Chapter 7

Estimation of Optimal Inlet Boundary Conditions for Blood Flow Assessment in Abdominal Aortic Aneurysm Using Variational Data Assimilation

Sara Paratico, Riccardo Munafò, Chiara Trenti, Petter Dyverfeldt, Simone Saitta, and Emiliano Votta

Abstract

Blood fluid dynamics impacts vessel wall cells and tissue biomechanics, influencing thrombus formation and vessel wall remodelling. Accurate in vivo quantification can thus aid in understanding these mechanisms and patient stratification. Computational fluid dynamics (CFD) and 4D flow magnetic resonance imaging (MRI) are both used for this but have limitations: CFD involves assumptions and boundary condition simplifications, while 4D flow MRI suffers from low spatial resolution and noise. This study employs variational data assimilation to integrate CFD and 4D flow MRI, yielding a high-resolution, noise-free flow field closely aligned with 4D flow MRI velocity data. To enhance alignment, the optimal inlet velocity profile is determined iteratively via an incremental pressure correction scheme. Previously tested

Sara Paratico e-mail: sara.paratico@mail.polimi.it
Politecnico di Milano, Milan, Italy

Riccardo Munafò
Politecnico di Milano, Milan, Italy

Chiara Trenti
Division of Cardiovascular Medicine, Department of Health, Medicine and Caring Sciences, Linköping University, Universitetssjukhuset, 581 83 Linköping, Sweden and
Center for Medical Image Science and Visualization (CMIV), Linköping University, Universitetssjukhuset, 581 83 Linköping, Sweden

Petter Dyverfeldt
Division of Cardiovascular Medicine, Department of Health, Medicine and Caring Sciences, Linköping University, Universitetssjukhuset, 581 83 Linköping, Sweden and
Center for Medical Image Science and Visualization (CMIV), Linköping University, Universitetssjukhuset, 581 83 Linköping, Sweden

Simone Saitta
Politecnico di Milano, Department of Electronics, Information, and Bioengineering, Milan, Italy

Emiliano Votta
Politecnico di Milano, Department of Electronics, Information, and Bioengineering, Milan, Italy

J. S. Dokken et al. (eds.), *The FEniCS Project*,
Simula SpringerBriefs on Computing 19,
https://doi.org/10.1007/978-3-032-17396-6_7

in simple synthetic geometries and later in a complex patient-specific abdominal aortic aneurysm, this approach demonstrates improved reliability in patient-specific haemodynamic evaluation.

7.1 Introduction

Alterations in blood fluid dynamics often contribute to the progress of cardiovascular pathological conditions (Bappoo et al., 2021; Guzzardi et al., 2015). Hence, quantifying blood fluid dynamics on a patient-specific basis and non-invasively can improve the understanding of pathological mechanisms or the stratification of patients based on the risk for adverse endpoints. With this aim, blood flow fields can be reconstructed from clinical imaging, namely, 4D flow magnetic resonance imaging (MRI) (Dyverfeldt et al., 2015), or computed through patient-specific computational fluid dynamics (CFD) models (Kheyfets et al., 2015). However, 4D flow MRI provides indirect and noisy velocity measurements with low spatiotemporal resolution that often violate mass conservation. CFD models solve discretized Navier–Stokes (NS) equations to compute well-resolved, noise-free velocity fields, but they are affected by numerical artefacts and rely on simplified boundary conditions (BCs), including inlet velocity BCs. Variational data assimilation (VarDA) integrates CFD-based NS equations with sparse, uncertain 4D flow MRI data by optimising BCs to minimise discrepancies. In cardiovascular flows, it refines inlet velocity profiles but requires computing both velocity and pressure gradient fields, which is challenging in high-velocity arterial flows. Pressure–velocity coupling or correction schemes address this issue, but MRI-induced noise can hinder proper pressure correction, affecting the accuracy of the solution.

