Quantitative Flow Components in 4D Flow MRI

Axel Jensen & Fredrik Hasselberg

June 2026

Master's Thesis in

Biomedical Engineering

Supervisors: Einar Heiberg, Magnus Cinthio

LUND
UNIVERSITY

Medviso

Faculty of Engineering LTH

Department of Biomedical Engineering

Abstract

Four-dimensional flow magnetic resonance imaging (4D flow MRI) enables time-resolved assessment of intracardiac blood transport, but quantitative analysis requires workflows that convert complex velocity data into consistent and interpretable measures. This thesis aimed to develop and evaluate a semi-automated workflow in Segment for quantitative left-ventricular (LV) flow-component analysis.

The workflow combines cine MRI-derived LV geometry with 4D flow MRI velocity data. An LV seed region is generated from cine short-axis segmentation and transferred to the 4D flow analysis space. From matched seed positions, pathlines are traced forward and backward in time using fourth-order Runge–Kutta (RK4) integration. Crossings of mitral and aortic analysis planes are used to classify particles into the four established LV flow components: Direct Flow, Retained Inflow, Delayed Ejection Flow, and Residual Volume. The implementation also provides visual output for quality control and calculates time-resolved LV kinetic energy as a complementary measure.

Technical evaluation was performed using two controlled phantoms. In a circular-flow phantom, RK4 reproduced the analytic particle trajectory with near-perfect agreement and performed better than forward Euler. In a three-cylinder vessel phantom, the implemented plane-crossing and classification logic exactly reproduced the predefined component assignments. The workflow was then applied to 60 retrospective human datasets, including healthy adults, patients with heart failure, and healthy children. In the heart-failure datasets, the group with reduced ejection fraction showed markedly lower Direct Flow and higher Residual Volume than groups with higher ejection fraction, consistent with previously reported changes in LV flow-component distributions in heart failure. Automatic and manual plane-placement analyses produced similar group-level internal-consistency measures, while dataset-specific failures showed that visual review and manual adjustment are still needed. A proof-of-concept analysis further demonstrated the technical feasibility of applying the general framework to right-ventricular data.

In conclusion, this thesis presents a practical Segment-based research workflow for quantitative LV flow-component analysis from 4D flow MRI. The method integrates particle tracing, rule-based classification, visualization, and kinetic-energy output within one analysis environment. Further development should focus on improving robustness, reducing user dependence, and strengthening the workflow for reproducible research use.

Acknowledgements

We would like to thank our supervisor, Einar Heiberg, for his guidance throughout this thesis. His ideas, feedback, and help in solving practical and methodological challenges were important to this work. We would also like to thank our co-supervisor, Magnus Cinthio, for his support and input during the project.

We would like to thank everyone at Medviso for welcoming us and for providing a supportive working environment. In particular, we would like to thank Helen Fransson for her help and feedback throughout the project, Fanny Månefjord for her guidance in working with Segment, Jelena Block for her help and input, and Anette Wolf for her help with the design.

We would also like to thank the Lund Cardiac MR group for being welcoming and for contributing to a helpful and encouraging research environment.

Finally, we would like to thank our peers, Julia Nilsson and Elisabeth Dierschke Passikhani, for their help and encouragement throughout the thesis work.

Use of Generative AI

Generative AI tools were used during the project to support language revision, text structuring, and code-related problem solving. For code-related work, the tools were used for debugging suggestions, explanations of programming concepts, and ideas for possible implementation approaches. The final implementation was written, reviewed, tested, and integrated by the authors.

Generative AI tools were not used as scientific sources. All references, methodological choices, results, figures, and interpretations were checked and approved by the authors.

The authors take full responsibility for the content of this thesis.

Sammanfattning

4D-flödes-MR – en metod för att följa blodets väg genom hjärtat

Hur väl hjärtat fungerar avgörs inte bara av hur mycket blod det pumpar ut, utan också av hur blodet rör sig genom hjärtat. Med 4D-flödes-MR går det att följa blodets väg genom hjärtats vänstra kammare och upptäcka förändringar som inte syns med vanliga mått på pumpförmåga.

Vid undersökning av hjärtat används ofta ejektionsfraktion, ett mått på hur stor andel av blodet i vänster kammare som pumpas ut vid varje hjärtslag. Måttet är viktigt vid exempelvis hjärtsvikt, men det beskriver inte hur blodet rör sig inne i hjärtat. Två personer kan därför ha liknande pumpförmåga, samtidigt som blodets väg genom vänster kammare ser olika ut. För att bättre förstå hjärtats funktion behövs metoder som inte bara mäter hur mycket blod som pumpas ut, utan också hur blodet transporteras.

Med 4D-flödes-MR kan blodets rörelse mätas i tre dimensioner under hela hjärtslaget. Genom att följa virtuella partiklar i blodflödet går det att dela in blodet i olika grupper: blod som kommer in i vänster kammare och pumpas ut direkt, blod som kommer in men stannar kvar, blod som redan finns i kammaren och pumpas ut, samt blod som blir kvar under längre tid. En sådan indelning kan ge information om hur effektivt blodet passerar genom hjärtat, men analysen kräver ett verktyg som kan omvandla de omfattande MR-data som samlas in till tydliga och jämförbara resultat.

I detta examensarbete utvecklades därför ett analysverktyg i bildanalysprogrammet Segment. Verktyget använder bilder av hjärtats form tillsammans med mätningar av blodets rörelse för att placera ut virtuella partiklar, följa deras väg och klassificera dem efter hur de rör sig genom vänster kammare. Metoden testades först i virtuella konstgjorda flödesmodeller där rätt resultat var känt i förväg. Därefter användes den på totalt 60 tidigare insamlade MR-undersökningar från friska vuxna, patienter med hjärtsvikt och friska barn.

Den tydligaste skillnaden sågs hos patienterna med mest nedsatt pumpförmåga. I denna grupp var andelen blod som passerade direkt genom vänster kammare betydligt mindre, medan en större andel blod blev kvar längre i kammaren. Detta överensstämmer med tidigare forskning om förändrat blodflöde vid hjärtsvikt. Arbetet visar hur 4D-flödes-MR kan användas för att beskriva hjärtats funktion på ett mer detaljerat sätt än genom att endast mäta hur mycket blod som pumpas ut.

Contents

1 List of Abbreviations

2CH Two-chamber

4CH Four-chamber

AV Atrioventricular

DE Delayed Ejection Flow

EDV End-diastolic volume

EF Ejection fraction

ESV End-systolic volume

HF Heart failure

HFmrEF Heart failure with mildly reduced ejection fraction

HFpEF Heart failure with preserved ejection fraction

HFrEF Heart failure with reduced ejection fraction

KE Kinetic energy

LAX Long-axis

LV Left ventricle / left-ventricular

LVEF Left-ventricular ejection fraction

MAE Mean absolute error

MRI Magnetic resonance imaging

PC-MRI Phase-contrast magnetic resonance imaging

RI Retained Inflow

RK4 Fourth-order Runge–Kutta

RV Right ventricle / right-ventricular

SAX Short-axis

VENC Velocity encoding

2 Introduction

2.1 Motivation

The left ventricle (LV) is often evaluated using measures such as ventricular volumes and ejection fraction. These measures are clinically useful, but they do not describe how blood is transported within the ventricle during the cardiac cycle. Two subjects can therefore have similar global ventricular function but different intraventricular flow patterns. This matters because filling, pre-systolic blood organization, and ejection are not only related to changes in shape or volume, but also to how blood moves through the ventricle [1, 2, 3].

Four-dimensional flow magnetic resonance imaging (4D flow MRI) provides time-resolved three-dimensional velocity data and makes it possible to study how blood moves through the heart over the full cardiac cycle [4, 5, 6, 7]. With this technique, blood transport can be followed from inflow through the mitral valve, through the LV cavity, and out through the aortic valve. This makes 4D flow MRI relevant both for understanding normal ventricular physiology and for studying altered flow in disease [7, 1].

Several clinical and research applications have been discussed in the literature, including heart failure, ventricular remodelling, and diastolic dysfunction [2, 1, 8]. 4D flow MRI is also relevant in complex cardiovascular anatomy, including congenital heart disease, where three-dimensional flow patterns may be difficult to assess with conventional two-dimensional measurements alone [7].

A strength of 4D flow MRI is that it describes motion in space and time rather than only anatomy. It can therefore be used to study flow patterns, transport pathways, and energetic behaviour inside the ventricle in a way that conventional global measures cannot. This makes 4D flow MRI useful for studying how anatomy, function, and blood flow are connected [7, 1].

However, visualizations alone are not sufficient for comparisons between subjects, patient groups, or processing methods. Although streamlines and pathlines provide an intuitive representation of intracardiac flow in an individual dataset, they do not by themselves provide a standardized measure that can be compared across analyses. Quantitative research therefore requires clearly defined and reproducible methods that translate the visualized flow patterns into numerical results. This is particularly important when studying differences between healthy and diseased hearts or when evaluating the influence of different processing approaches [6, 9].

One established way to summarize blood flow through the LV is flow-component analysis. In this approach, the LV blood volume is divided into physiologically meaningful groups depending on whether blood enters and/or leaves the ventricle during the analysed cardiac cycle [10, 11, 12]. This transforms a large set of particle trajectories into a smaller number of interpretable components, such as Direct Flow, Retained Inflow, Delayed Ejection Flow, and Residual Volume. The concept is introduced in more detail in Section 3.8 and illustrated in Figure 3.5.

For flow-component analysis to be useful, the workflow must be clear and repeatable. Important choices include how the LV seed volume is defined, how pathlines are integrated, how the mitral and aortic analysis planes are placed, and how plane crossings are interpreted. If these steps are not handled consistently, the final component percentages can become too dependent on local user choices or implementation details. This is also emphasized in the literature about 4D flow, where accurate analysis tools, transparent post-processing, and quality control are described as necessary for broader research and clinical use [6, 7, 9].

This thesis addresses that need by implementing a reproducible LV flow-component workflow in Segment [13]. The aim of this thesis is to develop and evaluate a research workflow in Segment for 4D flow MRI-based flow-component analysis. The workflow is intended for research use. It defines the main analysis steps, provides visual checks, and produces quantitative flow-component results for evaluation in phantom experiments and human datasets.

The work was evaluated in three steps. First, the implementation shows that the flow-component workflow can be built and used inside Segment. Second, the phantom experiments test important technical parts of the method in controlled cases where the expected result is known. These tests mainly evaluate the pathline integration and the plane-crossing classification. Third, the human datasets are used to show that the workflow can be applied to real 4D flow MRI data and that the results give plausible group-level trends. The phantom experiments should therefore be seen as technical validation of selected parts of the workflow, while the human-data analysis should be seen as a feasibility evaluation rather than full clinical validation.

3 Background

3.1 Cardiac anatomy and left-ventricular blood transport

The heart consists of four chambers: two atria and two ventricles, arranged as a right and a left side. The right side receives deoxygenated blood from the body and pumps it to the lungs, while the left side receives oxygenated blood from the lungs and pumps it into the systemic circulation [14]. On the left side, oxygenated blood enters the left atrium, passes through the mitral valve into the left ventricle (LV), and is later ejected through the aortic valve into the aorta [14]. The cardiac valves are shown schematically in Figure 3.1. In this thesis, the mitral and aortic valves are important anatomical landmarks, since they define the main inflow and outflow routes used in the LV flow-component analysis [10, 11].

The cardiac cycle is commonly divided into two main phases: diastole, when the heart muscle relaxes and the chambers fill with blood, and systole, when the heart muscle contracts and blood is ejected [14]. During diastole, the LV fills through the mitral valve. During systole, the LV contracts and ejects blood through the aortic valve into the aorta. These valve-defined transport routes are central to LV flow-component analysis [10, 11].

Earlier studies have shown that blood transport within the LV is not random, but organized over time in physiologically meaningful ways [10, 12]. This makes the LV well suited for time-resolved flow analysis, where the aim is not only to describe chamber shape, but also to follow how blood enters, moves within, and leaves the ventricle during the cardiac cycle. LV flow patterns are relevant both for understanding normal ventricular function and for studying how blood transport may change when cardiac function is altered, for example in ventricular remodelling, heart failure, or congenital heart disease [3, 2, 7].

An important concept behind LV flow-component analysis is that some blood is already favorably positioned for ejection before systole, while other blood remains in the ventricle for a longer time [12]. This pre-systolic organization of blood flow provides a physiological motivation for later dividing LV blood transport into separate flow components.

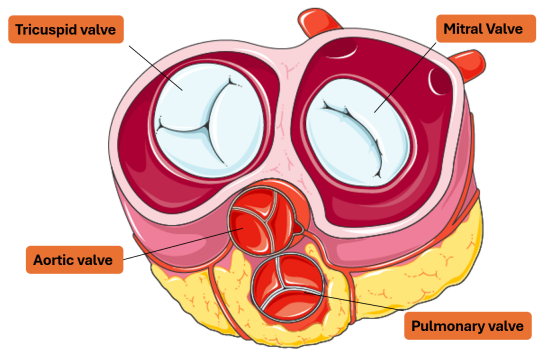

Figure 3.1: Schematic illustration of the cardiac valves. For the present work, the most important structures are the mitral valve and the aortic valve, since they define left-ventricular inflow and outflow. Adapted from Servier Medical Art [15], licensed under CC BY 4.0.

3.2 Heart failure and ejection fraction

Heart failure (HF) is a clinical syndrome in which structural or functional cardiac abnormalities impair the ability of the heart to fill or eject blood effectively. It is commonly associated with symptoms such as breathlessness, fatigue, and reduced exercise capacity [16]. In cardiovascular imaging, left-ventricular ejection fraction (LVEF) is one of the most widely used measures for describing global systolic function.

LVEF describes the fraction of the LV end-diastolic volume that is ejected during systole,

$$\mathrm{LVEF} = \frac{\mathrm{EDV} - \mathrm{ESV}}{\mathrm{EDV}}, \tag{3.1}$$

where EDV is the end-diastolic volume and ESV is the end-systolic volume. A lower LVEF generally indicates reduced systolic pump function, but LVEF alone does not describe how blood is transported inside the ventricle [16, 1].

Current guidelines commonly divide HF into EF-based categories. Heart failure with reduced ejection fraction (HFrEF) refers to patients with LVEF $\leq$ 40%, heart failure with mildly reduced ejection fraction (HFmrEF) refers to LVEF 41–49%, and heart failure with preserved ejection fraction (HFpEF) refers to LVEF $\geq$ 50% [16]. In clinical diagnosis, HFmrEF and HFpEF also require evidence of increased filling pressures or other supporting clinical findings, so EF should not be interpreted as the only diagnostic criterion [16].

These EF groups are relevant because they provide a simple way to relate LV flow-component results to ventricular function. LVEF describes how much blood is ejected from the LV during systole, but it does not describe how blood moves inside the ventricle. Flow-component analysis can provide additional information about whether the LV blood volume is mainly transported efficiently through the chamber or remains in the ventricle for a longer time.

In this thesis, the EF groups are used to compare flow-component results with patterns reported in previous literature. In particular, they make it possible to compare whether the results follow the same general trend as reported by Wong et al., where lower EF was associated with reduced Direct Flow and increased Residual Volume [17].

3.3 Basic principles of MRI

Magnetic resonance imaging (MRI) is a non-invasive imaging technique that can produce detailed images of internal anatomy without using ionizing radiation [18, 19]. MRI is based on the behaviour of hydrogen nuclei in a strong magnetic field. Since the human body contains a large amount of water, and water contains hydrogen, this makes it possible to obtain signals from many tissues in the body.

When a patient is placed in the scanner, the hydrogen nuclei tend to align with the main magnetic field. A radiofrequency pulse is then applied, which temporarily disturbs this alignment. As the nuclei return toward their original state, they emit signals that can be measured by the scanner [18]. Different tissues produce different signal behaviour, which makes it possible to create image contrast. Spatial information is added by magnetic field gradients, which allow the measured signals to be linked to specific positions in the body [18, 19].

In cardiovascular imaging, MRI is useful because it can provide both anatomical and functional information. Standard cine MRI can show cardiac motion and chamber geometry over time, which makes it valuable for assessing ventricular shape and function. Phase-contrast MRI extends this by using the MR signal to encode blood velocity.

3.4 Cine MRI and LV segmentation

Cine MRI is commonly used in cardiac imaging to show the motion of the heart throughout the cardiac cycle. Instead of producing a single static image, cine MRI gives a time-resolved image series that shows how the ventricles fill and contract. This makes it useful for assessing ventricular volumes, wall motion, and ejection fraction [20].

For LV analysis, cine short-axis images are often used because they cover the ventricle from base to apex. By outlining the LV endocardial border in each slice, the LV blood volume can be described over time. These contours can then be used to calculate LV volumes and to define the blood volume used for later particle placement [20, 13].

Long-axis cine images, such as two-chamber and four-chamber views, add information about the shape and orientation of the ventricle. These views are especially important because they show the basal part of the LV and the mitral valve region. This information is used when constructing the mitral analysis plane.

Figure 3.2 shows examples of the cine MRI views used in the workflow. The short-axis view is used to define the LV blood volume. The two-chamber and four-chamber views show the

ventricle in long-axis directions.

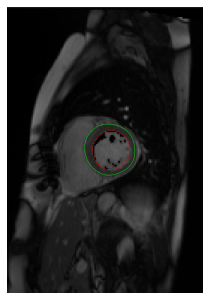

(a) Cine short-axis view.

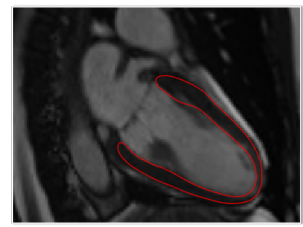

(b) Cine two-chamber view.

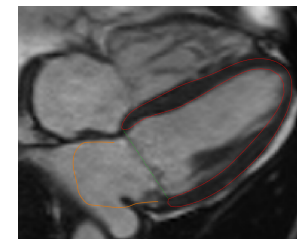

(c) Cine four-chamber view.

Figure 3.2: Examples of cine MRI views used in the workflow. The short-axis view is used to define the left-ventricular blood volume. The two-chamber and four-chamber views provide long-axis information about the ventricle and the mitral valve region. In the long-axis views, the red segmentation marks the LV myocardium. In the four-chamber view, the orange segmentation marks the left atrium and the green line marks the atrioventricular plane.

Using cine-derived segmentation also introduces a practical challenge. Cine images and 4D flow MRI are normally acquired as separate image series and may have different spatial resolution, slice positions, and timeframes. Information from the cine images must be transferred carefully to the 4D flow coordinate system before it can be used for particle placement or analysis-plane construction.

3.5 Phase-contrast MRI and 4D flow MRI

An important extension of conventional MRI is phase-contrast MRI, where the MR signal is used not only to create anatomical images, but also to measure motion. In phase-contrast MRI, velocity is encoded in the phase of the MR signal by applying bipolar gradient pulses. Stationary spins experience little or no net phase shift, while moving spins accumulate a phase shift related to their velocity along the encoding direction [21]. This makes it possible to estimate both the speed and direction of blood flow from the measured phase data.

An important acquisition parameter is the velocity encoding value, VENC. VENC defines the maximum velocity that can be measured without phase wrapping. If VENC is set too low, velocities above this range may appear incorrectly because of aliasing. If it is set too high, sensitivity to lower velocities is reduced [21]. Choosing an appropriate VENC is important for reliable flow quantification.

Four-dimensional flow MRI (4D flow MRI) extends phase-contrast MRI by measuring all three velocity components in a three-dimensional volume over time [5, 6]. This gives a time-resolved velocity field over the cardiac cycle, rather than flow information only in a single plane. This is useful in the heart, where blood changes direction over time and flow

may need to be evaluated in different regions.

A typical 4D flow MRI acquisition covers a three-dimensional volume and is acquired over several cardiac cycles. The acquisition is synchronized to the cardiac cycle, usually using ECG gating, so that the data can be reconstructed into timeframes representing different phases of the heartbeat [5]. Respiratory motion may also need to be managed, for example using respiratory gating or navigator techniques [5, 6]. In the later pathline analysis, the reconstructed cardiac cycle is treated as periodic, so the temporal sequence can wrap from the last timeframe back to the first.

After acquisition, the data are reconstructed into magnitude and velocity images. Before quantitative analysis, the velocity data may need correction for phase offsets, velocity aliasing, and other acquisition-related artifacts [6]. This is important because pathline visualization and flow-component analysis depend directly on the measured velocity field.

Compared with conventional two-dimensional flow imaging, 4D flow MRI gives more flexibility during analysis. Since velocity is measured in a full volume, analysis planes can be placed retrospectively during post-processing [5, 6]. This is useful in intracardiac applications, where blood transport is complex and the final analysis planes may not be fully known before acquisition [7].

Figure 3.3 shows an example of a phase-contrast MR image.

