Wu, Atsushi Kainuma, et al. A novel in vivo assessment of fluid dynamics on aortic valve leaflet using epi-aortic echocardiogram. *Echocardiography*, 37(2):323–330, 2020.
- [32] Sahar Hendabadi, Javier Bermejo, Yolanda Benito, Raquel Yotti, Francisco Fernández-Avilés, Juan C Del Álamo, and Shawn C Shadden. Topology of blood transport in the human left ventricle by novel processing of doppler echocardiography. *Annals of biomedical engineering*, 41:2603–2616, 2013.
- [33] Takashi Honda, Keiichi Itatani, Kagami Miyaji, and Masahiro Ishii. Assessment of the vortex flow in the post-stenotic dilatation above the pulmonary valve stenosis in an infant using echocardiography vector flow mapping. *European Heart Journal*, 35(5):306–306, 10 2013.
- [34] Geu-Ru Hong, Minji Kim, Gianni Pedrizzetti, and Mani A. Vannan. Current clinical application of intracardiac flow analysis using echocardiography. *Journal of cardiovascular ultrasound*, 21(4):155–162, 2013. 24459561[pmid] PMC3894365[pmcid].

- [35] M.F. Hutchinson. A stochastic estimator of the trace of the influence matrix for laplacian smoothing splines. *Communications in Statistics - Simulation and Computation*, 19(2):433–450, 1990.
- [36] Keiichi Itatani, Takashi Okada, Tokuhisa Uejima, Tomohiko Tanaka, Minoru Ono, Kagami Miyaji, and Katsu Takenaka. Intraventricular flow velocity vector visualization based on the continuity equation and measurements of vorticity and wall shear stress. *Japanese Journal of Applied Physics*, 52, 2013.
- [37] Jørgen Arendt Jensen, Svetoslav Ivanov Nikolov, Alfred C. H. Yu, and Damien Garcia. Ultrasound vector flow imaging—part i: Sequential systems. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 63(11):1704–1721, 2016.
- [38] G.E. Karniadakis, I.G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.
- [39] C. Kasai, K. Namekawa, A. Koyano, and R. Omoto. Real-time two-dimensional blood flow imaging using an autocorrelation technique. *IEEE Transactions on Sonics and Ultrasonics*, 32(3):458–464, 1985.
- [40] Alexandros Kontogiannis, Scott V Elgersma, Andrew J Sederman, and Matthew Juniper. Bayesian inverse navier-stokes problems: joint flow field reconstruction and parameter learning. 2024.
- [41] Roberto M. Lang, Luigi P. Badano, Victor Mor-Avi, Jonathan Afilalo, Anderson Armstrong, Laura Ernande, Frank A. Flachskampf, Elyse Foster, Steven A. Goldstein, Tatiana Kuznetsova, Patrizio Lancellotti, Denisa Muraru, Michael H. Picard, Ernst R. Rietzschel, Lawrence Rudski, Kirk T. Spencer, Wendy Tsang, and Jens-Uwe Voigt. Recommendations for cardiac chamber quantification by echocardiography in adults: An update from the american society of echocardiography and the european association of cardiovascular imaging. *Journal of the American Society of Echocardiography*, 28(1):1–39.e14, 2015.
- [42] Chen Li, Juqian Zhang, Xiaoqing Li, Can Zhou, Haihua Li, Hong Tang, and Li Rao. Quantification of chronic aortic regurgitation by vector flow mapping: a novel echocardiographic method. *European Journal of Echocardiography*, 11(2):119–124, 11 2010.
- [43] Michael Loecher, Eric Schrauben, Kevin M. Johnson, and Oliver Wieben. Phase unwrapping in 4d mr flow with a 4d single-step laplacian algorithm. *Journal of Magnetic Resonance Imaging*, 43(4):833–842, 2016.

- [44] Bahetihazi Maidu, Pablo Martinez-Legazpi, Manuel Guerrero-Hurtado, Cathleen M. Nguyen, Alejandro Gonzalo, Andrew M. Kahn, Javier Bermejo, Oscar Flores, and Juan C. del Alamo. Super-resolution left ventricular flow and pressure mapping by navier-stokes-informed neural networks. *bioRxiv*, 2024.