7.1.1 Related Works

Several studies have explored VarDA in haemodynamics, ranging from 2D steady-state to 3D transient conditions. D'Elia et al. (2012) have shown that VarDA allows flow fields to be reconstructed in geometrically complex 2D domains, such as the 2D representation of the aortic arch and carotid bifurcation, even with noisy velocity data. Subsequently, in D'Elia and Veneziani (2013), the same authors have shown that, in 2D domains, noisy velocity data can be effectively managed by properly managing inlet BCs. In particular, they show that the regularization of the inlet velocity profile through the use of a control variable also regularizes the velocity field over the whole domain and allows for successful pressure–velocity coupling. Tiago et al. (2017) have extended VarDA to a 3D saccular aneurysm, demonstrating its flexibility with various BCs and optimisation methods such as gradient-based and genetic algorithms to improve accuracy. Koltukluoğlu and Blanco (2018) show that VarDA applied to 4D flow MRI data outperforms traditional CFD methods by dynamically ad-

justing BCs in real time to maintain flow congruence near inlets. Funke et al. (2019) demonstrate the effectiveness of 4D (3D space + time) VarDA in capturing transient blood flow in aneurysms by using phase contrast MRI data. Finally, Dokken et al. (2020) propose a multimesh finite element method that enhances stability and accuracy by allowing flexible BC management across multiple mesh domains, which is key for simulating complex haemodynamics in realistic geometries.

7.1.2 Our Goal

This study aims to implement a method to compute in vivo blood fluid dynamics on a patient-specific basis with high spatial resolution without simplifications on the inlet BCs. To achieve this, VarDA is used to estimate an optimal inlet BC for CFD, starting from a noisy, uncertain 4D flow MRI-based velocity profile while enforcing consistency between the CFD-computed velocity field and sparse 4D flow MRI data in the bulk flow region. We benchmarked the method on ideal 2D and 3D geometries and then applied it to a patient-specific abdominal aortic aneurysm (AAA) geometry.

7.2 Methods

7.2.1 Data Assimilation Method

The VarDA approach was formulated as an optimisation problem constrained by the NS equations, using the dolfin-adjoint library for the adjoint problem. The process follows three steps: running a first numerical simulation with tentative inlet BCs (which we refer to as the *forward model* or *tape*), solving the optimisation problem to identify inlet BCs, and running a final numerical simulation yielding the refined velocity and pressure fields (Fig. 7.1).

7.2.2 Forward Problem Definition

The weak and discretized form of NS equations for an incompressible fluid (Stokes, 2009) has been solved using the finite element platform FEniCS (Alnæs et al., 2015) to compute the velocity field $\mathbf{u}$ over a domain Ω, given an initial condition (IC), defined as $\mathbf{u} = \mathbf{u_0}$ on Ω at time t_0, a zero pressure condition at the outlet section Ω_N, and a Dirichlet BC at the inlet section Ω_D in the form of a space- and time-dependent velocity profile $\mathbf{g}$. Through an in-house Python script, 2D and 3D fluid domains Ω were discretized into triangular and tetrahedral elements, respectively, with 1- to 1.5-mm characteristic sizes and linear and quadratic shape functions for nodal pressure and velocity, respectively. The semi-implicit Crank–Nicolson time integration scheme

was applied with a time increment of $\Delta t = 0.001$ s. The incremental pressure correction scheme (IPCS) proposed in (Goda, 1979) was implemented. A generalized minimal residual method was chosen as a linear solver, with tolerances of 1×10^{-4} for momentum and continuity equations.