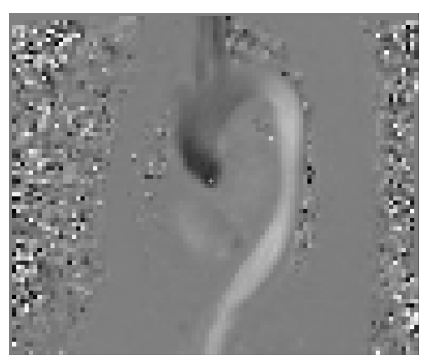

Figure 3.3: Example of a phase-contrast MR image with velocity encoding in the foot-head direction. In phase-contrast MRI, the pixel intensity in the phase image reflects velocity along the chosen encoding direction, which makes it possible to distinguish flow in different directions.

3.6 Segment as a platform for cardiovascular image analysis

Segment is a software platform developed for cardiovascular image analysis. It provides tools for visualization and quantitative analysis of cardiovascular MRI data, including cardiac segmentation and flow-related measurements [13].

Segment is relevant for this thesis because the workflow combines several types of data. Cine

MRI is used for ventricular geometry and segmentation, while 4D flow MRI provides the velocity field used for particle tracing. Working in the same software environment makes it possible to inspect the image data, segmentations, analysis planes, pathlines, and final flow-component results together.

The workflow developed in this thesis was implemented within Segment. It adds a structured procedure for ventricular flow-component analysis from 4D flow MRI, including particle placement, pathline tracing, plane-crossing based classification, visualization, and quantitative output.

3.7 Flow visualization in 4D flow MRI: streamlines and pathlines

Several methods can be used to visualize flow data from 4D flow MRI, including vector fields, streamlines, and pathlines [4, 7]. Streamlines and pathlines are often shown in a similar way, but they describe different things.

A streamline shows the instantaneous flow direction at one time point. It can be seen as a snapshot of the velocity field. A pathline follows a virtual massless particle through the time-varying velocity field [4, 9]. Pathlines show transport over time.

This difference is important in this thesis. LV flow-component analysis is based on whether blood enters the LV, moves within it, and later leaves it again. This requires particle tracing over time, so pathlines are used for the classification. Streamlines remain useful for visual inspection, but they cannot by themselves determine whether a particle entered or left the LV during the analysed cardiac cycle.

The visual difference between streamlines and pathlines is illustrated in Figure 3.4, and their roles in this thesis are summarized in Table 3.1.

Pathline-based analysis also depends on numerical choices such as interpolation, integration settings, temporal resolution, seeding strategy, and boundary handling [9, 6]. These choices need to be clearly defined when pathlines are used for quantitative analysis.

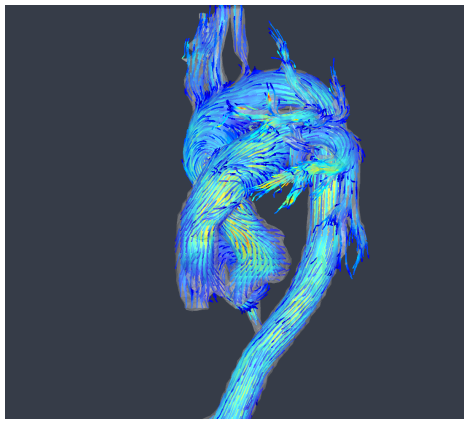

(a) Example of streamlines in Segment.

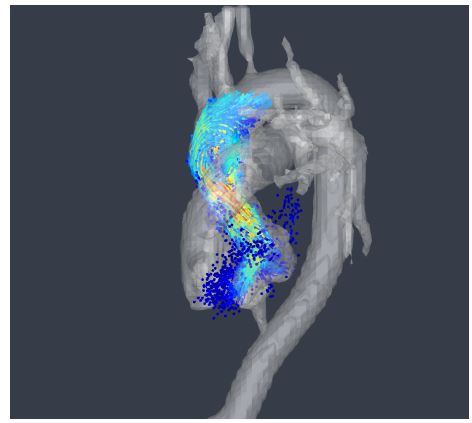

(b) Example of LV pathlines in Segment.

Figure 3.4: Examples of flow visualization in Segment. Streamlines show the instantaneous flow direction at a given time point, whereas pathlines show how virtual particles move through the velocity field over time. In this thesis, pathlines are used because the analysis focuses on blood transport through the left ventricle across the cardiac cycle.

Table 3.1: Comparison between streamlines and pathlines in the context of LV flow analysis.

Method	**What it shows**	**Role in this thesis**
Streamlines	Instantaneous flow direction at a single time point.	Useful for qualitative visualization of flow organization, but not enough for transport-based LV flow-component classification.
Pathlines	Transport of a virtual particle through the time-varying velocity field.	Forms the basis of the implemented workflow, since the analysis depends on tracking blood transport through the LV across the cardiac cycle.

3.8 Left-ventricular flow-component analysis

A commonly used way to summarize left-ventricular blood transport in 4D flow MRI is to divide the LV end-diastolic blood volume into four functional flow components, based on whether the blood enters and/or leaves the ventricle during the analysed cardiac cycle [10, 11]. In this framework, *Direct Flow* refers to blood that enters the LV during diastole and is ejected during the following systole. *Retained Inflow* enters during the cycle but remains in the ventricle at the end of the cycle. *Delayed Ejection Flow* is already present in the ventricle at the beginning of the cycle and leaves during that cycle. *Residual Volume* remains in the LV for at least two consecutive cycles. These four components are visualized in Figure 3.5 and summarized in Table 3.2.

In the present thesis, the same conceptual definition is used, but the classification is implemented as a rule-based particle-level test. Each particle is evaluated using backward tracing for mitral inflow and forward tracing for aortic outflow. The detailed implementation is described in Chapter 4.

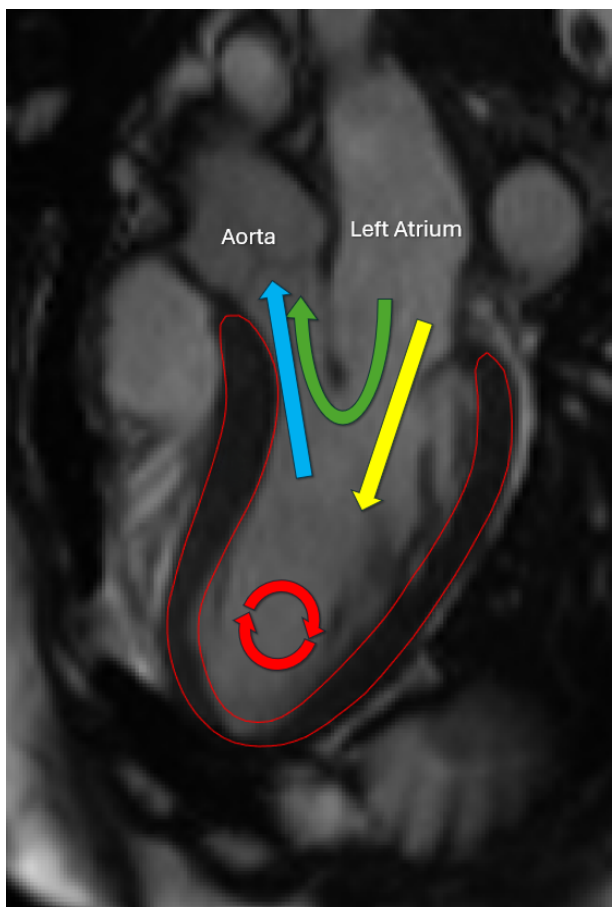

Figure 3.5: Representative visualization of the four left-ventricular flow components in a three-chamber view. Direct Flow is shown in green, Retained Inflow in yellow, Delayed Ejection Flow in blue, and Residual Volume in red. The arrows illustrate how the same LV blood volume can be separated into components depending on whether the blood enters through the mitral valve and/or leaves through the aortic valve during the analysed cardiac cycle.

Table 3.2: Definition of the four left-ventricular flow components used in this thesis.

Component	**Mitral inflow during cycle**	**Aortic outflow during cycle**	**Physiological interpretation**
Direct Flow	Yes	Yes	Blood that enters the LV during diastole and is ejected during the following systole.
Retained Inflow	Yes	No	Blood that enters the LV during the analysed cycle but remains in the ventricle at the end of the cycle.
Delayed Ejection Flow	No	Yes	Blood already present in the LV at the beginning of the cycle that is ejected during that cycle.
Residual Volume	No	No	Blood that remains in the LV for at least two consecutive cycles.

The four components describe different transport patterns. Direct Flow represents blood that enters and leaves the LV within the same cardiac cycle. Residual Volume represents blood that remains in the LV for at least two cycles. Retained Inflow and Delayed Ejection Flow describe intermediate patterns, where blood either enters but is not ejected during the same cycle, or is already present in the LV and leaves during that cycle.

This framework is useful because it turns complex particle trajectories into a smaller set of physiological categories [10, 1]. It also makes it easier to compare subjects and relate flow patterns to ventricular function.

Earlier studies in healthy subjects showed that Direct Flow is of particular interest. It has been described as a large part of the ejected blood volume and as being favorably positioned relative to the left-ventricular outflow tract before systole [10, 12]. This supports the idea that normal diastolic flow organization helps prepare part of the LV blood volume for ejection.

In healthy adults, Direct Flow is typically the largest left-ventricular flow component, accounting for roughly 35–40% of end-diastolic volume. Residual Volume is often about 30%, while Retained Inflow and Delayed Ejection Flow each usually contribute around 15–20% [11, 12, 22]. These approximate proportions are summarized in Table 3.3. The values should be interpreted as typical ranges, not diagnostic thresholds.

Table 3.3: Approximate distribution of LV flow components in healthy adults, based on values reported in previous 4D flow MRI studies [11, 12, 22]. The values should be interpreted as typical ranges rather than diagnostic thresholds.

Component	**Typical fraction of LV end-diastolic volume**
Direct Flow	35–40%
Retained Inflow	15–20%
Delayed Ejection Flow	15–20%
Residual Volume	approximately 30%

In heart failure, conventional measures such as LVEF describe global pump function, but not how blood is transported inside the LV. Wong et al. studied 4D-flow-derived LV flow components across the heart-failure spectrum and reported lower Direct Flow and higher Residual Volume in groups with lower EF [17]. Direct Flow also showed a positive relationship with LVEF, while Residual Volume showed a negative relationship with LVEF [17]. These findings support the use of LV flow-component analysis as a research marker of altered intraventricular blood transport.

3.9 Left-ventricular kinetic energy from 4D flow MRI

In addition to flow-component volumes, 4D flow MRI can be used to estimate the kinetic energy of the moving blood inside the LV. Kinetic energy describes the mechanical energy associated with motion. In this context, it is calculated from the blood velocity field and can be followed over the cardiac cycle. Earlier work from Lund has used 4D flow MRI to study LV kinetic-energy time curves in heart failure, and KE has also been described as part of the analysis of presystolic LV flow organization [23, 12]. A systematic review has also summarized the clinical relevance of LV blood-flow KE assessment by 4D flow CMR [24].

The basic calculation is performed voxel by voxel inside the LV. For one voxel, the kinetic energy can be written as

$$KE_{\text{voxel}} = \frac{1}{2}\rho_{\text{blood}} V_{\text{voxel}} v_{\text{voxel}}^2, \tag{3.2}$$

where ρ_{blood} is the density of blood, V_{voxel} is the voxel volume, and v_{voxel} is the velocity magnitude in that voxel. A commonly used blood density is 1.06 g/cm^3 [25]. The velocity magnitude is calculated from the three measured velocity components,

$$v_{\text{voxel}} = \sqrt{v_x^2 + v_y^2 + v_z^2}. \tag{3.3}$$

The total LV kinetic energy at one timeframe is then obtained by summing the voxel kinetic energy over all voxels inside the LV,

$$KE_{\text{LV}}(t) = \sum_{\mathbf{x}\in\text{LV}} KE_{\text{voxel}}(\mathbf{x}, t). \tag{3.4}$$

Repeating this calculation for all timeframes gives a time-resolved LV kinetic-energy curve.

In this thesis, LV kinetic energy is used as an additional output in the implemented workflow. It is not used to classify the four LV flow components. Instead, it provides a complementary measure of the velocity field inside the LV and can help the user inspect the overall time-resolved flow behaviour together with the component percentages.

3.10 Need for reproducible LV flow analysis workflows

Although 4D flow MRI provides a large amount of information about intracardiac blood transport, the final analysis depends on several processing and modelling choices. In practice, the result is dependent not only on the acquired data, but also on how the region of interest is defined, how pathlines are seeded and integrated, how analysis planes are positioned, and how classification rules are applied [9, 6].

This is important in LV flow-component analysis, where the aim is to divide blood transport into well-defined categories over the cardiac cycle. Small differences in implementation may affect whether a particle is considered to have entered or left the ventricle, and such differences can influence the resulting component volumes. Quantitative pathline-based analysis requires clear and reproducible workflow definitions rather than only qualitative visualization [11, 9].

Earlier 4D flow MRI literature has emphasized the importance of careful post-processing, quality control, and transparent analysis choices when quantitative measures are reported [6]. In a research setting, this is important both for internal consistency and for making results easier to review, compare, and reproduce across datasets and studies.

This need motivates the present thesis. Rather than proposing a stand-alone clinical decision tool, the aim is to define and implement a clear workflow for LV flow-component analysis that can be executed and evaluated in a consistent way.

4 Methods

4.1 Overview of the implemented workflow

A semi-automated workflow for left-ventricular (LV) flow-component analysis was implemented in Segment. The workflow combines cine-derived LV geometry with time-resolved three-dimensional velocity data from 4D flow MRI and produces particle-based LV flow-component fractions.

Three image inputs are used. Cine short-axis (SAX) data are used for LV seed generation, cine two-chamber (2CH) and four-chamber (4CH) data are used for analysis-plane construction, and the 4D flow MRI dataset provides the time-resolved velocity field used for pathline integration. The overall workflow is shown in Figure 4.1.

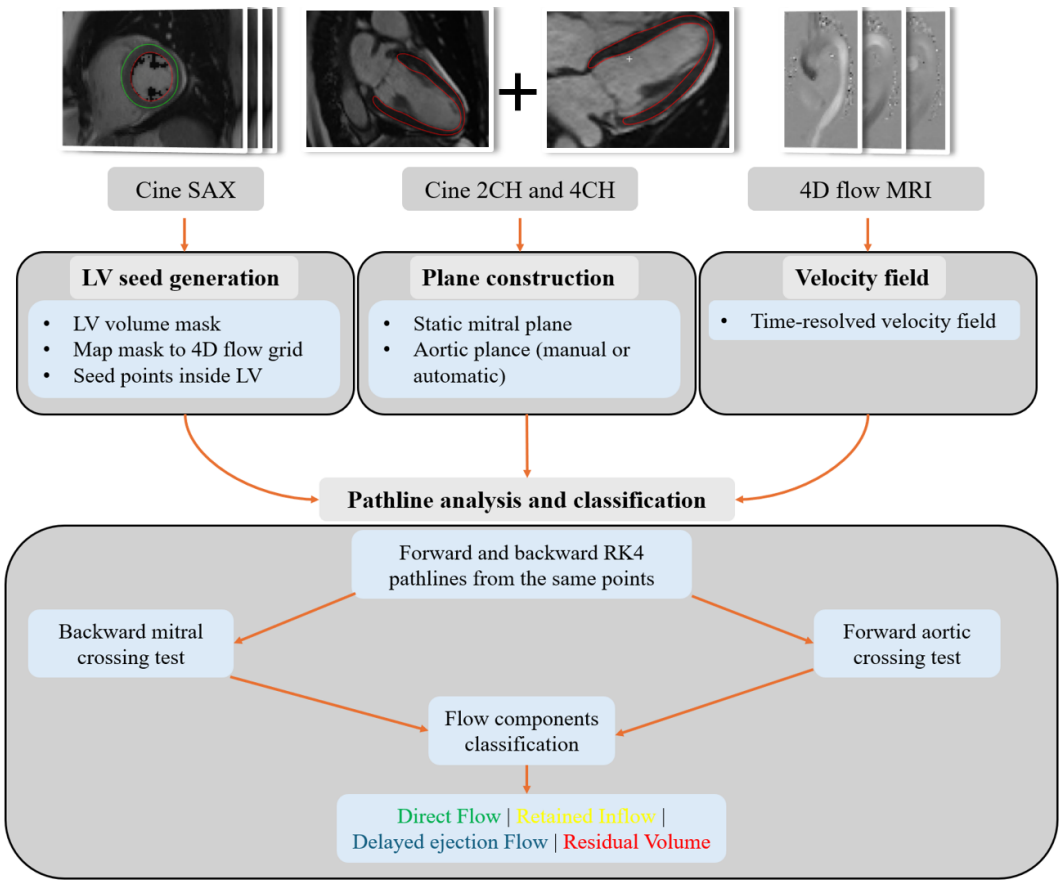

Figure 4.1: Overview of the implemented LV flow-component workflow. Cine SAX data are used to define the LV seed region, cine 2CH and 4CH data are used to construct the analysis planes, and 4D flow MRI provides the time-resolved velocity field for pathline integration. Forward and backward RK4 pathlines are generated from the same seed points. Classification is then based on backward mitral crossing and forward aortic crossing. Final results are reported as flow-component fractions of the seeded particles used for classification.

The analysis starts from the LV end-diastolic volume, just before systole. First, an LV volume mask is generated from the cine SAX segmentation and mapped to the 4D flow grid. The mapped seed mask is then reduced before particle placement so that seeds are placed inside the LV cavity and not in the most basal part of the transferred mask. Seed points are placed inside the remaining LV volume, and the total number of seeds is limited to keep the workflow practical.

From each seed point, one forward and one backward pathline are computed in the 4D flow velocity field using fourth-order Runge–Kutta (RK4) integration. The two pathlines always start from the same initial seed position. This makes it possible to combine backward information about where the blood came from with forward information about where it goes after the seed time.

Two analysis planes are used in the classification step. A mitral plane is constructed from cine 2CH and 4CH data, or an existing user-defined plane can be used. The aortic plane can either be placed manually or estimated automatically in the workflow. The pathlines are then tested for crossings of these planes over their traced lifetime.

Finally, each seeded particle is assigned to one of the four standard LV flow components: Direct Flow, Retained Inflow, Delayed Ejection Flow, or Residual Volume.

In Segment, the main analysis steps were connected through a callback-based workflow. This callback links LV seed generation, mitral- and aortic-plane handling, forward and backward pathline generation, plane-crossing detection, and flow-component classification. The resulting user-facing workflow and display options are presented in Section 5.4. The following sections describe the method steps in more detail.

The main implementation settings used in the workflow are summarized in Table 4.1.

Most settings in Table 4.1 were chosen during development to make the workflow stable and practical to run. The number of seeds and the RK4 time step were mainly chosen to balance computation time and pathline accuracy. The plane shifts, plane size, and seed-mask reduction were chosen after visual testing, to reduce errors in seed placement and plane crossings. These values should therefore be seen as practical implementation choices. Settings related to seed placement, plane placement, and crossing detection are expected to have the largest effect on the final flow-component percentages.

Table 4.1: Main implementation settings used in the LV flow-component workflow.

Setting	**Value used**	**Purpose**
Seed timeframe	End-diastole + 1 frame	Places the initial particles in the pre-systolic LV blood volume.
Basal seed-mask reduction	2 most basal occupied slices	Reduces the risk of seeding above the LV cavity after transfer from cine to 4D flow space.
Maximum analysis seeds	5000	Keeps pathline calculation practical on ordinary workstations.
Desired RK4 time step	5 ms	Gives a balance between pathline accuracy and computational cost.
Mitral-plane shift	2 mm	Moves the automatically created mitral plane slightly above the LV to reduce false crossings.
Aortic-plane shift	2 mm downstream	Places the automatically estimated aortic plane slightly in the local outflow direction.
Automatic aortic-plane size	60 mm	Gives a fixed plane size large enough to cover the detected outflow region.
Crossing criterion	At least one valid crossing	A particle is marked as crossing if its pathline intersects the full plane polygon at least once.
Reported output	Percent of particles	Makes the result independent of the exact number of seed particles used.
Seed-mask border erosion	0 voxels by default; 2 voxels for paediatric datasets	Moves seed points away from the LV boundary in each slice to reduce early pathline escape near the wall.

4.2 Input data, timeframe mapping, and coordinate handling

The implemented workflow combines information from cine MRI and 4D flow MRI. Cine data are used to define the LV seed region and to construct the analysis planes, while the 4D flow dataset provides the time-resolved three-dimensional velocity field used for pathline integration. Since these image sources have different spatial grids and do not necessarily contain the same number of cardiac frames, careful handling of both timeframes and coordinates is required.

Three image sources are used in the workflow. Cine short-axis (SAX) data are used for LV seed generation. Cine 2CH and 4CH long-axis data are used for mitral-plane construction. The 4D flow MRI dataset is used for pathline integration and for the final flow-component classification.