- [45] Barry J Maron, Paolo Spirito, Yolanda Wesley, and Julio Arce. Apical hypertrophic cardiomyopathy: clinical, echocardiographic and electrocardiographic features. *Journal of the American College of Cardiology*, 20(2):295–301, 1992.
- [46] Pablo Martinez-Legazpi, Lorenzo Rossini, Candelas Pérez del Villar, Yolanda Benito, Carolina Devesa-Cordero, Raquel Yotti, Antonia Delgado-Montero, Ana Gonzalez-Mansilla, Andrew M Kahn, Francisco Fernandez-Avilés, et al. Stasis mapping using ultrasound: a prospective study in acute myocardial infarction. *JACC: Cardiovascular Imaging*, 11(3):514–515, 2018.
- [47] Pablo Martínez-Legazpi et al. Contribution of the diastolic vortex ring to left ventricular filling. *Journal of the American College of Cardiology*, 64(16):1711–21, 2014.
- [48] Donato Mele, Vittorio Smarrazzo, Gianni Pedrizzetti, Fabio Capasso, Marco Pepe, Salvatore Severino, Giovanni Andrea Luisi, Marco Maglione, and Roberto Ferrari. Intracardiac flow analysis: Techniques and potential clinical applications. *Journal of the American Society of Echocardiography*, 32(3):319–332, 2019.
- [49] Brett A. Meyers, Craig J. Goergen, Patrick Segers, and Pavlos P. Vlachos. Colour-doppler echocardiography flow field velocity reconstruction using a streamfunction-vorticity formulation. *Journal of the Royal Society, Interface*, 17(173):20200741–20200741, 2020. 33259749[pmid] PMC7811584[pmcid].
- [50] Brett A. Meyers, Jiacheng Zhang, Jonathan Nyce, Yue-Hin Loke, and Pavlos P. Vlachos. Enhanced echocardiographic assessment of intracardiac flow in congenital heart disease. *PLOS ONE*, 19(3):1–14, 03 2024.
- [51] Mahesh S. Nagargoje, Eneko Lazpita, Jesús Garicano-Mena, and Soledad Le Clainche. Review on vortex dynamics in the left ventricle as an early diagnosis marker for heart diseases and its treatment outcomes. *arXiv:2411.11441*, 2024. Accessed: 2025-08-20.
- [52] Gianni Pedrizzetti and Giovanni Tonti. Method of transforming a doppler velocity dataset into a velocity vector field, October 18 2012. US Patent App. 13/449,634.
- [53] Piotr Ponikowski and Adriaan A. Voors. Role of echocardiography in the diagnosis of dilated cardiomyopathy. In *Heart Failure*. Elsevier, 2022. Chapter discussing echo limitations in DCM.

- [54] Andrea Postigo, Federica Viola, Christian Chazo, Pablo Martínez-Legazpi, Ana González-Mansilla, Elena Rodríguez-González, Francisco Fernández-Avilés, Juan C Del Álamo, Tino Ebbers, and Javier Bermejo. Assessment of blood flow transport in the left ventricle using ultrasound. validation against 4-d flow cardiac magnetic resonance. *Ultrasound in medicine & biology*, 48(9):1822–1832, 2022.
- [55] Richard Ro, Dan Halpern, David J Sahn, Peter Homel, Milla Arabadjian, Charles Lopresto, and Mark V Sherrid. Vector flow mapping in obstructive hypertrophic cardiomyopathy to assess the relationship of early systolic left ventricular flow and the mitral valve. *Journal of the American College of Cardiology*, 64(19):1984–1995, 2014.
- [56] Herbert Robbins. An empirical bayes approach to statistics. *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability*, 1:157–163, 1956.
- [57] Christian P. Robert and George Casella. *Monte Carlo Statistical Methods*. Springer, New York, 2nd edition, 2004.
- [58] Olaf Rodevand, Reidar Bjornerheim, Thor Edvardsen, Otto A. Smiseth, and Halfdan Ihlen. Diastolic flow pattern in the normal left ventricle. *Journal of the American Society of Echocardiography*, 12(6):500–507, 1999.