The cine and 4D flow datasets do not necessarily contain the same number of cardiac frames. Therefore, when a cine-defined timeframe had to be transferred to the 4D flow dataset, the mapping was based on relative position within the cardiac cycle. First, the cine timeframe was expressed as a fraction of the full cine cycle,

$$f = \frac{t_{\text{cine}} - 1}{T_{\text{cine}} - 1}, \tag{4.1}$$

where t_{cine} is the cine frame index and T_{cine} is the total number of cine frames. The same fraction was then applied to the 4D flow cycle, which gave the mapped 4D flow frame

$$t^*_{\text{flow}} = 1 + f\,(T_{\text{flow}} - 1) = 1 + (t_{\text{cine}} - 1)\frac{T_{\text{flow}} - 1}{T_{\text{cine}} - 1}, \tag{4.2}$$

where T_{flow} is the total number of 4D flow frames and t^*_{flow} is the mapped frame index before rounding. This mapping preserves the relative position in the cardiac cycle: the first cine frame maps to the first 4D flow frame, the last cine frame maps to the last 4D flow frame, and the frames in between are placed proportionally spaced. The mapped value was then rounded to the nearest valid frame.

Spatially, contours and landmark points were transferred between cine stacks and the 4D flow dataset through a common anatomical coordinate system. A point in one image stack is initially defined only by its voxel location within that stack. However, the same voxel indices cannot be used directly in another dataset, because different image stacks may have different voxel sizes, slice spacing, positions, and orientations.

To handle this, each point was first converted from image-grid coordinates to its anatomical position in space using the geometry information stored for that stack. The same anatomical position was then converted to image-grid coordinates in the target stack. In this way, cine-derived contours and landmark points could be used in the 4D flow analysis space even though the original image stacks were not sampled on the same grid.

The procedure can be viewed as a two-step method: first from image coordinates to anatomical space, and then from anatomical space to the target image grid. This was necessary because the cine and 4D flow datasets did not share the same image grid, even when they represented the same anatomy.

For the LV seed region, this coordinate transfer was used to map the cine SAX LV mask to the 4D flow grid. For analysis-plane construction, it was used to transfer long-axis contour information to the 4D flow analysis environment. This coordinate handling step was necessary because the final workflow combines geometry from cine MRI with transport information from 4D flow MRI in one common analysis space.

4.3 LV seed generation

Seeding was based on the LV end-diastolic volume, just before systole. This step defines the initial particle region used for both forward and backward pathline integration.

4.3.1 LV mask construction and transfer to 4D flow space

The seed region was first defined in the cine short-axis stack. At the selected cine timeframe, each short-axis slice with a valid LV endocardial contour was converted to a binary mask representing the LV cavity in that slice. Pixels inside the endocardial contour were included in the mask, while pixels outside the contour were excluded.

If the LV cavity mask is denoted by M_{LV}, this can be written as

$$M_{\mathrm{LV}}(x, y, z) = \begin{cases} 1, & \text{if } (x, y) \text{ lies inside the LV endocardial contour on slice } z, \\ 0, & \text{otherwise.} \end{cases} \tag{4.3}$$

If LV endocardial contours were missing at the required cine timeframe, existing LV segmentation tools in Segment were used before the mask was created. In this way, the workflow used the existing LV segmentation tools in Segment, while the mask construction and the following seed-generation steps were part of the implemented method.

The LV cavity mask was first defined in the cine short-axis stack, but the later pathline analysis was performed in the 4D flow dataset. Therefore, the mask had to be transferred from cine space to the 4D flow grid.

For each voxel in the 4D flow grid, the corresponding anatomical location in the cine short-axis stack was identified and sampled from the cine LV mask. This transferred the cine-defined LV cavity region to the 4D flow grid.

Since the cine mask was binary, nearest-neighbor resampling was used so that the mapped mask remained binary and did not get artificial intermediate values at the boundary.

The result of this step was a binary LV cavity mask defined directly on the 4D flow grid. This mapped mask was then used in the following mask-reduction and seed-placement steps. Figure 4.2 shows an example of the cine-based LV definition and how the LV region and seed points were visualized in the 4D flow module.

4.3.2 Mask reduction and seed placement

The mapped LV mask was reduced before seed placement. This was mainly done to avoid placing seeds in the most basal part of the transferred mask, where small differences between cine and 4D flow data could otherwise place seed points close to the left atrium or the valve plane region.

In the current implementation, the two most basal occupied slices of the mapped LV mask were excluded before seed placement. The basal direction was determined from the anatomical foot-head direction of the 4D flow grid. This reduction was used because the cine images and the 4D flow MRI data are acquired separately, and small patient motion or registration differences can make the transferred basal LV boundary uncertain.

An optional border erosion was also implemented. This erosion was applied after the cine-derived LV mask had been transferred to the 4D flow grid. In the implementation, the outer border voxels of the mapped LV mask were removed slice by slice. The default setting was no border erosion. For the paediatric datasets, a two-voxel in-plane erosion was applied before seed placement.

This setting was used because visual review showed that, without erosion, seeds close to the LV boundary produced pathlines that left the ventricle early in the paediatric datasets. The two-voxel erosion moved the seed region further inward in each slice and reduced this problem in the smaller hearts.

The remaining mask was then used as the final seed region. One seed point was placed at the center of each included voxel. If the number of candidate seed points was larger than 5000, the total number of analysis seeds was limited to 5000 in order to keep the workflow practical.

The seed limitation was performed by randomly selecting a subset of the candidate seed points using a fixed random seed. This makes the selection reproducible between runs while still sampling the full candidate seed set without manually favoring a specific LV region. These seed points were then used as the common starting positions for both forward and backward pathline integration.

4.3.3 LV contour object for visualization and quality control

In addition to the seed mask itself, an LV contour object was implemented in the 4D flow module. The object was built from cine short-axis endocardial contours and displayed as an LV surface in the same environment as the seeds, pathlines, and analysis planes.

The main purpose of this object was visual and practical. It provided a clear representation of the LV geometry in the 4D flow analysis space and made it easier to assess whether the seed region and later analysis objects were placed in a reasonable way.

The LV contour object was therefore not used directly for the flow-component classification itself. Instead, it was included as part of the implementation to support visualization and quality control during development, testing, and analysis.

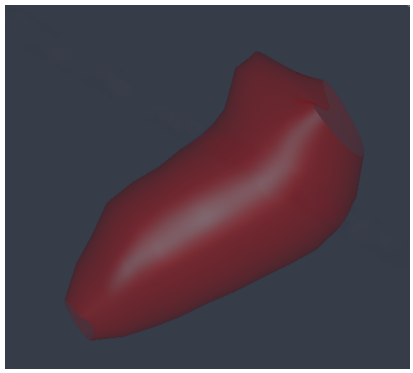

(a) LV region visualized in the 4D flow module with the implemented LV contour object.

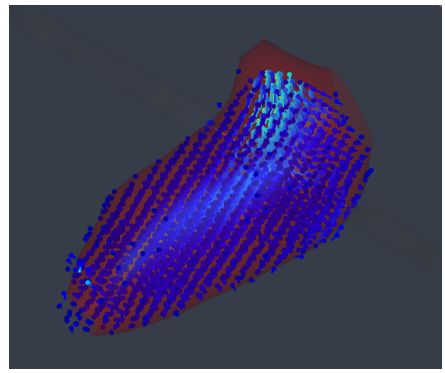

(b) Example of seed points placed inside the LV cavity before pathline integration.

Figure 4.2: Example of the implemented LV seed-generation workflow. The LV cavity was first defined from cine short-axis segmentation. The corresponding LV region was then visualized in the 4D flow analysis environment, and seed points were placed inside the final LV seed region before forward and backward pathline integration.

4.4 Time-resolved pathline integration

4.4.1 Pathlines

As introduced in Section 3.7, pathlines describe how particles move through the time-resolved velocity field. In this workflow, they are used to follow blood transport from the LV seed region.

For each seed $\mathbf{x}_0$, particle transport was modeled by the ordinary differential equation

$$\frac{d\mathbf{x}(t)}{dt} = \tilde{\mathbf{v}}(\mathbf{x}(t), t), \tag{4.4}$$

where $\mathbf{x}(t)$ is the particle position and $\tilde{\mathbf{v}}(\mathbf{x}, t)$ is the time-resolved 3D velocity field expressed in image-grid coordinates. Here, $\tilde{\mathbf{v}}(\mathbf{x}, t)$ denotes the field used in the numerical integration, expressed in the target image grid and scaled to displacement per integration step.

Because the 4D flow data are available only at some voxel locations and cardiac frames, interpolation is required in both space and time. Let the discrete velocity field be written

$$\mathbf{v}_{ijkm} = \mathbf{v}(x_i, y_j, z_k, t_m), \tag{4.5}$$

with $m = 1, \ldots, T_{\text{flow}}$. A continuous approximation $\tilde{\mathbf{v}}(\mathbf{x}, t)$ was obtained by linear interpolation between neighboring samples.

4.4.2 Forward and backward tracing

For each seed, one forward and one backward trajectory were computed. This one-to-one correspondence is central to the later classification step, since each final particle class depends on combining backward inflow information with forward outflow information for the same initial seed.

To write both cases compactly, a direction parameter $\sigma \in \{+1, -1\}$ can be introduced,

$$\frac{d\mathbf{x}(t)}{dt} = \sigma\, \tilde{\mathbf{v}}(\mathbf{x}(t), t), \tag{4.6}$$

where $\sigma = +1$ corresponds to forward tracing and $\sigma = -1$ corresponds to backward tracing. In this way, the same numerical scheme can be used for both directions.

4.4.3 RK4 integration and boundary constraints

Numerical integration was performed using fourth-order Runge–Kutta (RK4). This method was chosen because it gives a good balance between accuracy and computational cost, which is important when pathlines are computed inside a practical workflow in Segment and not only in a separate offline test. Fourth-order Runge–Kutta integration has previously been used for cardiac 4D-flow particle tracing [26, 27, 9]. For example, de Koning et al. reported particle tracing using RK4 with a time step of 10 ms [27], while earlier LV flow-component work described the use of adaptive step length [11].

Let $\mathbf{x}_n$ be the particle position at integration step n, and let Δt be the integration time step. The velocity field is sampled at the current particle position and time, and RK4 then updates the particle position through four intermediate evaluations:

$$\mathbf{k}_1 = \sigma\, \tilde{\mathbf{v}}(\mathbf{x}_n, t_n), \tag{4.7}$$

$$\mathbf{k}_2 = \sigma\, \tilde{\mathbf{v}} \left(\mathbf{x}_n + \frac{\Delta t}{2}\mathbf{k}_1,\ t_n + \frac{\Delta t}{2} \right), \tag{4.8}$$

$$\mathbf{k}_3 = \sigma\, \tilde{\mathbf{v}} \left(\mathbf{x}_n + \frac{\Delta t}{2}\mathbf{k}_2,\ t_n + \frac{\Delta t}{2} \right), \tag{4.9}$$

$$\mathbf{k}_4 = \sigma\, \tilde{\mathbf{v}} \left(\mathbf{x}_n + \Delta t\, \mathbf{k}_3,\ t_n + \Delta t \right), \tag{4.10}$$

after which the new position is computed as

$$\mathbf{x}_{n+1} = \mathbf{x}_n + \frac{\Delta t}{6} \left(\mathbf{k}_1 + 2\mathbf{k}_2 + 2\mathbf{k}_3 + \mathbf{k}_4 \right). \tag{4.11}$$

In the present implementation, the time step was chosen relative to the temporal resolution of the 4D flow data instead of forcing the same fixed number of integration steps between all datasets. A desired time step of 5 ms was used, and the actual integration step was then adjusted so that each cardiac frame was divided into an integer number of equal substeps:

$$N_{\text{sub}} = \left\lceil \frac{T_{\text{frame}}}{\Delta t_{\text{des}}} \right\rceil, \qquad \Delta t = \frac{T_{\text{frame}}}{N_{\text{sub}}}, \tag{4.12}$$

where T_{frame} is the time between two measured 4D flow frames and $\Delta t_{\text{des}} = 5$ ms.

This means that the actual Δt may vary slightly between datasets, but it stays close to the desired value while still fitting exactly into each cardiac frame. This gives more even integration across datasets with different numbers of frames. At the same time, the method remains simple, reproducible, and easy to store and evaluate frame by frame.

Thus, each dataset used a fixed integration time step adapted to its temporal resolution.

During integration, particle positions were constrained to remain inside the valid 4D flow volume. This was done to avoid interpolation outside the available velocity field and to make the numerical integration robust in practical use. If a particle reached the edge of the image volume, its position was kept inside the valid grid. This was mainly a numerical boundary constraint and was not intended to change the classification logic. The traced pathlines were later used both for flow-component classification and for visual quality control.

4.5 Plane-crossing detection and flow-component classification

4.5.1 Analysis planes used for classification

The flow-component classification was based on two valve-related analysis planes: one mitral plane and one aortic plane. The mitral plane was used to determine whether a backward-traced particle had entered the LV through the mitral valve. The aortic plane was used to determine whether a forward-traced particle left the LV through the aortic outflow.

The workflow supports different levels of user interaction. Both planes can be placed manually, both can be handled automatically, or one plane can be manual while the other is automatic. This makes the workflow usable both when the automatic placement works well and when the user needs more control because of dataset-specific anatomy or registration uncertainty.

Mitral plane

The mitral plane could be reused if an existing user-defined plane named *Mitral valve* was already present. If no such plane was available and automatic plane handling was selected, the mitral analysis plane was constructed from cine long-axis data.

The 2CH and 4CH views were used to identify two atrioventricular (AV) plane points in each view from the LV contour. If a valid contour was available at end-diastole, that timeframe was used. Otherwise, the first timeframe with a valid contour was used. This gave four annular reference points in total.

All four points were transferred to the common anatomical coordinate system used in Seg-

ment. If these points are denoted by $\mathbf{p}_1, \ldots, \mathbf{p}_N$, their centroid is

$$\bar{\mathbf{p}} = \frac{1}{N} \sum_{i=1}^{N} \mathbf{p}_i. \tag{4.13}$$

A plane was then fitted through the points by minimizing the squared orthogonal distance from the points to the plane,

$$\min_{\|\mathbf{n}\|=1} \sum_{i=1}^{N} \left[\mathbf{n}^\top (\mathbf{p}_i - \bar{\mathbf{p}}) \right]^2, \tag{4.14}$$

where $\mathbf{n}$ is the plane normal. This gives the plane that best matches the four annular points in a least-squares sense.

After the plane orientation had been determined, a square analysis plane was created around the fitted annular center. Let $d_{2\mathrm{CH}}$ and $d_{4\mathrm{CH}}$ denote the distances between the two AV-plane points in the 2CH and 4CH views. The plane side length was then set to

$$L = 1.5 \, \max(d_{2\mathrm{CH}}, d_{4\mathrm{CH}}). \tag{4.15}$$

This gave a plane that covered the mitral opening with some margin.

The mitral plane was kept static over the full cardiac cycle. If the mitral plane was created automatically, it was shifted 2 mm along its normal direction before the final pathline classification. This shifted plane was the one used in the crossing analysis. The purpose of this shift was to place the plane slightly above the LV and reduce false crossings caused by motion near the basal LV boundary.

Aortic plane

The workflow supports two alternatives for the aortic plane: manual placement and automatic estimation. If a user-defined plane named *Aortic valve* already existed, it could be reused. Otherwise, the automatic method could be used to estimate the aortic plane from the 4D flow velocity data.

In the automatic method, the mitral plane was used as a starting reference. A single enlarged version of the mitral plane was created by scaling the mitral-plane vertices around their center while keeping the same orientation. In the current implementation, this search plane was two times larger than the original mitral plane. For each cardiac timeframe, the speed magnitude was sampled on this enlarged search plane from the 4D flow dataset,

$$s(\mathbf{x}, t) = \|\mathbf{v}(\mathbf{x}, t)\| = \sqrt{v_x^2 + v_y^2 + v_z^2}. \tag{4.16}$$

The highest velocity values on the enlarged search plane were then used to score how likely each timeframe was to correspond to aortic systolic flow. The timeframe with the highest score was selected as an estimate of peak systole.

Speed maps from the selected timeframe and its neighboring frames were averaged before the high-speed region was identified. The main high-speed region on the selected plane was

then identified, and its center was estimated as a weighted average of the sampled points. If the sampled plane points are denoted by $\mathbf{r}_i$ and their corresponding weights by w_i, the estimated center is

$$\mathbf{c} = \frac{1}{\sum_i w_i} \sum_i w_i \mathbf{r}_i. \tag{4.17}$$

This gives a center that is pulled towards the strongest part of the flow region.

After this, the local vessel direction was estimated around the detected center. The implementation sampled several temporary cross-sections near the detected aortic region. In each cross-section, the faster part of the flow was identified, and a weighted center point was calculated. A local vessel axis was then estimated from these center points. The final aortic plane was placed orthogonally to this estimated local flow direction.

The automatically created aortic plane had a fixed side length of 60 mm and was shifted 2 mm downstream in the local flow direction. This automatic method worked well in many cases, but it was not fully robust in all datasets. For this reason, manual aortic-plane placement remained available in the workflow.

4.5.2 Plane-crossing detection

For each pathline, consecutive trajectory points were tested against the relevant analysis plane. The purpose of this step was to determine whether the pathline passed through the plane at any point during its traced lifetime.

Let $\mathbf{x}_0$ be one point on the plane and let $\mathbf{n}$ be the plane normal. The signed distance from a trajectory point $\mathbf{x}(t)$ to the plane was computed as

$$d(t) = \mathbf{n} \cdot (\mathbf{x}(t) - \mathbf{x}_0)\,. \tag{4.18}$$

The sign of $d(t)$ indicates on which side of the plane the point lies. If two consecutive points lie on opposite sides of the plane, then the connecting pathline segment must cross the plane. If $\mathbf{x}_i$ and $\mathbf{x}_{i+1}$ are two consecutive points with signed distances d_i and d_{i+1}, this condition can be written as

$$d_i d_{i+1} \leq 0. \tag{4.19}$$

In other words, the segment was treated as a crossing candidate when it crossed the plane or touched it.

Before this test was performed, non-finite pathline points were removed. This made the crossing detection more robust and avoided invalid calculations.

When a crossing candidate had been found, the intersection point between the pathline segment and the plane was estimated by linear interpolation between the two consecutive trajectory points. The interpolation factor was computed as

$$\alpha = \frac{d_i}{d_i - d_{i+1}}, \qquad \mathbf{p}_{\text{int}} = \mathbf{x}_i + \alpha(\mathbf{x}_{i+1} - \mathbf{x}_i), \tag{4.20}$$

with $0 \leq \alpha \leq 1$. This gives the point where the segment meets the plane.

The intersection point was then expressed in the local two-dimensional coordinate system of the plane. If $\mathbf{u}$ and $\mathbf{v}$ are the two in-plane basis vectors, the projected coordinates were

$$\xi = \mathbf{u}^\top(\mathbf{p}_{\text{int}} - \mathbf{x}_0), \qquad \eta = \mathbf{v}^\top(\mathbf{p}_{\text{int}} - \mathbf{x}_0). \tag{4.21}$$

The point (ξ, η) was then tested against the full polygon of the plane object.

Only intersections within the boundary of the full analysis plane were counted as valid crossings. In the current implementation, backward pathlines were tested against the mitral plane and forward pathlines against the aortic plane. A particle was marked positive for a given event if its pathline crossed the corresponding plane at least once during its traced lifetime.

4.5.3 Flow-component assignment

For each seed i, two outcomes were defined:

$$E_i = \begin{cases} 1, & \text{if the backward pathline crosses the mitral plane,} \\ 0, & \text{otherwise,} \end{cases} \tag{4.22}$$

and

$$L_i = \begin{cases} 1, & \text{if the forward pathline crosses the aortic plane,} \\ 0, & \text{otherwise.} \end{cases} \tag{4.23}$$

Here, E_i represents whether the seeded particle entered the LV, and L_i represents whether it left the LV. The four standard LV flow components were then defined as

$$\text{Direct Flow}: \quad E_i = 1,\ L_i = 1, \tag{4.24}$$

$$\text{Retained Inflow}: \quad E_i = 1,\ L_i = 0, \tag{4.25}$$

$$\text{Delayed Ejection Flow}: \quad E_i = 0,\ L_i = 1, \tag{4.26}$$

$$\text{Residual Volume}: \quad E_i = 0,\ L_i = 0. \tag{4.27}$$

The decision logic is illustrated in Figure 4.3.

If N_c is the number of particles in a component class c, the component fraction becomes

$$P_c = \frac{N_c}{N_{\text{total}}}, \tag{4.28}$$

where N_{total} is the total number of particles used for classification. The final result was reported as percentages, $100P_c$, for the four flow components.

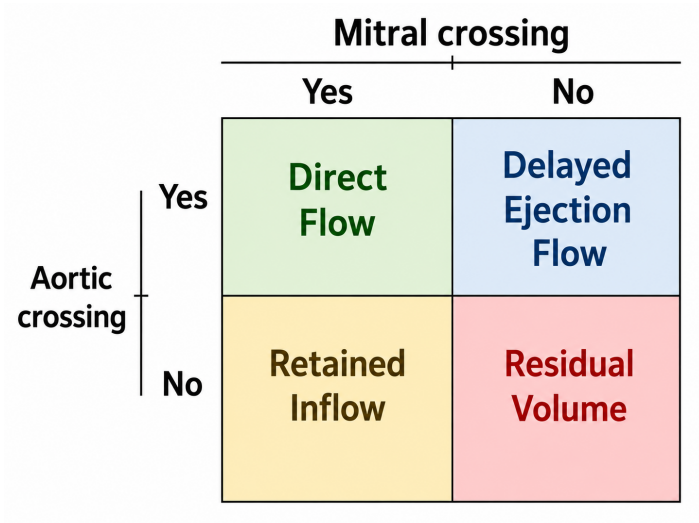

Figure 4.3: Decision table for LV flow-component classification.

4.6 LV kinetic energy calculation

The implementation also calculates an LV kinetic-energy curve. This curve is not used in the flow-component classification. It is included as an additional output that describes the amount of kinetic energy in the LV blood volume over the cardiac cycle.

For each timeframe, an LV mask was generated using the same cine-to-4D flow transfer procedure as in the seed-generation workflow. The kinetic energy was then calculated voxel by voxel inside this LV mask. The velocity magnitude in each voxel was computed from the three measured velocity components,

$$v = \sqrt{v_x^2 + v_y^2 + v_z^2}. \tag{4.29}$$

The voxel volume was calculated from the 4D flow voxel dimensions,

$$V_{\text{voxel}} = \Delta x \, \Delta y \, \Delta z, \tag{4.30}$$

where Δx, Δy, and Δz are the voxel dimensions. In the implementation, the voxel volume was converted from mm^3 to ml. A blood density of 1.06 g/ml was used [25].

The kinetic energy contribution from one voxel was then calculated as

$$KE_{\text{voxel}} = \frac{1}{2} m_{\text{voxel}} v^2, \tag{4.31}$$

where

$$m_{\text{voxel}} = \rho_{\text{blood}} V_{\text{voxel}}. \tag{4.32}$$

Since the velocity data are expressed in cm/s, the result was converted to microjoule [μJ] using the corresponding unit conversion. The total LV kinetic energy at one timeframe was calculated as the sum over all voxels inside the LV mask,

$$KE_{\mathrm{LV}}(t) = \sum_{\mathbf{x} \in LV} KE_{\mathrm{voxel}}(\mathbf{x}, t). \tag{4.33}$$

Repeating this calculation for all timeframes gives a time-resolved LV kinetic-energy curve. The implementation reports the peak and mean LV kinetic energy together with the curve in the Flow Components result tab.

4.7 Phantom validation

The implemented method was evaluated using two controlled synthetic phantom setups. These experiments were used for methodological validation of two central parts of the workflow: numerical pathline integration and plane-crossing based flow-component classification.

4.7.1 Pathline integration validation in a circular-flow phantom

Pathline integration was evaluated in a synthetic circular-flow phantom with analytically defined rotational flow around the central axis of a straight cylinder. In the transverse plane, the exact particle trajectory is therefore a circle with constant radius, providing a ground-truth trajectory against which numerically integrated pathlines can be compared.

A single pathline was seeded at a fixed position 8 pixels from the center of the cylinder. Since the phantom center was located at $(64, 64)$ in the image plane, the analytic reference trajectory was a circle of radius $r = 8$ pixels centered at $(64, 64)$. The same seed position and integration duration were used for both fourth-order Runge–Kutta (RK4) and forward Euler integration.

Two integration settings were evaluated. The user specifies a desired integration time step, `desiredDt`, and the solver determines the actual step size used for integration as

$$\Delta t = \frac{T_{\mathrm{incr}}}{\lceil T_{\mathrm{incr}} / \texttt{desiredDt} \rceil}, \tag{4.34}$$

where the phantom time-frame spacing was $T_{\mathrm{incr}} = 0.05$ s. For the evaluated settings, this resulted in

$$\texttt{desiredDt} = 0.05 \text{ s} \;\Rightarrow\; \Delta t = 0.05 \text{ s}, \qquad \texttt{desiredDt} = 0.005 \text{ s} \;\Rightarrow\; \Delta t = 0.005 \text{ s}. \tag{4.35}$$

Agreement between the numerical pathlines and the analytic circular trajectory was evaluated using two measures: the percentage of trajectory points with radial error within ± 0.1 pixels

of the analytic circle, and a normalized mean-absolute-error-based agreement score, where 100% corresponds to zero mean absolute radial error. The tolerance of ±0.1 pixels was selected to provide a strict comparison of numerical integration accuracy.

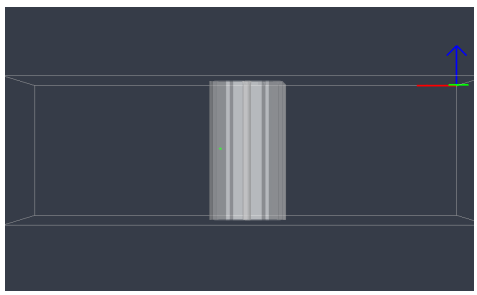

(a) Side view of the circular-flow phantom used for pathline validation. The green marker shows the fixed seed point inside the cylinder.

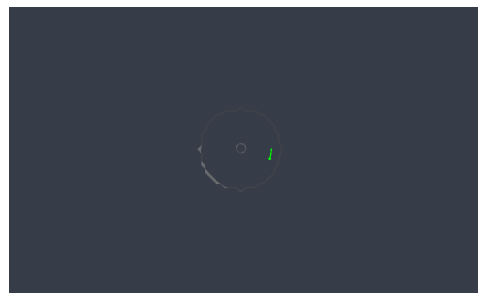

(b) Transverse view of the phantom. The fixed seed point is placed at radius 8 pixels from the phantom center.

Figure 4.4: Geometry of the synthetic circular-flow phantom used for validation of pathline integration.

4.7.2 Flow-component classification validation in a three-cylinder vessel phantom

Flow-component classification was evaluated in a synthetic three-cylinder vessel phantom with temporally constant through-plane flow. The phantom was constructed so that each cylinder represented a predefined crossing pattern. By controlling which analysis plane each cylinder intersected, the expected flow-component classification was known before the analysis was performed.

Forward and backward pathlines were seeded from exactly the same voxel positions in all three cylinders. Seeding was performed in the middle of the volume at $z = 20$.

Two analysis planes were placed at different z-positions. Their placement was selected so that one cylinder intersected only the forward plane, one cylinder intersected only the backward plane, and one cylinder intersected both planes. A pathline crossing only the forward plane should be classified as Delayed Ejection Flow, a pathline crossing only the backward plane should be classified as Retained Inflow, and a pathline crossing both planes should be classified as Direct Flow.

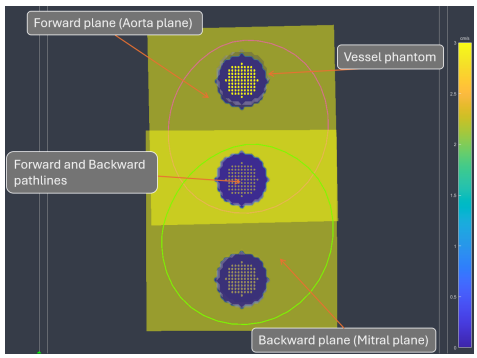

(a) Top view showing the three cylinders and the finite coverage of the analysis planes.

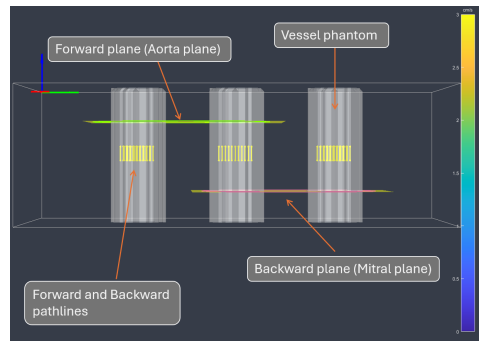

(b) Side view showing the plane positions relative to the seeded pathlines.

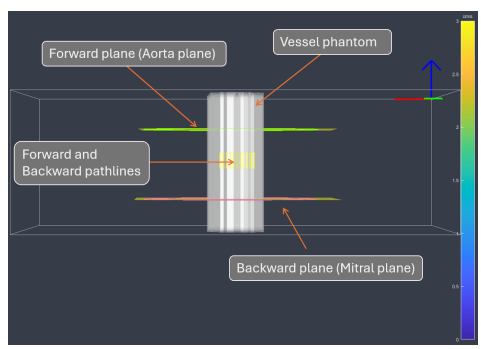

(c) Side view showing the separation between the two analysis planes.

Figure 4.5: Three-cylinder vessel phantom used for validation of flow-component classification. Forward and backward pathlines were seeded from identical voxel positions at $z = 20$ in all three cylinders. The plane placements were chosen to create known crossing patterns: one cylinder intersecting only the forward plane, one intersecting only the backward plane, and one intersecting both planes.

4.8 Selection and quality control of human datasets

All human datasets used in this thesis were retrospective and had been previously approved. The datasets were anonymized before being made available for the work. No new human data were acquired, and no subject contact was performed. The use of the datasets in this project was limited to secondary analysis of anonymized research data. Since these data had been acquired before the development of the present workflow, they were not necessarily optimized for particle-based flow-component analysis. The available datasets could therefore differ in image quality, VENC, cine-to-4D-flow alignment and velocity-field quality.

Before any flow-component results were interpreted, each dataset was assessed according to a set of quality-control criteria. The purpose of this was to determine whether the required input data were available and whether the particle-based analysis was reliable enough to include in the evaluation.

4.8.1 Dataset eligibility and quality-control criteria

The assessment was performed in three stages. First, the required image data and segmentation information were checked. Second, the spatial agreement between the cine-derived LV geometry and the 4D flow data was checked, together with the quality of the velocity information. Third, the resulting seed placement, pathline behaviour, and analysis-plane positions were visually inspected in Segment.

If an automatically estimated analysis plane was poorly placed, manual plane placement could be used instead. A bad automatic plane position therefore did not require exclusion, given that a good manual correction could be made and the analysis passed visual review. Datasets with missing required inputs, severe velocity-field artifacts, major spatial mismatch, or clearly unreliable pathline behaviour were not included in the reported evaluation.

Table 4.2: Dataset usability and quality-control criteria used before interpreting human-dataset flow-component results.

Criterion	Purpose
Available cine SAX data and LV segmentation	Required to define the LV seed region and calculate LV-based outputs.
Available cine LAX data	Required for mitral-plane construction.
Cine-to-4D-flow alignment	The transferred LV geometry and seed region had to match the LV position in the 4D flow data.
Velocity-field quality and aliasing	The velocity field was checked for severe noise, velocity aliasing, or non-physiological flow-direction changes.
Motion or gating artifacts	Datasets with visible motion-related artifacts were excluded from analysis.
Good seed placement	Seed points had to be located inside the intended LV cavity after transfer to the 4D flow grid.
Reasonable pathline behaviour	Pathlines had to move in intracardiac flow regions without clear non-physiological drift.
Plausible mitral and aortic plane placement	The analysis planes had to be placed at reasonable inflow and outflow locations.

4.8.2 Analysis strategy for included human datasets

Datasets that passed the quality review were used for the flow-component evaluation. The evaluation included healthy adult datasets, heart-failure datasets, and healthy paediatric datasets.

For the heart-failure comparison, datasets were grouped according to left-ventricular ejection fraction. The groups were defined as EF < 40% for heart failure with reduced ejection fraction (HFrEF), EF 40–50% for heart failure with mildly reduced ejection fraction (HFmrEF), and EF > 50% for heart failure with preserved ejection fraction (HFpEF). Manual analysis-plane results were used for this EF-group comparison in order to focus on differences in flow-component distributions rather than on errors caused by automatic plane placement.

For the healthy paediatric datasets, both automatic and manual plane-placement results were evaluated to assess how the workflow behaved in smaller hearts. Automatic and manual plane placement were also compared using the internal consistency measure described below.

4.8.3 Included cohorts and demographic characteristics

After applying the dataset quality-control criteria described in Section 4.8, 60 human datasets were included in the LV flow-component evaluation. These consisted of 15 healthy adult datasets, 39 heart-failure datasets, and 6 healthy paediatric datasets. The heart-failure datasets were further divided into 7 HFrEF datasets, 5 HFmrEF datasets, and 27 HFpEF datasets.

Table 4.3 summarizes the cohort sizes, age ranges, and recorded sex distribution for the

datasets included in the LV evaluation. Age and sex information were not available for every dataset. The age ranges therefore refer only to datasets with recorded age information, and missing sex information is reported explicitly.

Table 4.3: Characteristics of the human datasets included in the LV flow-component evaluation. Age ranges are based on datasets with available age information.

Cohort	Number of datasets	Age range [years]	Female	Male	Sex not recorded
Healthy adults	15	24–66	3	9	3
HFrEF	7	58–75	1	5	1
HFmrEF	5	53–75	1	2	2
HFpEF	27	47–78	9	15	3
Healthy paediatric	6	6–13	2	4	0
Total	60	6–78	16	35	9

4.8.4 Internal consistency measure and visual review

An internal consistency measure was calculated from the flow-component percentages. The difference between Retained Inflow and Delayed Ejection Flow was defined as

$$\Delta_{\mathrm{RI-DE}} = P_{\mathrm{RI}} - P_{\mathrm{DE}}, \tag{4.36}$$

where P_{RI} is the percentage of particles classified as Retained Inflow and P_{DE} is the percentage classified as Delayed Ejection Flow. The absolute value,

$$|\Delta_{\mathrm{RI-DE}}|\,, \tag{4.37}$$

was used as a practical quality-control indicator.

Retained Inflow consists of particles that enter the LV during the analysed cycle but are not ejected during the same cycle. Delayed Ejection Flow consists of particles that are already present in the LV and are ejected during that cycle. These exchange components are expected to be similar. A value of $|\Delta_{\mathrm{RI-DE}}|$ close to zero therefore suggests that the backward mitral-crossing analysis and the forward aortic-crossing analysis are reasonably balanced.

A large $|\Delta_{\mathrm{RI-DE}}|$ does not identify the exact source of the error, but indicates that the analysis should be reviewed. Possible causes could be inaccurate plane placement, bad cine-to-4D-flow alignment, seed points close to uncertain LV boundaries, velocity-field noise or aliasing. The measure was therefore used as a quality-control indicator rather than as a physiological marker.

5 Results

This chapter presents the results from the implemented flow-component workflow. The results are organized from controlled technical validation to real human datasets. First, pathline integration was evaluated in an analytic circular-flow phantom. Second, the plane-crossing logic and flow-component classification were evaluated in a numerical three-cylinder vessel phantom. Third, the implemented Segment workflow is shown to illustrate how the method can be used in practice. The workflow was then applied to retrospective human datasets to evaluate automatic plane placement, quality, and flow-component patterns across EF groups, healthy humans, children and the RV.

5.1 Implemented contributions and scope

The implemented work resulted in a research workflow in Segment for LV flow-component analysis from 4D flow MRI. The main implemented contributions were:

- generation of an LV seed region from cine-derived LV segmentation and transfer of this region to the 4D flow analysis space,
- forward and backward pathline integration from the same LV seed points using RK4 integration,
- construction and handling of mitral and aortic analysis planes for pathline crossing detection,
- classification of valid particles into the four LV flow components based on mitral and aortic plane crossings,
- visualization of LV contours, seed points, analysis planes, pathlines, and component-colored particles for visual review,
- quantitative output of LV flow-component results and LV kinetic energy curves,
- verification of pathline integration and classification logic using controlled phantom experiments,
- application of the workflow to retrospective human datasets, including healthy subjects, heart-failure datasets, and paediatric datasets,

- comparison between automatic and manual analysis-plane placement,
- a proof-of-concept application of the workflow to the right ventricle.

The implemented method should be viewed as a research workflow. The focus of the work was to define, implement, visualize, and test the main processing steps needed for flow-component analysis. The following sections present the validation experiments, representative workflow output, and the results from applying the method to human datasets.

5.2 Pathline integration validation in the circular-flow phantom

Numerical pathlines obtained using RK4 and forward Euler were compared with the analytic circular trajectory described in Section 4.7.1. The phantom geometry and seed placement used for this validation are shown in Figure 4.4.

5.2.1 Qualitative comparison between RK4 and forward Euler

Figures 5.1 and 5.2 compare the numerical pathlines with the analytic circular trajectory for RK4 and forward Euler.

For RK4, the numerical trajectory closely overlapped the analytic circle at both tested time steps. Reducing the time step from $\Delta t = 0.05$ s to $\Delta t = 0.005$ s produced no visible improvement at the plotting scale.

Forward Euler was more step-size dependent. At $\Delta t = 0.05$ s, the trajectory drifted outward from the analytic circle, indicating accumulated numerical error. At $\Delta t = 0.005$ s, the result improved substantially, but a visible deviation from the analytic solution remained.

Overall, RK4 was more stable and accurate than forward Euler in the circular-flow phantom.

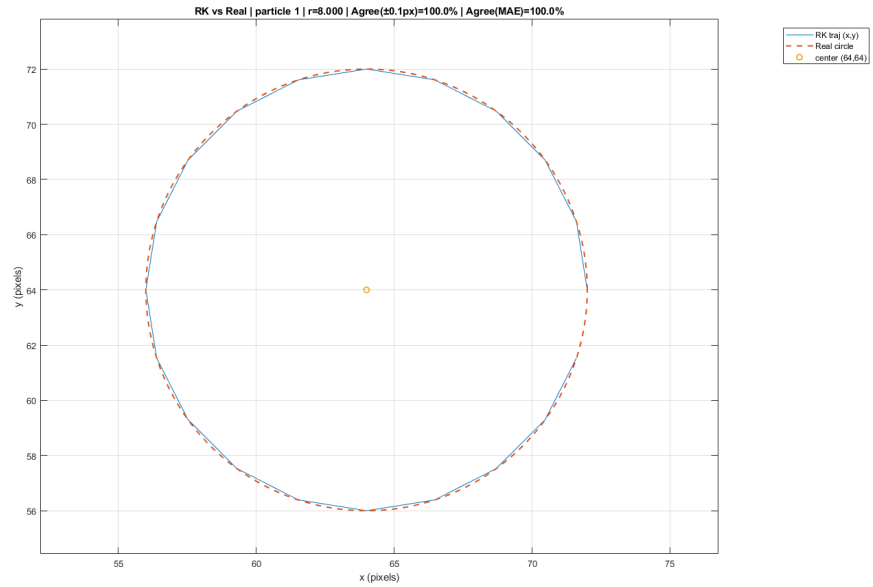

(a) RK4 versus analytic circle for $\Delta t = 0.05$ s.

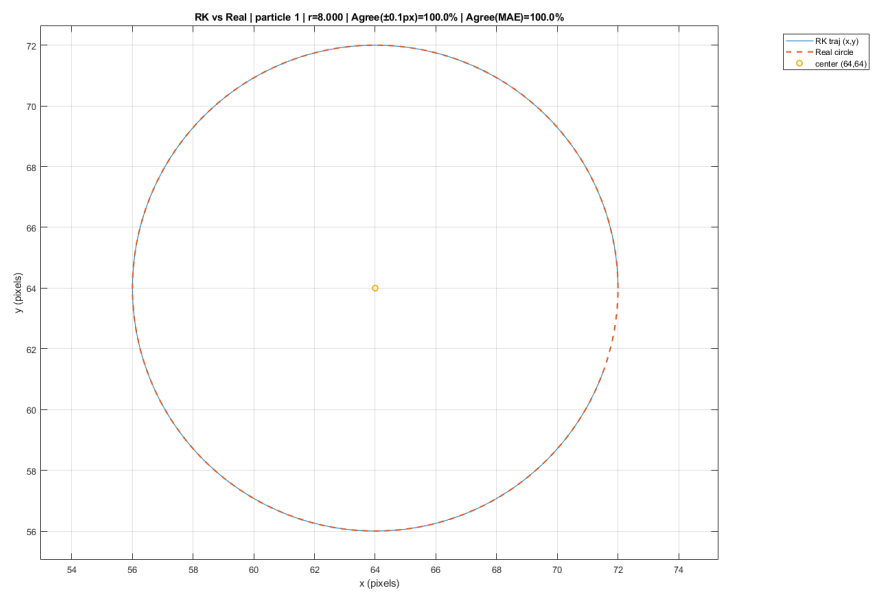

(b) RK4 versus analytic circle for $\Delta t = 0.005$ s.

Figure 5.1: Validation of pathline integration in the circular-flow phantom using RK4. The dashed curve is the analytic ground-truth orbit and the solid curve is the numerically integrated pathline.