- [59] Elena Rodríguez-González, Pablo Martínez-Legazpi, Ana González-Mansilla, M Ángeles Espinosa, Teresa Mombiela, Juan A Guzmán De-Villoria, Maria Guadalupe Borja, Fernando Díaz-Otero, Rubén Gómez de Antonio, Pilar Fernández-García, et al. Cardiac stasis imaging, stroke, and silent brain infarcts in patients with nonischemic dilated cardiomyopathy. *American Journal of Physiology-Heart and Circulatory Physiology*, 327(2):H446–H453, 2024.
- [60] Elena Rodríguez-González, Pablo Martínez-Legazpi, Teresa Mombiela, Ana González-Mansilla, Antonia Delgado-Montero, Juan A Guzmán-De-Villoria, Fernando Díaz-Otero, Raquel Prieto-Arévalo, Miriam Juárez, María Del Carmen García Del Rey, et al. Stasis imaging predicts the risk of cardioembolic events related to acute myocardial infarction: the isbitami study. *Revista Española de Cardiología (English Edition)*, 2024.
- [61] Daniel Rodríguez Muñoz, Cristina Lozano Granero, and José Luis Zamorano. Vector flow mapping in mitral valve disease: A novel method for the assessment of flow mechanics and their potential implications for mitral valve repair. *Current Cardiovascular Imaging Reports*, 8:41, 2015.
- [62] Daniel Rodríguez Muñoz, José Luis Moya Mur, Javier Moreno, Covadonga Fernández-Golfín, Eduardo Franco, Juan Manuel Monteagudo, Roberto Matía, Antonio Hernández-Madrid, and José Luis Zamorano. Energy dissipation in resynchronization therapy:

- Impact of atrioventricular delay. *Journal of the American Society of Echocardiography*, 32(6):744–754.e1, 2019. American Society of Echocardiography 30th Annual Scientific Sessions.
- [63] Sandro Rossi, Laura Rosa, Valentina Ravetta, Alessandro Cascina, Pietro Quaretti, Andrea Azzaretti, Paola Scagnelli, Carmine Tinelli, Paolo Dionigi, and Fabrizio Calliada. Contrast-enhanced versus conventional and color doppler sonography for the detection of thrombosis of the portal and hepatic venous systems. *American Journal of Roentgenology*, 186(3):763–773, 2006. PMID: 16498104.
- [64] Lorenzo Rossini, Oscar Ö Braun, Michela Brambatti, Yolanda Benito, Adam Mizeracki, Marissa Miramontes, Cathleen Nguyen, Pablo Martinez-Legazpi, Shone Almeida, Megan Kraushaar, et al. Intraventricular flow patterns in patients treated with left ventricular assist devices. *ASAIO Journal*, 67(1):74–83, 2021.
- [65] Lorenzo Rossini, Pablo Martinez-Legazpi, Yolanda Benito, Candelas Pérez del Villar, Ana Gonzalez-Mansilla, Alicia Barrio, María-Guadalupe Borja, Raquel Yotti, Andrew M Kahn, Shawn C Shadden, et al. Clinical assessment of intraventricular blood transport in patients undergoing cardiac resynchronization therapy. *Meccanica*, 52(3):563, 2017.
- [66] Lorenzo Rossini, Pablo Martinez-Legazpi, Vi Vu, Leticia Fernández-Friera, Candelas Pérez del Villar, Sara Rodríguez-López, Yolanda Benito, María-Guadalupe Borja, David Pastor-Escuredo, Raquel Yotti, María J. Ledesma-Carbayo, Andrew M. Kahn, Borja Ibáñez, Francisco Fernández-Avilés, Karen May-Newman, Javier Bermejo, and Juan C. del Álamo. A clinical method for mapping and quantifying blood stasis in the left ventricle. *Journal of Biomechanics*, 49(11):2152–2161, 2016. Selected Articles from the International Conference on CFD in Medicine and Biology (Albufeira, Portugal – August 30th - September 4th, 2015).
- [67] Fulvio Scarano and Michael L Riethmuller. Iterative multigrid approach in piv image processing with discrete window offset. *Experiments in Fluids*, 26:513–523, 1999.