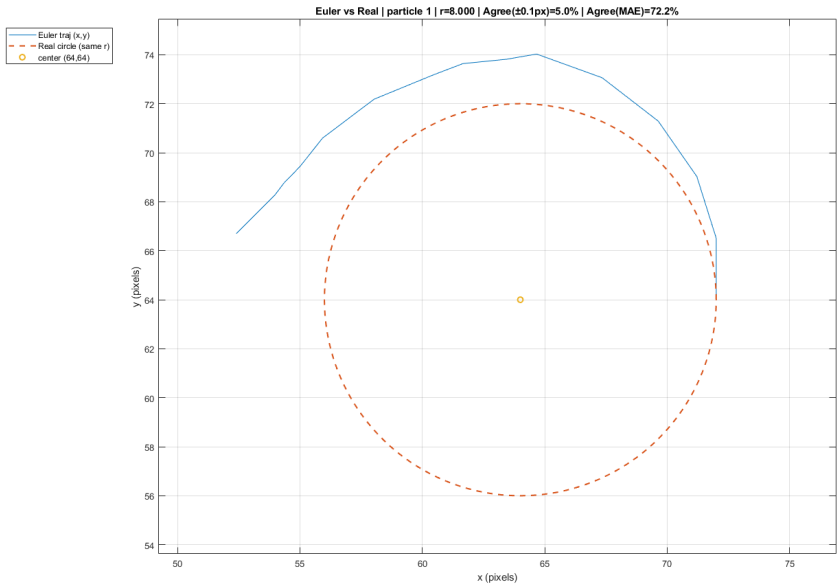

(a) Forward Euler versus analytic circle for $\Delta t = 0.05$ s.

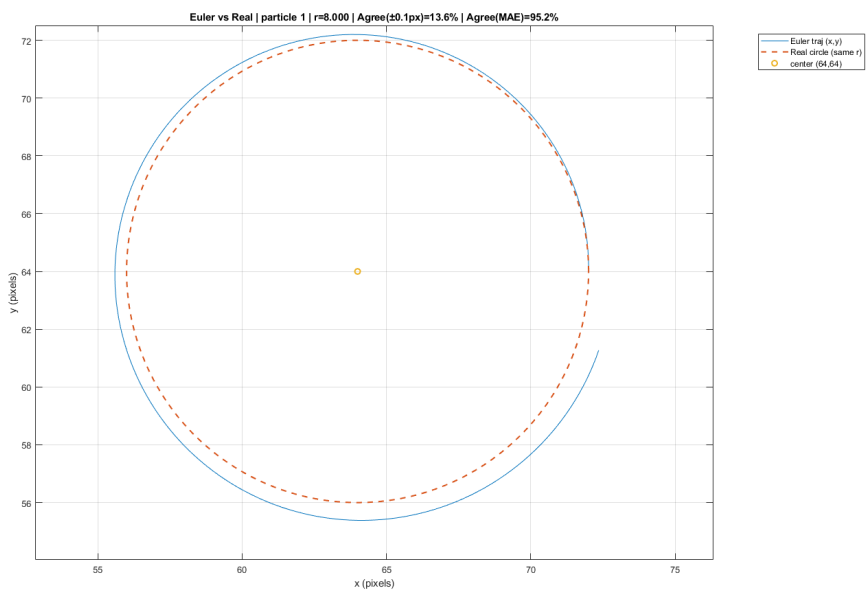

(b) Forward Euler versus analytic circle for $\Delta t = 0.005$ s.

Figure 5.2: Validation of pathline integration in the circular-flow phantom using forward Euler. The dashed curve is the analytic ground-truth orbit and the solid curve is the numerically integrated pathline. Forward Euler showed stronger time-step dependence than RK4.

5.2.2 Quantitative agreement with ground truth

The comparison is summarized in Table 5.1. For RK4, both tested cases gave 100.0% agreement within ±0.1 pixels and 100.0% in the MAE-based agreement score, consistent with the near-perfect overlap in Figure 5.1.

For forward Euler, the coarser time step $\Delta t = 0.05$ s gave 5.0% agreement within ±0.1 pixels and an MAE-based agreement of 72.2%. Reducing the time step to $\Delta t = 0.005$ s improved these values to 13.6% and 95.2%, respectively, but forward Euler remained less accurate than RK4.

Together, these results support the use of RK4 as the default pathline integration method in the implemented workflow.

Table 5.1: Agreement between numerical pathlines and analytic ground truth in the circular-flow phantom.

Method	Case	r (px)	Δt (s)	Agreement within ±0.1 px [%]	MAE-based agreement [%]	Comment
RK4	1	8.000	0.05	100.0	100.0	Near-perfect overlap
RK4	2	8.000	0.005	100.0	100.0	Near-perfect overlap
Euler	1	8.000	0.05	5.0	72.2	Clear outward drift
Euler	2	8.000	0.005	13.6	95.2	Improved, but still less accurate than RK4

5.3 Flow-component classification in the three-cylinder vessel phantom

The classification outcome obtained in the three-cylinder vessel phantom is summarized in Table 5.2. The phantom configuration used for this validation is shown in Figure 4.5 and described in Section 4.7.2.

The measured classification matched the expected outcome exactly. Of the 243 seeded pathlines, 81 were classified as Direct Flow, 81 as Retained Inflow, and 81 as Delayed Ejection Flow. No pathlines were classified as Residual Volume, since no cylinder was configured to remain within the phantom without intersecting either analysis plane.

These results confirm that the implemented plane-crossing logic and particle-level flow-component classification rules behave as intended when the expected crossing patterns are known in advance.

Table 5.2: LV flow-component classification results in the three-cylinder vessel phantom.

Component	Count (seeds)	Fraction [%]
Direct Flow	81	33.3
Retained Inflow	81	33.3
Delayed Ejection Flow	81	33.3
Residual Volume	0	0.0
Total	243	100.0

Together, the two phantom experiments validate technical parts of the workflow. The circular-flow phantom tests whether the pathline integration follows a known trajectory, and the three-cylinder vessel phantom tests whether the plane-crossing and classification rules give the expected component-results. However, these tests do not validate the complete workflow in real cardiac data. They do not include important sources of uncertainty such as image noise, segmentation errors or cine-to-4D-flow mismatch. The phantom results should therefore be interpreted as technical validation of key substeps, not as full validation of the complete method.

5.4 Implemented Segment workflow and representative output

The implemented method was integrated into the 4D flow module in Segment as a user-facing Flow Components workflow. The workflow makes it possible to generate LV seed points, define the mitral and aortic analysis planes, compute forward and backward pathlines, classify the particles into flow components, and inspect the resulting component distribution within the same analysis environment.

Figure 5.3 shows an overview of the implemented Flow Components workflow in the 4D flow module in Segment. The figure highlights user interface elements that were added and/or used in the implemented workflow, including the Flow Components button, the 4D flow object list, and the Flow Components results tab. This overview shows how the method was integrated into the existing Segment environment and how the workflow outputs can be inspected in practice.

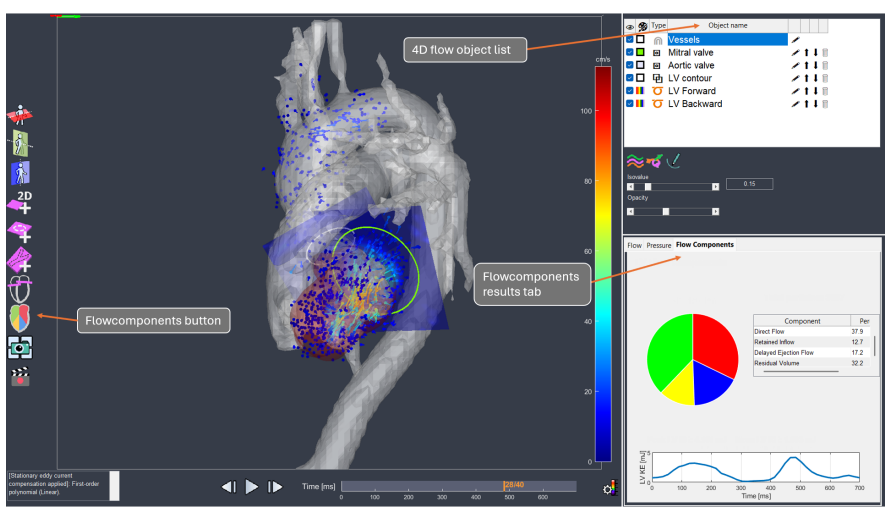

Figure 5.3: Overview of the implemented Flow Components workflow in the 4D flow module in Segment. The figure shows user interface elements added or used in the workflow, including the Flow Components button, the 4D flow object list, and the Flow Components results tab.

The workflow supports flexible analysis-plane handling. The mitral and aortic planes can both be created automatically, both be placed manually, or placed using a mixed setup where one plane is manual and the other automatic. In the automatic mitral-plane mode, the mitral plane is constructed from 2CH and 4CH cine contours. In the automatic aortic-plane mode, the aortic plane is estimated from the 4D flow velocity data. Manual placement remains available for both planes, which is important when automatic placement is poor or when the user wants to inspect a specific plane definition.

Figure 5.4 illustrates how the implemented workflow combines the two tracing directions with the analysis planes. Forward pathlines are evaluated relative to the aortic plane, while backward pathlines are evaluated relative to the mitral plane. The combined view shows how both tracing directions and both planes can be inspected together in Segment.

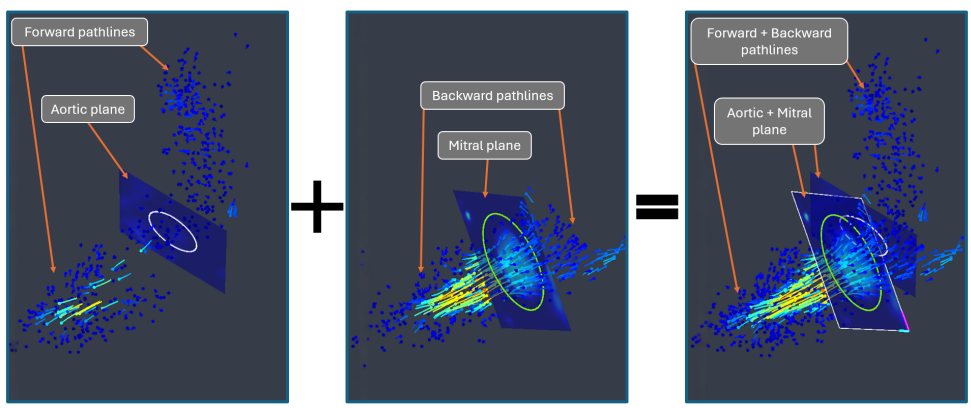

Figure 5.4: Visualization of forward and backward pathlines together with the analysis planes in Segment. Left: forward pathlines shown together with the aortic plane. Middle: backward pathlines shown together with the mitral plane. Right: combined visualization of forward and backward pathlines together with both analysis planes.

The workflow supports two display modes. In dynamic display mode, a reduced number of pathlines is shown to make the particle transport visually inspectable without displaying all the particles used for analysis. In static display mode, the classified pathlines are displayed as component-colored objects corresponding to Direct Flow, Retained Inflow, Delayed Ejection Flow, and Residual Volume. The full classification is performed using hidden analysis pathlines generated from all LV seed points, while the displayed pathlines are downsampled for easier visualization and faster computation.

Figure 5.5 shows a representative example of the static component display in Segment. In this mode, the classified pathlines are shown as separate component-colored objects, which makes it easier to visually inspect the spatial distribution and characteristic trajectory pattern

of each flow component individually.

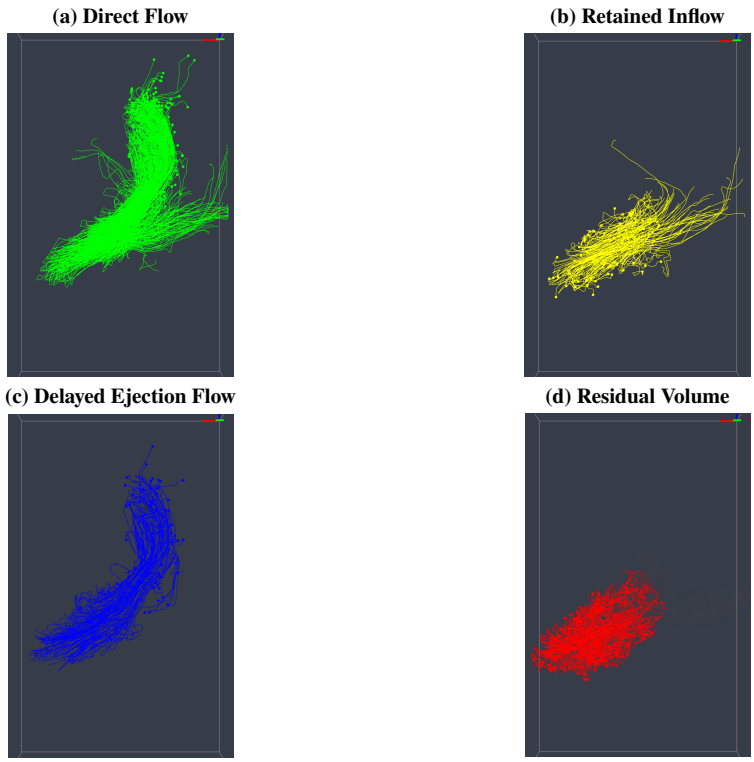

Figure 5.5: Representative example of the static flow-component display in Segment. The classified pathlines are shown separately for Direct Flow, Retained Inflow, Delayed Ejection Flow, and Residual Volume.

The final result is displayed as the percentage of particles assigned to each of the four flow components, as shown in Figure 5.6. These percentages are the main user-facing output, since they allow comparison between analyses of different datasets. In addition to the flow-component distribution, the implementation also displays an LV kinetic-energy curve based on the LV mask and the 4D flow velocity data. This curve is an auxiliary output and is not used in the flow-component classification itself.

Figure 5.7 shows two additional visualization outputs from the implemented workflow. The LV contour object provides a visual representation of the cine-derived LV endocardium after transfer to the 4D flow space. The seeded pathline view shows how particles are released from the LV seed region and traced through the velocity field.

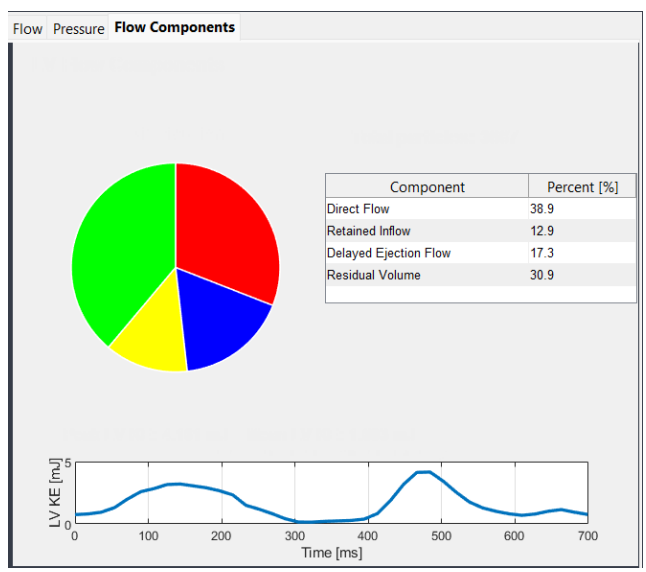

Figure 5.6: User-facing output in the Flow Components tab after running the workflow.

(a) LV contour object displayed in the 4D flow environment.

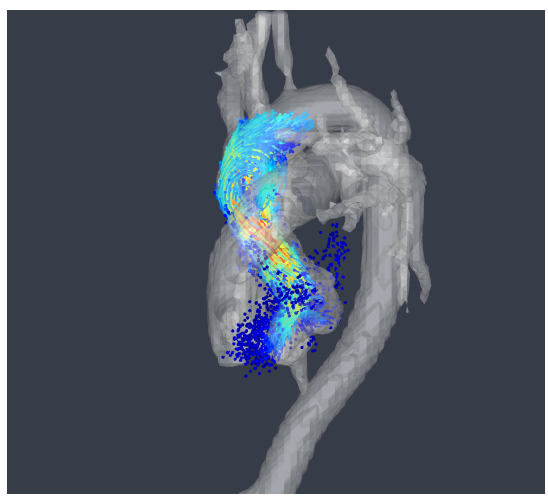

(b) Example of LV seed points traced forward in time.

Figure 5.7: Additional visualization outputs from the implemented workflow. The LV contour object shows the transferred LV geometry, while the pathline view provides a visual check of seed placement and particle transport in Segment.

5.5 Workflow evaluation using human datasets

5.5.1 Application on healthy subject datasets

The implemented workflow was applied to 15 healthy subject datasets. For each dataset, the LV blood volume was classified into Direct Flow, Retained Inflow, Delayed Ejection Flow, and Residual Volume.

The mean flow-component distribution in the healthy cohort is as follows: Residual Volume was the largest component, with a mean value of 35.7%. Direct Flow had a mean value of 26.1%. Retained Inflow and Delayed Ejection Flow were similar in magnitude, with mean values of 19.4% and 18.9%, respectively.

Figure 5.8 shows the healthy cohort results. The pie chart shows the mean flow-component distribution, while the scattered boxplot shows the individual subject values for each flow component.

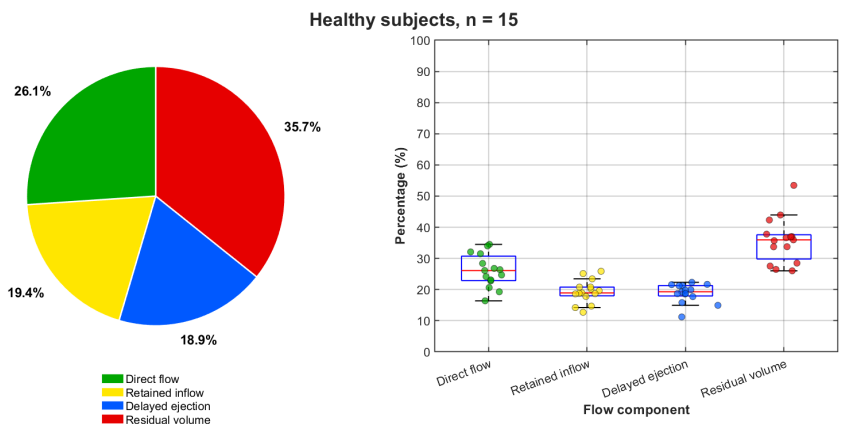

Figure 5.8: LV flow-component results for the healthy subjects ($n = 15$). The pie chart shows the mean distribution, and the scattered boxplot shows the individual subject values for each flow component.

The individual subject values showed variation within each flow component. Residual Volume showed the highest mean value and the widest spread across subjects. Direct Flow was the second largest component, while Retained Inflow and Delayed Ejection Flow had lower and similar mean values.

5.5.2 Application on heart-failure datasets

The workflow was also applied to anonymised datasets from patients with varying degrees of heart failure. This section presents the results obtained using manually placed analysis planes. Manual-plane results were used for this EF-group comparison in order to focus on

the flow-component distributions rather than on errors caused by automatic plane placement.

The datasets were grouped according to left-ventricular ejection fraction (EF). EF describes the percentage of blood ejected from the LV during one cardiac cycle, and lower EF corresponds to more reduced systolic function. The three groups were EF < 40%, corresponding to heart failure with reduced ejection fraction (HFrEF), EF 40–50%, corresponding to heart failure with mildly reduced ejection fraction (HFmrEF), and EF > 50%, corresponding to heart failure with preserved ejection fraction (HFpEF). The manual-plane analysis included 39 unique datasets: 7 HFrEF datasets, 5 HFmrEF datasets, and 27 HFpEF datasets. For each EF group, the flow-component distribution was calculated.

Figures 5.9, 5.10, and 5.11 show the results graphically. Each figure combines the mean group distribution as a pie chart with the individual patient values shown as scattered boxplots.

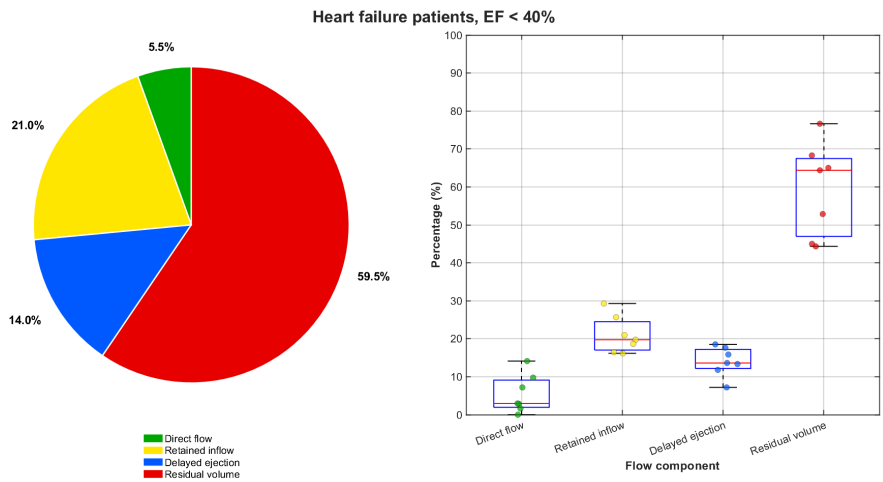

Figure 5.9: Manual-plane LV flow-component results for HFrEF patients with EF < 40% ($n = 7$).

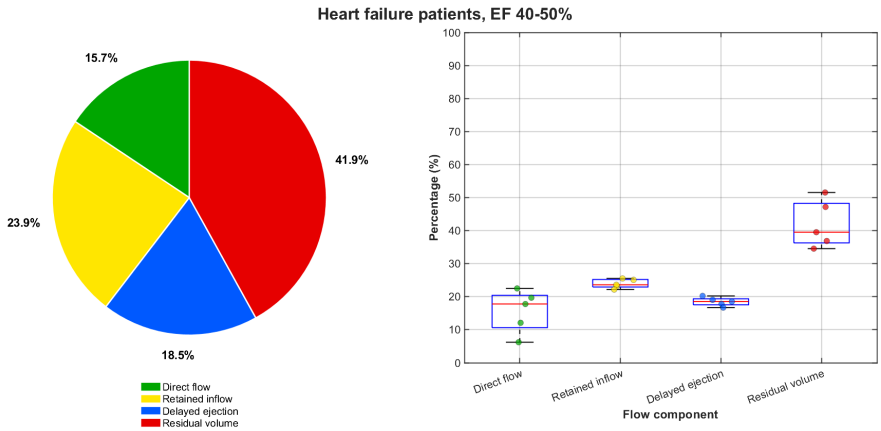

Figure 5.10: Manual-plane LV flow-component results for HFmrEF patients with EF 40–50% ($n = 5$).

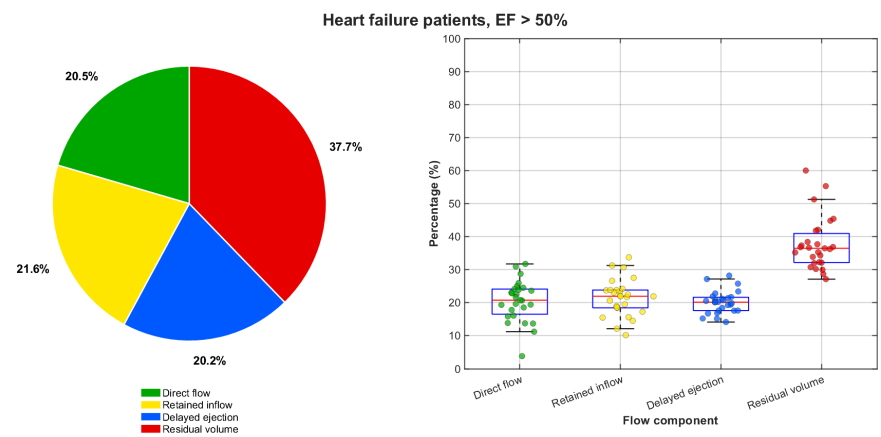

Figure 5.11: Manual-plane LV flow-component results for HFpEF patients with EF > 50% ($n = 27$).

The HFrEF group showed the most distinct flow-component pattern. Residual Volume was the largest component, with a mean value of 59.5%, while Direct Flow was the smallest component, with a mean value of 5.5%. In the HFmrEF and HFpEF groups, Residual Volume was lower, with mean values of 41.9% and 37.7%, respectively. Direct Flow was higher in these groups than in the HFrEF group, with mean values of 15.7% in HFmrEF and 20.5% in HFpEF. Retained Inflow and Delayed Ejection Flow were more similar across the HFmrEF and HFpEF groups.

5.5.3 Application on paediatric datasets

The workflow was also applied to healthy paediatric datasets ($n = 6$) to evaluate how the implementation performed in smaller hearts. Both automatic and manual plane-placement results were evaluated. For these datasets, a two-voxel in-plane erosion of the mapped LV seed mask was used before seed placement.

The automatic and manual analyses produced similar mean flow-component distributions. Direct Flow was 31.2% with automatic plane placement and 29.6% with manual plane placement. Residual Volume was 21.9% in the automatic analysis and 24.1% in the manual analysis.

Figures 5.12 and 5.13 show the corresponding pie charts and scattered boxplots. The pie charts show the mean flow-component distribution, while the scattered boxplots show the individual subject values for each component.

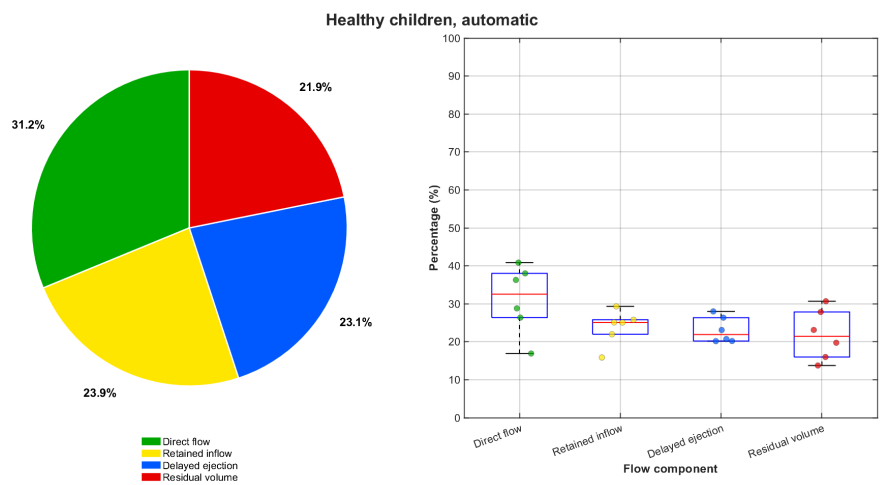

Figure 5.12: Automatic-plane LV flow-component results for healthy children ($n = 6$).

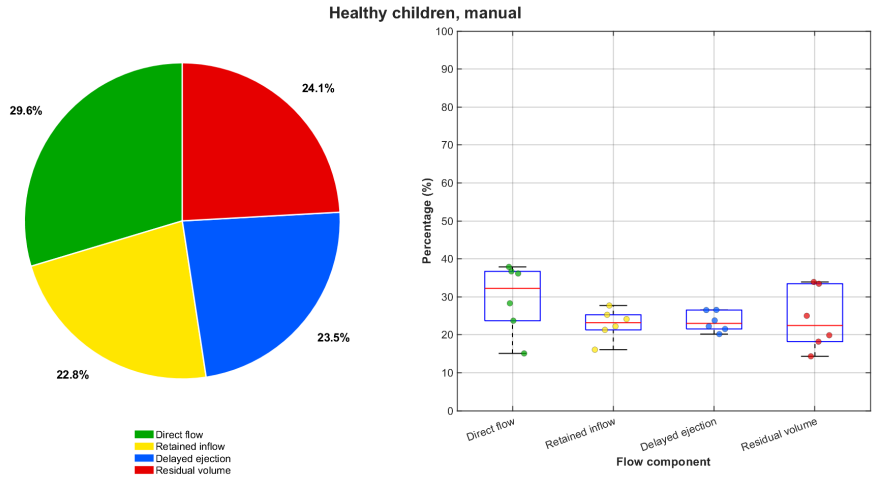

Figure 5.13: Manual-plane LV flow-component results for healthy children ($n = 6$).

The individual subject values showed variation within each component, but the mean distributions were similar between automatic and manual plane placement. Direct Flow was the largest component in both analyses, while the remaining components were distributed more evenly.

5.5.4 Internal consistency with automatic and manual plane placement

To compare internal consistency under automatic and manual plane placement, the RI–DE consistency measure was calculated for both analyses. For each dataset, the absolute RI–DE value was defined as

$$|\Delta_{\mathrm{RI-DE}}| = |P_{\mathrm{RI}} - P_{\mathrm{DE}}|, \tag{5.1}$$

where P_{RI} is the percentage of particles classified as Retained Inflow and P_{DE} is the percentage classified as Delayed Ejection Flow.

Table 5.3 summarizes the absolute RI–DE values for manual and automatic plane placement. Values are reported as mean ± SD in percentage points. The final column shows the difference between the automatic and manual mean values. A difference close to zero indicates that automatic and manual plane placement produced similar absolute RI–DE values.

The Healthy cohort showed similar mean absolute RI–DE values for manual and automatic plane placement, with a difference of 0.2 percentage points. The Heart failure cohort had higher absolute RI–DE values than the Healthy cohort for both manual and automatic plane placement, with an automatic-minus-manual difference of 0.6 percentage points. The Paediatric cohort showed intermediate absolute RI–DE values, with the largest difference between automatic and manual plane placement at 1.1 percentage points.

Table 5.3: Absolute RI–DE consistency measure for manual and automatic plane placement. Values are reported as mean ± SD in percentage points. The final column shows the difference between the automatic and manual mean values.

Cohort	Manual |RI − DE| [%p]	Automatic |RI − DE| [%p]	Difference [%p]
Healthy	2.9 ± 2.4	3.1 ± 2.3	0.2
Heart failure	6.2 ± 4.0	6.8 ± 4.6	0.6
Paediatric	3.5 ± 1.9	4.6 ± 1.8	1.1

5.5.5 Left-ventricular kinetic-energy output

In addition to the flow-component percentages, the implemented workflow exported maximum and average LV kinetic energy for each analysed dataset. To account for differences in ventricular size, the average LV kinetic energy was normalized by the end-diastolic volume (EDV) for each dataset:

$$\text{Average LV KE}_{\text{norm}} = \frac{\text{Average LV KE}}{\text{EDV}}. \tag{5.2}$$

Table 5.4 summarizes the average LV kinetic energy and EDV-normalized average LV kinetic energy for each dataset group.

Table 5.4: Average LV kinetic energy and EDV-normalized average LV kinetic energy in the analysed human datasets. Values are reported as mean ± SD.

Dataset group	n	Average LV KE [μJ]	Average LV KE/EDV [μJ/ml]
Healthy adults	15	1710 ± 740	9.6 ± 3.5
HFrEF	7	2670 ± 1320	8.3 ± 2.3
HFmrEF	5	1970 ± 720	10.0 ± 2.9
HFpEF	27	1840 ± 730	10.9 ± 4.1
Healthy children	6	1490 ± 890	11.7 ± 3.5

The HFrEF group had the highest absolute average LV kinetic energy, with a mean value of 2670 μJ. After EDV normalization, the HFrEF group had the lowest average value of 8.3 μJ/ml. The highest EDV-normalized average value was observed in the healthy paediatric datasets.

5.5.6 Proof-of-concept application to the right ventricle

As a final proof-of-concept test, the implemented flow-component workflow was applied to the right ventricle (RV). The purpose of this analysis was to evaluate whether the same general workflow could be used outside the LV by using the RV segmentation as the seed region and adapting the analysis planes to the tricuspid inflow and pulmonary outflow geometry.

This analysis was performed as a technical feasibility test and was not intended as a formal physiological validation of RV flow components.

The RV analysis was performed in eight high-EF datasets selected from healthy subjects and patients with HFpEF, all with EF above 60%. For each dataset, pathlines were generated from the RV seed region and classified using the same general crossing-based logic as in the LV analysis.

Figure 5.14 shows the RV flow-component distribution for the eight datasets. The pie chart shows the mean distribution, while the scattered boxplot shows the individual dataset values for each component.

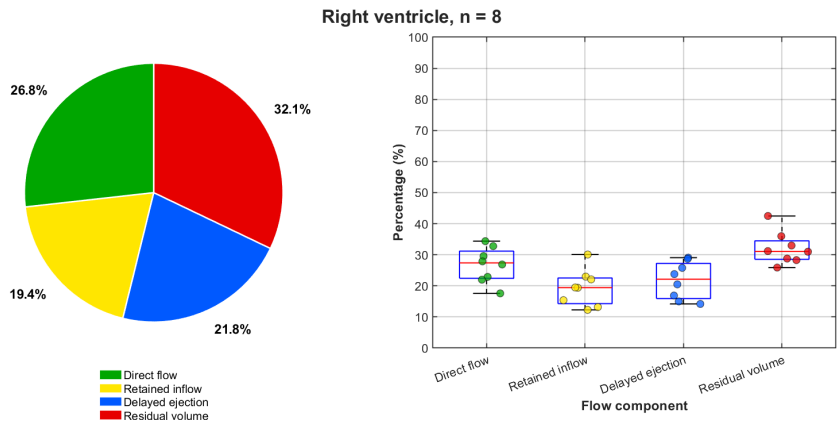

Figure 5.14: Right-ventricular flow-component results for the proof-of-concept analysis ($n = 8$). The pie chart shows the mean distribution, and the scattered boxplot shows the individual dataset values for each flow component.

Figure 5.15 shows representative visual output from the RV proof-of-concept analysis. The examples show that RV seed points could be generated from the RV segmentation and that pathlines could be visualized in relation to both anatomical image data and adapted RV analysis planes. The pathline patterns also illustrate transport through the RV and toward the pulmonary outflow region.

The mean RV flow-component distribution was 26.8% Direct Flow, 19.4% Retained Inflow, 21.8% Delayed Ejection Flow, and 32.1% Residual Volume. The main result of this analysis was that the implemented workflow could be extended to RV datasets and produce classified pathline output.

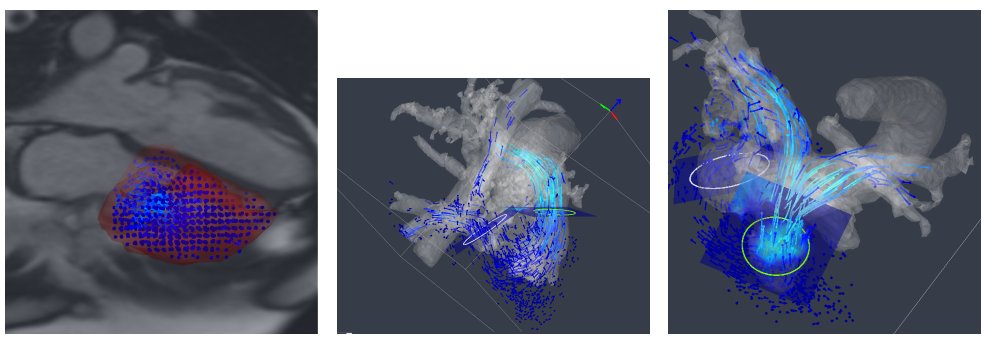

(a) RV seed points and pathlines shown together with anatomical image information.

(b) RV pathlines visualized together with the adapted analysis planes.

(c) Pathlines visualized in relation to the pulmonary outflow anatomy.

Figure 5.15: Proof-of-concept application of the flow-component workflow to the right ventricle. The examples show RV seed points, adapted analysis planes, and pathlines visualized in relation to RV anatomy and pulmonary outflow.

6 Discussion

Conventional measures such as left-ventricular ejection fraction describe global pump function, but not how blood is transported through the ventricle. LV flow-component analysis using 4D flow MRI addresses this by tracing blood transport through the cardiac cycle and dividing the blood volume present in the LV into Direct Flow, Retained Inflow, Delayed Ejection Flow, and Residual Volume, based on whether the particles enter through the mitral valve and/or leave through the aortic valve [11, 12]. For this type of analysis to be useful, the workflow must produce quantitative results that can be reviewed and applied consistently across datasets. This discussion therefore considers both the performance of the implemented workflow and the sources of uncertainty that affect its output.

6.1 Main findings

This thesis developed a semi-automated Segment workflow for quantitative left-ventricular flow-component analysis from 4D flow MRI. The workflow combines cine-derived LV seed generation, forward and backward RK4 pathline tracing, mitral and aortic plane-crossing detection, particle-based flow-component classification, visual review, and LV kinetic-energy calculation.

The phantom experiments tested two central parts of the implementation. In the circular-flow phantom, RK4 reproduced the analytic trajectory more accurately than forward Euler under the tested conditions. In the three-cylinder vessel phantom, the plane-crossing and classification logic reproduced the expected component assignments when the crossing pattern was known in advance. These results support the numerical integration and classification steps under controlled conditions. They do not, however, validate the full workflow in human data, where image quality, segmentation, cine-to-4D-flow alignment, and plane placement add further uncertainty.

The workflow was then applied to retrospective human datasets. Across these datasets, the flow-component distributions were generally plausible, but they did not fully match previously reported healthy reference values. In the healthy adult cohort, Residual Volume was the largest component and Direct Flow was the second largest component, whereas Direct Flow is often reported as the largest component in earlier healthy-adult studies [11, 12, 22]. This suggests that the results are sensitive to cohort composition and implementation choices.

The clearest physiological trend was seen in the heart-failure analysis. The HFrEF group

showed markedly lower Direct Flow and higher Residual Volume than the HFmrEF and HFpEF groups, in line with previously reported trends across the heart-failure spectrum [17]. In the paediatric datasets, automatic and manual plane placement gave similar mean flow-component distributions, although the smaller heart size made these analyses more sensitive to spatial errors and plane-placement settings.

The proof-of-concept RV analysis showed that the same basic seed–pathline–plane-crossing approach could be applied using an RV seed region and analysis planes adapted to tricuspid inflow and pulmonary outflow. This supports the flexibility of the implemented framework, but the RV results were exploratory and were not used for physiological interpretation.

Automatic and manual plane placement gave similar mean Retained Inflow–Delayed Ejection Flow differences at group level, but differences occurred in individual datasets, especially in the heart-failure data. This suggests that automatic plane placement can be a useful starting point and can reduce manual work. The analysis planes, seed region, and pathlines still need to be reviewed before the final component fractions are interpreted.

6.2 Pathline integration and choice of RK4

Pathline integration is a central part of the workflow because the final flow-component classification depends on the estimated motion of each seeded particle. A particle may change component if its pathline is displaced enough to change whether it crosses the mitral or aortic analysis plane. The numerical integration method is therefore part of the quantitative measurement chain, not only a visualization choice.

In the circular-flow phantom, RK4 followed the analytic trajectory more accurately than forward Euler under the tested conditions. Forward Euler showed visible outward drift, especially at the larger time step, whereas RK4 remained close to the expected path at both tested time steps. This is consistent with previous cardiac 4D-flow particle-tracing work, where RK4 has been used as a practical integration method for pathline calculation [26, 27, 9]. These results support the use of RK4 as the default integration method in the workflow.

The phantom result should be interpreted as a test of the numerical integration step. It does not show that pathlines in human datasets are equally accurate, since those pathlines depend on the measured 4D flow MRI velocity field. In human data, uncertainty can still come from spatial and temporal resolution, interpolation, velocity noise, aliasing, and low-flow regions. These factors can affect the estimated particle trajectories even when the integration method itself is stable.

The chosen time-step strategy was a practical compromise between accuracy, reproducibility, and computation time. A target time step of 5 ms was used, and the final step size was adjusted so that each 4D flow MRI timeframe was divided into an integer number of substeps. This gave a consistent integration strategy across datasets with different temporal resolutions.

Boundary handling was also needed for practical use. During integration, particles were kept inside the valid 4D flow volume to avoid interpolation outside the available velocity field. This makes the calculation more robust, but pathlines close to the edge of the image volume

should still be interpreted carefully. Velocity-field limitations and observed pathline failure modes in human datasets are discussed further in Section 6.6.

6.3 Flow-component classification and plane-crossing logic

The flow-component classification was implemented as a rule-based plane-crossing test. For each seeded particle, the backward pathline was used to determine mitral inflow, and the forward pathline was used to determine aortic outflow. These two crossing results were then combined to assign the particle to Direct Flow, Retained Inflow, Delayed Ejection Flow, or Residual Volume.

This structure is one of the main strengths of the workflow. Each particle is classified using the same rule, rather than by visual interpretation of its trajectory. The three-cylinder phantom showed that the implementation reproduced the expected component assignments when the crossing pattern was known in advance. This supports the plane-crossing and classification logic under controlled conditions.

The same rule-based structure also makes the method sensitive to how the analysis planes are defined. A particle was counted as crossing a valve plane if its pathline intersected the full mitral or aortic plane at least once. This gives a clear and reproducible crossing rule, but it also means that the final classification depends on plane position, orientation, and size. If a plane is too large, particles near the valve region may be counted as crossings even if they do not represent true inflow or outflow. If a plane is too small, tilted incorrectly, or placed away from the valve opening, true crossings may be missed.

This sensitivity is especially important for the aortic plane. Direct Flow and Retained Inflow differ only by the aortic crossing result, and Delayed Ejection Flow and Residual Volume also differ only by the aortic crossing result. Errors in the aortic plane can therefore shift particles between ejecting and non-ejecting components. In the same way, errors in the mitral plane can shift particles between blood classified as newly entered and blood classified as already present in the LV.

The difference between Retained Inflow and Delayed Ejection Flow was used as an internal consistency measure. In an approximately cyclic analysis, these two exchange components are expected to be similar in magnitude. A large RI–DE difference may indicate imbalance between inflow and outflow classification, for example due to plane placement, seed-volume definition, finite particle sampling, cine-to-4D-flow mismatch, velocity-field artefacts, interpolation, or boundary handling. The RI–DE difference should not be interpreted as a physiological variable by itself, but as a practical indicator of cases that need visual review.