- [68] S. Shahriari and D. Garcia. Meshfree simulations of ultrasound vector flow imaging using smoothed particle hydrodynamics. *Phys Med Biol*, 63(20):205011, 2018. 1361-6560 Shahriari, Shahrokh Garcia, Damien Journal Article Research Support, Non-U.S. Gov't England 2018/09/25 Phys Med Biol. 2018 Oct 17;63(20):205011. doi: 10.1088/1361-6560/aae3c3.
- [69] Jang-Won Son, Won-Jong Park, Jung-Hyun Choi, Helene Houle, Mani A. Vannan, Geu-Ru Hong, and Namsik Chung. Abnormal left ventricular vortex flow patterns in association with left ventricular apical thrombus formation in patients with anterior

- myocardial infarction: a quantitative analysis by contrast echocardiography. *Circulation Journal*, 76(11):2640–2646, 2012.
- [70] Ze-Zhou Song. Echocardiography in the diagnosis of left ventricular noncompaction. *Cardiovascular Ultrasound*, 6:64, 2008.
- [71] Jie Sun, Fernando J Quevedo, and Erik Bollt. Bayesian optical flow with uncertainty quantification. *Inverse Problems*, 34(10):105008, August 2018.
- [72] Tomohiko Tanaka, R. Asami, Ken-ichi Kawabata, Keiichi Itatani, T. Uejima, T. Nishiyama, and T. Okada. Intracardiac vfm technique using diagnostic ultrasound system. 64:488–492, 2015.
- [73] Tomohiko Tanaka, Rei Asami, Ken-ichi Kawabata, Kunio Hashiba, Takashi Okada, and Tomohide Nishiyama. A posteriori accuracy estimation of ultrasonic vector-flow mapping (vfm). *Journal of Visualization*, 20, 02 2017.
- [74] Tomohiko Tanaka, Takashi Okada, Tomohide Nishiyama, and Yoshinori Seki. Relative pressure imaging in left ventricle using ultrasonic vector flow mapping. *Japanese Journal of Applied Physics*, 56(7S1):07JF26, jun 2017.
- [75] Jok M. Tang and Yousef Saad. A probing method for computing the diagonal of a matrix inverse. *Numerical Linear Algebra with Applications*, 19(3):485–501, 2012.
- [76] Richard B Thompson and Elliot R McVeigh. Fast measurement of intracardiac pressure differences with 2d breath-hold phase-contrast mri. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 49(6):1056–1066, 2003.
- [77] Tokuhi Uejima, Akira Koike, Hitoshi Sawada, Tadanori Aizawa, Shigeo Ohtsuki, Motonao Tanaka, Tetsushi Furukawa, and Alan G. Fraser. A new echocardiographic method for identifying vortex flow in the left ventricle: Numerical validation. *Ultrasound in Medicine and Biology*, 36(5):772–788, May 2010.
- [78] Toshinori Utsunomiya, Toshio Ogawa, S.William King, Eric Sunada, G.Wayne Moore, Walter L. Henry, and Julius M. Gardin. Effect of machine parameters on variance display in doppler color flow mapping. *American Heart Journal*, 120(6, Part 1):1395–1402, 1990.
- [79] Florian Vixège, Alain Berod, Pierre-Yves Courand, Simon Mendez, Franck Nicoud, Philippe Blanc-Benon, Didier Vray, and Damien Garcia. Full-volume three-component intraventricular vector flow mapping by triplane color doppler. *Physics in Medicine & Biology*, 67(9):095002, 2022.

- [80] Florian Vixège, Alain Berod, Yunyun Sun, Simon Mendez, Olivier Bernard, Nicolas Ducros, Pierre-Yves Courand, Franck Nicoud, and Damien Garcia. Physics-constrained intraventricular vector flow mapping by color doppler. page arXiv:2107.08104, July 01, 2021 2021.
- [81] Liang Wang, Patrick Clarysse, Zhengjun Liu, Bin Gao, Wanyu Liu, Pierre Croisille, and Philippe Delachartre. A gradient-based optical-flow cardiac motion estimation method for cine and tagged mr images. *Medical Image Analysis*, 57:136–148, 2019.
- [82] Ke Zou, Zhihao Chen, Xuedong Yuan, Xiaojing Shen, Meng Wang, and Huazhu Fu. A review of uncertainty estimation and its application in medical imaging, 2023.