Another simplification is that the implemented analysis planes are static, while the mitral and aortic valve regions move during the cardiac cycle. Static planes make the workflow easier to reproduce and avoid the need for time-resolved valve tracking, but they simplify the true valve motion. This may affect classifications close to the valve planes, especially in datasets with through-plane motion, unusual anatomy, or imperfect agreement between cine MRI and 4D flow MRI.

A related methodological choice is whether particles should also be checked against the LV segmentation at later cardiac phases. In the present implementation, classification was based on mitral and aortic plane crossings from the pre-systolic seed region. Particles were not excluded based on whether they remained inside the LV segmentation at end-systole. This keeps the classification rule simple and directly tied to the two valve-related crossing tests, but it may also contribute to differences compared with earlier LV flow-component methods, where the relation between particle position and ventricular segmentation over the cardiac cycle has been used more explicitly [11, 12]. Testing an additional segmentation-based inclusion rule would therefore be a useful future comparison.

The component percentages should be interpreted as method-dependent quantitative estimates, not as absolute physiological volumes. The method can be useful for reproducible comparisons between datasets, as long as seed generation, plane placement, crossing logic, and visual review are applied consistently.

6.4 Practical implementation in Segment

The method was implemented directly in Segment so that the analysis could be performed in the same environment as the image data and visual review. This was important because the workflow depends on several objects that need to be checked together: the LV contour, seed points, forward and backward pathlines, mitral and aortic analysis planes, component-coloured pathline objects, and the LV kinetic-energy curve.

A practical benefit of the Segment implementation is that numerical output and visual output are produced together. The component percentages give the quantitative result, while the visual objects help the user detect problems in seed placement, plane position, pathline behaviour, or cine-to-4D-flow alignment. These problems may not be visible from the final percentages alone.

The LV contour object is useful for checking whether the cine-derived LV geometry has been transferred correctly to the 4D flow analysis space. Spatial mismatch between cine MRI and 4D flow MRI may otherwise cause pathlines to start from an incorrectly positioned seed region. This is especially important because the workflow combines anatomical information from cine MRI with velocity information from 4D flow MRI.

The mitral and aortic analysis planes can be created automatically or placed manually. Automatic plane handling reduces user interaction and gives a consistent starting point for the analysis. Manual placement remains necessary when the automatic result is not anatomically plausible. The automatic mitral-plane construction depends on the quality of the cine long-axis contours and the agreement between cine and 4D flow MRI. The automatic aortic-plane method depends on a clear outflow region in the velocity data and may be affected by noise, aliasing, unclear outflow patterns, small anatomy, or imperfect image alignment.

The human-dataset evaluation showed that automatic plane placement was useful in many datasets, but visual review was still needed. Implausible plane positions could occur, and manual adjustment was necessary in some cases. The automatic functions should therefore be viewed as tools for reducing manual work, not as replacements for user review.

It is also important to distinguish between analysis pathlines and displayed pathlines. Classification is performed using the analysis seed set, while the displayed pathlines may be downsampled to keep the visualisation clear and the computation practical. The pathlines shown in the interface and figures are used for qualitative review. The reported component percentages are based on the analysis particles.

The number of seed points was limited to keep computation doable on ordinary workstations. When the number of candidate positions exceeded this limit, a fixed random selection was used so that the selected seed set was reproducible between repeated runs. The LV blood volume is still represented by a finite particle sample, so the component fractions should be interpreted together with the visual output.

Overall, the Segment implementation brings particle tracing, flow-component classification, quantitative output, and visual review into one analysis environment. This makes the method practical for research use, while the remaining sensitivity to data quality, seed placement, and plane placement motivates further development toward a more robust and less user-dependent workflow.

6.5 Application to retrospective human datasets

The application of the workflow to retrospective human datasets tested how the method performed outside the controlled phantom experiments. These datasets included variation in image quality, segmentation quality, ventricular anatomy, velocity-field quality, and cine-to-4D-flow agreement. The results should therefore be interpreted both as flow-component output and as a practical test of how robust the workflow was in real analysis situations.

6.5.1 Healthy adult datasets

In the healthy adult cohort, Residual Volume was the largest component, with a mean value of 35.7%, while Direct Flow had a mean value of 26.1%. Retained Inflow and Delayed Ejection Flow were similar in magnitude, with mean values of 19.4% and 18.9%, respectively. The similar Retained Inflow and Delayed Ejection Flow values suggest that the inflow and outflow classification was reasonably balanced at group level.

The distribution differs slightly from previous LV flow-component studies in healthy adults, where Direct Flow is often reported as the largest component and typically accounts for approximately 35–40% of the LV end-diastolic volume [11, 12, 22]. Residual Volume was higher and Direct Flow was lower than these reference values. This difference should be acknowledged, but not overinterpreted. The results still produced a physiologically plausible distribution, and the exchange components were closely balanced.

Several factors may explain the lower Direct Flow and higher Residual Volume in the present healthy cohort. The cohort was relatively small, and flow-component values are known to vary with age and sex [28]. The results are also method-dependent and may be affected by seed-region definition, basal-slice removal, LV mask erosion, plane placement, temporal

resolution, velocity-field quality, and the crossing logic used for classification. In addition, the present implementation did not exclude particles based on their position relative to the LV segmentation at end-systole, which may contribute to differences compared with earlier methods.

The healthy adult results should therefore be interpreted as plausible method-dependent estimates rather than as reference values. The main finding is that the workflow produced stable and interpretable component distributions in healthy datasets, while also showing a tendency toward higher Residual Volume and lower Direct Flow than reported in some previous healthy-adult studies.

6.5.2 Heart-failure EF groups

The heart-failure analysis showed the clearest physiological trend in the human-data evaluation. In the manual-plane analysis, the HFrEF group had markedly lower Direct Flow than the groups with higher EF. Direct Flow was 5.5% in HFrEF, compared with 15.7% in HFmrEF and 20.5% in HFpEF. Residual Volume showed the opposite pattern, with a mean value of 59.5% in HFrEF, compared with 41.9% in HFmrEF and 37.7% in HFpEF. This pattern is consistent with a less efficient direct through-flow in patients with reduced systolic function.

The HFrEF result agrees well with previous 4D-flow CMR studies of altered LV flow patterns in ventricular dysfunction and heart failure [2, 3, 17]. In particular, Wong et al. reported a progressive reduction in Direct Flow and an increase in Residual Volume across the heart-failure spectrum [17]. Their reported Direct Flow fraction in HFrEF was 5.2%, which is very close to the 5.5% found in the present HFrEF group. This supports the ability of the workflow to reproduce a previously reported heart-failure-related flow pattern.

Residual Volume in the present HFrEF group was higher than the value reported by Wong et al. This difference should not be overinterpreted, since the present analysis used retrospective datasets, a different implementation, manual plane placement, and a limited number of HFrEF cases. It may also reflect methodological differences in particle inclusion, segmentation handling, plane placement, and crossing logic. The important finding is therefore not the exact Residual Volume value, but the direction of the trend: lower Direct Flow and higher Residual Volume in the group with reduced EF.

The separation between HFmrEF and HFpEF was less clear than the separation between HFrEF and the other groups. This is expected, since the HFmrEF group contained only five datasets, making the mean value sensitive to individual cases. In addition, preserved EF alone is not enough to define HFpEF clinically, since guideline definitions also require evidence of increased filling pressures or other supporting clinical findings [16]. The present analysis grouped datasets by EF and did not include detailed clinical information such as heart-failure aetiology, symptom status, medication, comorbidities, age, sex, or disease duration.

For these reasons, the heart-failure results should be interpreted as descriptive flow-component patterns across EF groups, not as a clinical comparison between heart-failure phenotypes. The strongest result is the clear HFrEF pattern, which agrees with previous literature and suggests that the implemented workflow can capture relevant differences in intraventricular

flow-component distribution.

6.5.3 Paediatric datasets

The workflow was also applied to healthy paediatric datasets. Automatic and manual plane placement produced similar mean flow-component distributions. Direct Flow was 31.2% with automatic plane placement and 29.6% with manual plane placement, while Residual Volume was 21.9% and 24.1%, respectively. These results suggest that the workflow can produce plausible group-level output also in smaller hearts, although the small number of datasets means that the results should not be interpreted as paediatric reference values.

For the paediatric analysis, a two-voxel in-plane erosion of the mapped LV seed mask was used before seed placement. Without this erosion, seeds close to the LV boundary produced pathlines that left the ventricle early. Moving the seed region inward reduced this problem, but it also means that the paediatric component fractions were obtained using a seed-mask setting that was not used by default in the adult datasets.

The paediatric results are relevant in relation to previous work on normal paediatric LV flow dynamics assessed with 4D flow MRI [29]. The present analysis was not designed to reproduce or establish normal paediatric ranges, but it shows that the implemented workflow can be applied to paediatric data when image quality, segmentation, seed placement, and plane placement are acceptable.

Paediatric datasets should still be interpreted carefully. Because the heart and vessels are smaller in children, a small absolute error in segmentation, seed placement, or plane placement corresponds to a larger relative error than in adults. Seed-mask erosion and basal-slice removal may also have a larger relative effect in smaller ventricles, since the excluded region represents a larger fraction of the total LV volume.

Overall, the paediatric analysis supports the technical applicability of the workflow in smaller hearts, but it also shows that these datasets are more sensitive to spatial and parameter-related uncertainties. This is especially important when automatic plane placement is used, since fixed search-region and plane-size settings may not scale well across different heart sizes. The effect of seed-mask erosion should therefore be tested further before the paediatric results are interpreted as physiological reference values.

6.5.4 Internal consistency with automatic and manual plane placement

The comparison between automatic and manual plane placement was evaluated using the absolute difference between Retained Inflow and Delayed Ejection Flow, $|\Delta_{\mathrm{RI-DE}}|$. This measure was used as an internal consistency indicator, since Retained Inflow and Delayed Ejection Flow are expected to be similar in magnitude in an approximately cyclic analysis. A low value suggests that inflow and outflow classification are reasonably balanced, while a high value indicates that the analysis should be reviewed more carefully.

At the larger level, automatic and manual plane placement produced similar mean $|\Delta_{\mathrm{RI-DE}}|$

values. In the healthy adult cohort, the mean value was 2.9 percentage points for manual placement and 3.1 percentage points for automatic placement. In the heart-failure cohort, the corresponding values were 6.2 and 6.8 percentage points, while the paediatric cohort showed values of 3.5 and 4.6 percentage points. In these data, automatic plane placement did not substantially increase the average internal-consistency difference in any of the three groups.

The higher $|\Delta_{\mathrm{RI-DE}}|$ values in the heart-failure cohort suggest that these datasets were more challenging overall. This may be related to altered ventricular geometry, more complex intracardiac flow, reduced or less clearly directed outflow, variable image quality, or imperfect agreement between cine MRI and 4D flow MRI. Since the difference was elevated also with manual plane placement, it cannot be explained only by the automatic aortic-plane method.

Similar mean $|\Delta_{\mathrm{RI-DE}}|$ values should not be interpreted as proof that automatic and manual plane placement produce equivalent flow-component distributions. Two analyses may have similar RI–DE balance while still differing in Direct Flow, Retained Inflow, Delayed Ejection Flow, or Residual Volume. The present comparison evaluates average internal consistency, not full agreement between plane-placement methods.

The paediatric comparison should also be interpreted together with the seed-mask setting used for these datasets. Both automatic and manual paediatric analyses used the same two-voxel in-plane erosion of the mapped LV seed mask. The comparison therefore reflects the difference between plane-placement methods with this seed-generation setting and does not evaluate the effect of erosion itself.

Overall, the results support automatic plane placement as a practical starting point, especially when the anatomy and velocity field are clearly visualized. Visual inspection of the analysis planes, seed region, and pathline behaviour remains necessary before interpreting the final flow-component fractions. This is especially important in datasets with smaller anatomy, uncertain cine-to-4D-flow alignment, noisy or aliased velocity data, or an automatically estimated aortic plane that appears anatomically implausible.

6.5.5 Left-ventricular kinetic-energy output

The kinetic-energy output was included as an additional measure of the LV velocity field. It was not used for flow-component classification, but it gives another way to inspect the time-resolved flow behaviour in the LV.

In the results, the HFrEF group had the highest absolute average LV kinetic energy. When the values were normalized by EDV, the HFrEF group was instead lower than several of the other groups. This shows that absolute KE and EDV-normalized KE describe different things. Absolute KE depends on both velocity and LV volume, while EDV-normalized KE reduces the direct effect of ventricular size. This is important when comparing healthy adults, heart-failure patients, and paediatric datasets, since these groups differ in heart size.

Previous 4D-flow MRI studies have used LV kinetic energy to study pre-systolic LV flow organization and changes in heart failure [12, 3, 23, 17]. Reference values and variation with age and sex have also been reported in healthy subjects [22, 28]. In this thesis, KE should be seen as a complementary output, not as a main result.

The KE values should also be interpreted carefully. They depend on the LV segmentation, cine-to-4D-flow alignment, VENC setting, velocity noise, and normalization method. For this reason, the KE results are mainly useful as an additional way to inspect the workflow output, rather than as values that can be directly compared with previous studies.

6.5.6 Right-ventricular proof-of-concept

The RV proof-of-concept analysis showed that the general workflow concept could also be used outside the LV. By replacing the LV seed region with an RV segmentation and adapting the analysis planes to tricuspid inflow and pulmonary outflow, the workflow produced classified RV pathline output in eight datasets with EF above 60%. This supports the flexibility of the implemented seed–pathline–plane-crossing framework.

The RV results should still be interpreted only as a technical feasibility test. The RV has a more complex shape than the LV, with a thinner wall, crescent-shaped cavity, and different inflow and outflow directions. The RV analysis was also not optimized or validated in the same way as the LV workflow. Reliable RV flow-component analysis would require RV-specific seed settings, tricuspid and pulmonary plane handling, visualization tools, review criteria, and validation in a larger dataset.

The RV proof-of-concept can be related to previous 4D-flow MRI work on right-ventricular blood transport. Fredriksson et al. showed that RV blood flow can be separated into functional components related to inflow, outflow, and residence time [30]. This supports the general idea that a seed–pathline–classification framework can be useful also for right-sided ventricular flow.

In the present proof-of-concept analysis, the RV distribution showed 26.8% Direct Flow and 32.1% Residual Volume. These values should not be interpreted as reference values or as validation against Fredriksson et al., since the present RV implementation was adapted from the LV workflow and was not optimized for RV anatomy [30]. However, the ability to generate RV seed points, trace pathlines through the chamber, and classify them into four transport components suggests that the framework could be developed further for RV analysis.

6.5.7 Overall interpretation of the human-data results

Overall, the human-data results show that the workflow can produce interpretable flow-component distributions in healthy adults, heart-failure patients, healthy paediatric datasets, and RV proof-of-concept analyses. The clearest physiological trend was the reduced Direct Flow and increased Residual Volume in the HFrEF group, which agrees with previous heart-failure literature [17].

At the same time, the healthy adult cohort showed lower Direct Flow and higher Residual Volume than typically reported in healthy reference data [11, 12, 22]. This shows that the results are still influenced by cohort composition and implementation choices. The paediatric results also depended on the two-voxel in-plane seed-mask erosion used to improve pathline stability near the LV wall.

The human-data results should therefore be interpreted as a feasibility and review-based evaluation of the implemented workflow, not as diagnostic evidence. The observed failure modes and methodological limitations are discussed further in the next section.

6.6 Limitations and observed failure modes

Several limitations remain in the implemented workflow. Some are general limitations of the method, while others were observed directly during the retrospective human-dataset analysis. The final flow-component percentages depend not only on the classification logic, but also on image quality, velocity-field reliability, seed-region definition, analysis-plane placement, and the agreement between cine MRI and 4D flow MRI.

6.6.1 Velocity-field and acquisition-related limitations

One important limitation is that the pathlines are calculated from measured 4D flow MRI velocity data, not from a perfect representation of blood motion. The data can be affected by spatial resolution, temporal resolution, VENC, signal-to-noise ratio, gating quality, and correction of phase offsets and aliasing [6]. Errors in the velocity field can propagate into the pathlines and then into the flow-component classification.

VENC selection is especially important. If VENC is too low, velocities above the encoded range may wrap and appear with the wrong direction [21, 6]. This is most problematic close to the aortic valve, because forward pathlines and aortic-plane crossings depend directly on the measured outflow direction. Aliasing in the LV outflow tract or aortic valve region may therefore reduce the number of particles classified as leaving the LV, which can affect Direct Flow and Delayed Ejection Flow. Examples of this type of velocity-field error are shown in Figures 6.1a and 6.1b.

The opposite problem may also occur. If VENC is set too high, lower intraventricular velocities are measured with reduced sensitivity and may become noisier [21, 6]. This can make particle tracing less reliable in low-flow regions. In some datasets, pathlines appeared to drift into regions where little or no meaningful flow was visible. An example of this behaviour is shown in Figure 6.1c. These velocity-field problems can also affect the KE calculation, since KE depends directly on the measured velocity magnitude.

6.6.2 Cine-to-4D-flow alignment and seed-region limitations

LV seed generation depends on the spatial agreement between cine images and the 4D flow data. These image stacks are normally acquired separately, so patient motion, breathing, slice-position differences, or differences in cardiac phase matching can affect the transfer of the LV segmentation to the 4D flow space. If this happens, the transferred seed region may not match the LV cavity in the 4D flow data. Pathlines may then appear to start partly outside the LV, even though the seed mask was generated from the cine-derived LV segmentation.

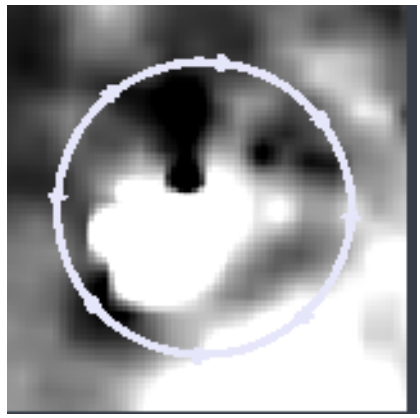

(a) Aliasing in the selected region, where black areas indicate velocities exceeding the selected VENC and wrapping to the opposite flow direction.

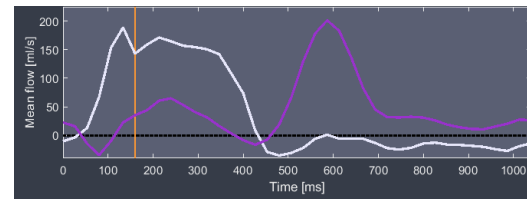

(b) Flow curve showing abrupt direction changes caused by aliasing.

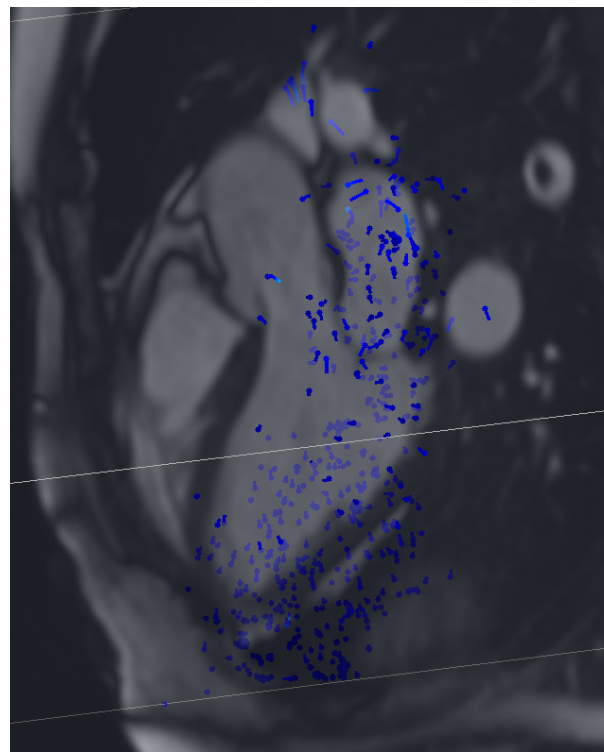

(c) Pathlines drifting into a region with little visible flow.

Figure 6.1: Examples of velocity-field-related failure modes observed during human-dataset analysis. Aliasing can affect the measured direction of high-velocity outflow, while unsuitable VENC settings or noisy low-flow regions can make pathline tracing less reliable.

Figure 6.2a shows an example where cine–4D-flow mismatch caused the pathlines to be displaced relative to the LV anatomy.

To reduce the risk of false crossings close to the valve planes, two basal slices were removed before seed placement. This was used as a practical default setting. It reduces the risk of placing seeds close to the left atrium or valve-plane region, but it also means that the seed volume is not exactly equal to the full anatomical LV end-diastolic volume. Some blood close to the base may therefore be excluded from the analysis.

Seed-mask erosion creates a similar trade-off. Erosion can move seed points away from the endocardial boundary, where segmentation uncertainty, partial-volume effects, limited 4D flow resolution, and small cine-to-4D-flow mismatches may make pathlines less reliable. However, erosion also removes part of the anatomical LV volume from the seed region. Too little erosion may give unstable pathlines near the wall, while too much erosion may exclude relevant blood volume and affect the final component percentages.

This issue was most evident in the paediatric datasets. In these cases, seeds placed close to the LV wall often produced pathlines that rapidly left the ventricle. A two-voxel in-plane erosion of the mapped LV seed mask was therefore used for the paediatric analysis. This improved the visual stability of the pathlines, but it also introduced a dataset-specific seed-region choice. The effect of this choice on the final component percentages was not tested systematically.

The workflow also uses a discrete particle representation of the LV blood volume. The continuous blood volume is represented by a finite number of seed points, and each seed contributes to the final component percentages. If the seed density is too low, smaller flow structures may be missed. If the number of candidate seeds is larger than the maximum allowed number, only a subset is used. This keeps the computation practical, but it also means that some regions may be represented less completely than others.

6.6.3 Analysis-plane and crossing-related limitations

The mitral and aortic analysis planes are simplified representations of moving anatomical structures. The mitral plane is constructed from cine long-axis contour information, while the aortic plane is either estimated from the velocity data or adjusted manually by the user. Both approaches can work well, but neither removes the need for review. The mitral-plane construction may be affected by contour quality, long-axis image orientation, and cine-to-4D-flow alignment. The aortic-plane estimate may be affected by noisy velocity data, unclear outflow jets, unusual anatomy, or nearby high-velocity structures.

The planes are also static, while the mitral and aortic valve regions move during the cardiac cycle. This simplification makes the workflow practical and easier to reproduce, but it may introduce crossing errors close to the valve planes. Since the whole flow-component classification depends on mitral and aortic crossings, uncertainty in either plane can affect the final component percentages.

Automatic aortic-plane placement showed dataset-dependent limitations. In the automatic method, the aortic search starts from an enlarged version of the mitral plane. In some paediatric datasets, this search region covered too large a relative anatomical area because

of the smaller heart size. High-velocity structures outside the intended LV outflow region could then influence the detection. Figure 6.2b shows an example where the automatically estimated aortic plane was placed in the descending aorta instead of the intended aortic outflow location. This shows why automatic aortic-plane placement was more sensitive in paediatric datasets and why visual review remained necessary.

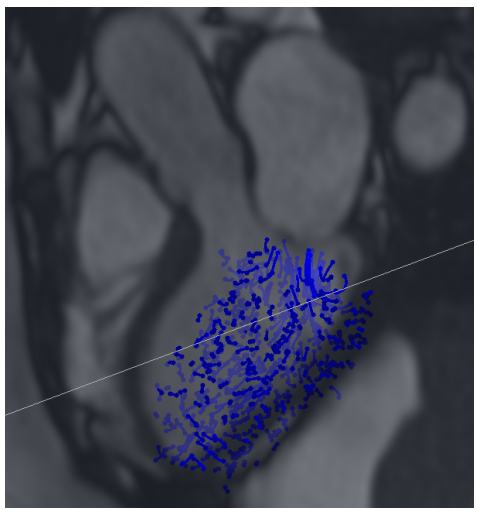

(a) Cine–4D-flow mismatch causing pathline displacement outside the LV.

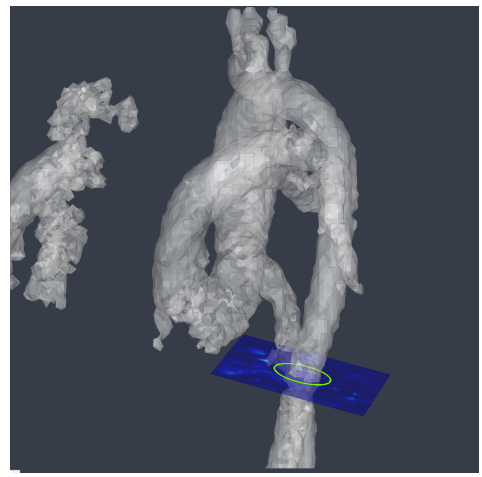

(b) Automatic aortic plane placed in the descending aorta.

Figure 6.2: Examples of geometry- and plane-placement-related failure modes. Spatial mismatch between cine and 4D flow data can affect seed placement, while the automatic aortic-plane search can fail when the search region includes high-velocity structures outside the intended LV outflow location.

The binary crossing rule is another simplification. A particle is counted as crossing a valve plane if its pathline intersects the full plane polygon at least once. Multiple crossings are not counted separately. This makes the classification easier to interpret, but it also means that the result depends on plane size, position, orientation, and the chosen crossing definition. The RI–DE measure was useful as an internal consistency check, but it should not be interpreted as a full validation of the flow-component classification.

These examples show why visual review remains part of the analysis. The final component percentages should be interpreted only after checking LV seed placement, pathline behaviour, velocity-field quality, and analysis-plane positions. In datasets with severe aliasing, poor cine-to-4D-flow agreement, implausible automatic aortic-plane placement, or non-physiological pathline drift, the flow-component output should be treated as technically unreliable even if the algorithm produces a complete numerical result.

6.6.4 Dataset and validation limitations

The available data were retrospective, and only limited clinical information was used in the present analysis. Although the heart-failure datasets included EF values for subgrouping, detailed clinical information such as heart-failure aetiology, medication, comorbidities, symptom status, disease duration, age, sex, and scan-specific acquisition differences was not included. This limits the physiological interpretation of the observed flow-component distributions. Differences between EF groups may reflect real differences in intraventricular flow, but they may also be influenced by cohort composition, image quality, segmentation quality, and post-processing choices.

The retrospective nature of the data also means that the acquisitions were not necessarily optimized for the present workflow. Some datasets had limited image quality, imperfect segmentation, or velocity-field artifacts. Scanner- and protocol-specific reconstruction or correction steps may also influence the measured velocity field, but these effects were not evaluated systematically in the present work [6].

The phantom validation also does not cover all sources of error in patient data. The circular-flow phantom tested numerical integration against a known trajectory, and the three-cylinder phantom tested the plane-crossing and classification logic under controlled conditions. These tests are valuable because they isolate key parts of the method. They do not, however, validate the full workflow under realistic physiological conditions. They do not include complex LV geometry, moving valve planes, turbulent or vortical intracardiac flow, cine-to-4D-flow mismatch, patient motion, limited image quality, clinical variability, uncertainty in contour quality, or uncertainty in automatic aortic-plane detection. Broader validation on human data is needed before the workflow can be considered clinically reliable.

6.6.5 Clinical and ethical considerations

The workflow was developed for research use and should not be used for clinical decision-making without further validation. The human datasets were retrospective and anonymized, and no new patient data were collected for this thesis.

A numerical output can look objective, but the final flow-component percentages can be affected by image quality, segmentation, cine-to-4D-flow alignment, seed placement, and plane placement. The user therefore has to review the visual output before interpreting the results. If the seed region, pathlines, velocity field, or analysis planes look unreliable, the numerical result should not be trusted.

This is especially important in a semi-automatic workflow. The software can support the analysis, but the user is still responsible for checking whether the result is technically reasonable. Until repeatability, user dependence, and external validation have been tested more thoroughly, individual results should not be overinterpreted.

6.7 Future work

Future work should first test how sensitive the results are to the main analysis choices. The final component fractions may be influenced by seed-region definition, basal seed-mask reduction, seed-mask erosion, integration settings, analysis-plane placement, and crossing criteria. A systematic sensitivity analysis should be performed to determine which settings have the largest influence on the reported results. The results could then be used to define a more standardized analysis protocol.

The effect of seed-mask erosion should also be tested systematically. In the present work, a two-voxel in-plane erosion of the mapped LV seed mask was used for the paediatric datasets to reduce early pathline escape from seeds placed close to the LV wall. Future work should compare no erosion, two-voxel erosion, and possibly resolution-based erosion distances to determine how strongly this choice affects the flow-component fractions.

Plane handling should also be developed further. In the current workflow, both planes can be created automatically or placed manually. Manual correction is useful when the automatic result is not anatomically plausible, but it also introduces user dependence. Future development should make automatic plane placement more reliable while still allowing manual adjustment when needed. For the mitral plane, this could include better handling of incomplete or uncertain long-axis contour information. For the aortic plane, the automatic search method should be made more robust in datasets with noisy velocity fields, unclear outflow patterns, unusual geometry, or smaller ventricular dimensions. The paediatric results suggest that fixed search-region and plane-size settings may not work equally well for hearts of different sizes. Future versions could adapt the search region and plane dimensions to the size of the ventricle or detected outflow region.

The erosion setting should also be made available in the Segment user interface. At present, the setting is implemented in the code but is not exposed as a user-adjustable option. Making it visible in the interface would make the analysis easier to document, repeat, and adjust for datasets with different spatial resolution or ventricular size.

Automatic warning criteria could also be added. The software could flag analysis planes that are placed in anatomically implausible positions, seed regions that do not match the expected ventricular cavity, pathlines that show clear non-physiological drift, suspected velocity aliasing, or large differences between Retained Inflow and Delayed Ejection Flow. These warnings would not replace visual inspection, but they could help identify cases that need manual review. They could also reduce the risk that technically unreliable output is interpreted without review.

Repeatability should also be tested by applying the workflow to the same datasets more than once. The same user could repeat the analysis to test intra-observer repeatability. Different users could also repeat the analysis to test inter-observer variability. This would show how strongly the output depends on manual plane placement and review decisions. It would also be useful when evaluating future improvements to automatic plane placement.

Future work could also test whether moving valve planes improve classification compared with the current static planes. Static planes make the workflow practical and easy to repeat,

but they simplify the motion of the mitral and aortic valve regions during the cardiac cycle. Time-varying planes could reduce crossing errors near the valves, but they would also make the analysis more complex and introduce new sources of uncertainty. Their value should be tested by comparing component fractions, repeatability, and visual plausibility with the current static-plane approach.

The kinetic-energy output should also be evaluated further. In the current implementation, LV kinetic energy is calculated and displayed, but it is not used for flow-component classification. Future work could test how sensitive this output is to segmentation, cine-to-4D-flow alignment, VENC, and normalization choices, and whether it gives stable additional information together with the flow-component fractions.

The RV proof-of-concept shows that the general seed–pathline–plane-crossing framework can be adapted beyond the LV. A dedicated RV workflow would require RV-specific seed settings, tricuspid and pulmonary plane handling, visualization tools, and review criteria. It should be treated as a separate methodological extension, not as a direct application of the LV settings.

Overall, future work should focus on standardizing the workflow, improving robustness, and reducing user dependence while keeping visual review as part of the analysis. These developments would make the workflow more robust for future studies of intracardiac flow using 4D flow MRI.

7 Conclusion

Conventional measures such as ejection fraction describe overall LV pump function, but do not show how blood is transported through the ventricle. This thesis therefore focused on implementing a workflow for LV flow-component analysis from 4D flow MRI, where particle paths are used to quantify Direct Flow, Retained Inflow, Delayed Ejection Flow, and Residual Volume based on mitral inflow and aortic outflow during the cardiac cycle.

A semi-automated workflow was developed in Segment that combines cine-derived LV seed generation, forward and backward RK4 pathline tracing, mitral and aortic plane-crossing classification, visual review, and LV kinetic-energy calculation. The phantom tests supported two central parts of the implementation: RK4 closely followed the known circular trajectory, and the classification method gave the expected result in the three-cylinder phantom.

The workflow was applied to retrospective human datasets from healthy adults, patients with heart failure, and healthy children. The clearest physiological trend was seen in the heart-failure analysis, where the HFrEF group had markedly lower Direct Flow and higher Residual Volume than the groups with higher ejection fraction. This agrees with previous findings in heart failure. In the healthy adult cohort, Direct Flow was lower and Residual Volume higher than values commonly reported in previous studies, which suggests that the results are influenced by both dataset characteristics and analysis choices.

In the paediatric datasets, automatic and manual plane placement gave similar mean flow-component distributions. A two-voxel in-plane erosion of the mapped LV seed mask was used for these datasets to improve pathline stability near the LV wall. The paediatric results were therefore plausible, but should not be interpreted as reference values.

Automatic and manual plane placement also gave similar mean RI–DE values at group level. This supports automatic plane placement as a practical starting point, but it does not show that the two approaches give the same flow-component fractions in every dataset. The final output should still be reviewed visually, since seed placement, cine-to-4D-flow alignment, velocity-data quality, and the positions of the mitral and aortic planes can affect the classification. LV kinetic energy was included as an additional output.

The RV proof-of-concept analysis showed that the same basic seed–pathline–plane-crossing approach could be applied using an RV seed region and RV-adapted planes. This part of the work was exploratory, and the RV results should not be interpreted further without separate RV-specific method development.

Overall, the workflow should be viewed as a research method, not as a clinically validated tool. Future work should focus on improving automatic plane handling, testing the sensitivity of the main analysis settings, evaluating repeatability within and between users, making seed-mask erosion available in the user interface, and adding automatic warnings for results that need closer review.

References

[1] J. J. Qin, B. Indja, A. Gholipour, M. Gök, and S. M. Grieve. "Evaluation of Left Ventricular Function Using Four-Dimensional Flow Cardiovascular Magnetic Resonance: A Systematic Review". In: *Journal of Cardiovascular Development and Disease* 9.9 (2022).

[2] V. M. Stoll, A. T. Hess, C. T. Rodgers, et al. "Left Ventricular Flow Analysis: Novel Imaging Biomarkers and Predictors of Exercise Capacity in Heart Failure". In: *Circulation: Cardiovascular Imaging* 12.5 (2019).

[3] E. Svalbring, A. Fredriksson, J. Eriksson, et al. "Altered Diastolic Flow Patterns and Kinetic Energy in Subtle Left Ventricular Remodeling and Dysfunction Detected by 4D Flow MRI". In: *PLOS ONE* 11.8 (2016).

[4] L. Wigström, L. Sjöqvist, B. Wranne, et al. "Particle Trace Visualization of Intracardiac Flow Using Time-Resolved 3D Phase Contrast MRI". In: *Magnetic Resonance in Medicine* 41.4 (1999).

[5] M. Markl, A. Frydrychowicz, S. Kozerke, M. Hope, and O. Wieben. "4D Flow MRI". In: *Journal of Magnetic Resonance Imaging* 36.5 (2012).

[6] M. M. Bissell, F. Raimondi, L. Ait Ali, et al. "4D Flow Cardiovascular Magnetic Resonance Consensus Statement: 2023 Update". In: *Journal of Cardiovascular Magnetic Resonance* 25.1 (2023).

[7] A. Demirkiran, P. van Ooij, J. J. M. Westenberg, et al. "Clinical Intra-Cardiac 4D Flow CMR: Acquisition, Analysis, and Clinical Applications". In: *European Heart Journal - Cardiovascular Imaging* 23.2 (2022).

[8] Z. A. Z. Ashkir, S. M. Myerson, S. Neubauer, C. J. Carlhäll, T. Ebbers, and B. Raman. "Four-Dimensional Flow Cardiac Magnetic Resonance Assessment of Left Ventricular Diastolic Function". In: *Frontiers in Cardiovascular Medicine* 9 (2022).

[9] P. R. Roos, F. M. Rijnberg, J. J. M. Westenberg, and H. J. Lamb. "Particle Tracing Based on 4D Flow Magnetic Resonance Imaging: A Systematic Review into Methods, Applications, and Current Developments". In: *Journal of Magnetic Resonance Imaging* 57.5 (2023).

[10] A. F. Bolger, E. Heiberg, M. Karlsson, et al. "Transit of Blood Flow Through the Human Left Ventricle Mapped by Cardiovascular Magnetic Resonance". In: *Journal of Cardiovascular Magnetic Resonance* 9.5 (2007).

[11] J. Eriksson, C. J. Carlhäll, P. Dyverfeldt, J. Engvall, A. F. Bolger, and T. Ebbers. "Semi-Automatic Quantification of 4D Left Ventricular Blood Flow". In: *Journal of Cardiovascular Magnetic Resonance* 12.1 (2010).

[12] J. Eriksson, P. Dyverfeldt, J. Engvall, A. F. Bolger, T. Ebbers, and C. J. Carlhäll. "Quantification of Presystolic Blood Flow Organization and Energetics in the Human Left Ventricle". In: *American Journal of Physiology - Heart and Circulatory Physiology* 300.6 (2011).

[13] E. Heiberg, J. Sjögren, M. Ugander, M. Carlsson, H. Engblom, and H. Arheden. "Design and Validation of Segment—Freely Available Software for Cardiovascular Image Analysis". In: *BMC Medical Imaging* 10 (2010).

[14] J. G. Betts, K. A. Young, J. A. Wise, et al. *Anatomy and Physiology 2e*. Accessed: 2026-05-21. Houston, Texas: OpenStax, 2022.

[15] Servier Medical Art. *Heart Valves*. `https://smart.servier.com`. Licensed under Creative Commons Attribution 4.0 International (CC BY 4.0). Accessed: 2026-05-26. 2026.

[16] P. A. Heidenreich, B. Bozkurt, D. Aguilar, et al. "2022 AHA/ACC/HFSA Guideline for the Management of Heart Failure". In: *Circulation* 145.18 (2022).

[17] H. W. Wong, H. Wang, C. T. Kwan, et al. "Cardiovascular magnetic resonance left ventricular 4D-flow: differences in flow components and kinetic energy across heart failure spectrum". In: *European Heart Journal - Imaging Methods and Practice* 3.1 (2025).

[18] V. P. B. Grover, J. M. Tognarelli, M. M. E. Crossey, I. J. Cox, S. D. Taylor-Robinson, and M. J. W. McPhail. "Magnetic Resonance Imaging: Principles and Techniques: Lessons for Clinicians". In: *Journal of Clinical and Experimental Hepatology* 5.3 (2015).

[19] R. A. Pooley. "Fundamental Physics of MR Imaging". In: *RadioGraphics* 25.4 (2005).

[20] J. Schulz-Menger, D. A. Bluemke, J. Bremerich, et al. "Standardized Image Interpretation and Post-Processing in Cardiovascular Magnetic Resonance – 2020 Update". In: *Journal of Cardiovascular Magnetic Resonance* 22.1 (2020).

[21] K. S. Nayak, J.-F. Nielsen, M. A. Bernstein, et al. "Cardiovascular Magnetic Resonance Phase Contrast Imaging". In: *Journal of Cardiovascular Magnetic Resonance* 17.1 (2015).

[22] V. M. Stoll, M. Loudon, J. Eriksson, et al. "Test-Retest Variability of Left Ventricular 4D Flow Cardiovascular Magnetic Resonance Measurements in Healthy Subjects". In: *Journal of Cardiovascular Magnetic Resonance* 20.1 (2018).

[23] M. Kanski, P. M. Arvidsson, J. Töger, et al. "Left Ventricular Fluid Kinetic Energy Time Curves in Heart Failure from Cardiovascular Magnetic Resonance 4D Flow Data". In: *Journal of Cardiovascular Magnetic Resonance* 17.1 (2015).

[24] H. Kaur, H. Assadi, S. Alabed, et al. "Left Ventricular Blood Flow Kinetic Energy Assessment by 4D Flow Cardiovascular Magnetic Resonance: A Systematic Review of the Clinical Relevance". In: *Journal of Cardiovascular Development and Disease* 7.3 (2020).

[25] A. Elhawaz, G. T. Archer, H. Zafar, et al. “Left Ventricular Blood Flow Kinetic Energy Is Associated with the Six-Minute Walk Test and Left Ventricular Remodelling Post Valvular Intervention in Aortic Stenosis”. In: *Quantitative Imaging in Medicine and Surgery* 11.4 (2021).

[26] M. Kanski, J. Töger, K. Steding-Ehrenborg, et al. “Whole-heart four-dimensional flow can be acquired with preserved quality without respiratory gating, facilitating clinical use: a head-to-head comparison”. In: *BMC Medical Imaging* 15.1 (2015).

[27] P. J. de Koning, R. J. van der Geest, and J. J. M. Westenberg. “Objective Method for Assessment of Reliability of Particle Tracing Visualization in 4D Flow MRI”. In: *Journal of Cardiovascular Magnetic Resonance* 15.Suppl 1 (2013).

[28] X. Zhao et al. “Age- and Sex-Specific Reference Values of Biventricular Flow Components and Kinetic Energy by 4D Flow Cardiovascular Magnetic Resonance in Healthy Subjects”. In: *Journal of Cardiovascular Magnetic Resonance* 25.1 (2023).

[29] F. M. Callaghan, B. Burkhardt, J. Geiger, E. Valsangiacomo Buechel, and C. Kellenberger. “Normal left ventricular flow dynamics in a paediatric population assessed by 4D Flow MRI”. In: *Proceedings of the International Society for Magnetic Resonance in Medicine*. Vol. 31. 2023.

[30] A. G. Fredriksson, J. Zajac, J. Eriksson, et al. “4-D Blood Flow in the Human Right Ventricle”. In: *American Journal of Physiology - Heart and Circulatory Physiology* 301.6 (2011).