7.2.3 Optimisation Problem Definition

The optimisation problem (7.1), constrained by the NS equations, aims to minimise a functional $J(\mathbf{u})$ (7.2), defined as the difference between the computed and observed velocities:

$$\min_{\mathbf{c}} J(\mathbf{u}) + R(\mathbf{c}) \quad \text{s.t.} \quad F(\mathbf{u}, \mathbf{c}) = 0, \tag{7.1}$$

$$J(\mathbf{u}) = \|\mathbf{u} - \mathbf{u}_{\text{obs}}\|_{L^2(\Omega)}. \tag{7.2}$$

To address the ill-posedness of the problem, a Tikhonov regularization term $\mathbf{R}(\mathbf{c})$ is introduced with respect to the controlled variable defined as c. It accounts for two terms that are scaled by parameters α and β, where β is set to zero for steady-state conditions. This reformulation transforms the problem into an unconstrained optimisation scenario, which is more suitable for gradient descent methods:

$$\begin{aligned} R(\mathbf{c}) &= \|\mathbf{c}\|_{L^2(\Omega)}, \\ &\text{where} \\ \|c\|_{\Gamma\times(0,T]} &= \left(\int_0^T \int_\Omega \frac{\alpha}{2} \left(\left(|\mathbf{g}_D|^2 + |\nabla \mathbf{g}_D|^2 \right) + \frac{\beta}{2} \left(|\dot{\mathbf{g}}_D|^2 + |(\nabla \mathbf{g})_D|^2 \right) \right) dx\, dt \right)^{\frac{1}{2}}. \end{aligned} \tag{7.3}$$

The adjoint approach efficiently computes the total derivative of the functional, yielding the adjoint NS equations that facilitate optimisation. The iterative Broyden–Fletcher–Goldfarb–Shanno (BFGS) algorithm, in its L-BFGS variant (Liu and Nocedal, 1989), serves as an optimiser. The L-BFGS algorithm is already implemented in the SciPy library, which is automatically called by the dolfin-adjoint library and provides many user-friendly numerical routines, such as the routine for optimisation.

Convergence was ensured through Wolfe conditions (Nocedal and Wright, 2006), with a maximum of 10 iterations and ftol = 1×10^{-9} and gtol = 1×10^{-12} as tolerances. The performance of the method was assessed by the $J(\mathbf{u})$ values before and after optimisation, the root mean squared error (RMSE) between $\mathbf{u}$ and $\mathbf{u}_{\text{obs}}$, and qualitative analysis of the effect on the velocity field through the software ParaView.

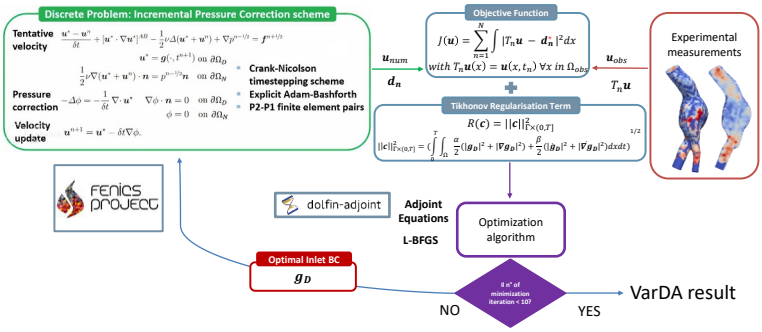

Fig. 7.1: Overview of the data assimilation workflow used in this study. Top left: The finite element solver computes the NS equations with an initial guess for the inlet BCs. Top right: Experimental velocity measurements are taken at discrete points in the domain. Centre: Discrepancy between the CFD velocity field and experimental data is minimised by iteratively refining inlet velocity profile with a gradient-based method.

7.2.4 Benchmark Tests

Preliminary Tests. First, preliminary tests were performed to compare the computational efficiency of the IPCS versus an alternative coupled scheme (Figueroa et al., 2006) in a 2D straight conduit (which represents a case of 2D VarDA and can thus be addressed using the term we define as *2DVar*). Simulations were run under laminar conditions at both low and high Reynolds numbers (Re) to evaluate the method in laminar and transitionally unstable regimes. The conduit was a longitudinal section of a cylinder with a diameter of 41 mm and a length of 200 mm, made of 4,967 mesh elements. Synthetic observations ($\mathbf{u}_{\text{obs}}$) were generated by an auxiliary CFD simulation, prescribing a parabolic velocity profile at the inlet with peak velocity $U_{\text{max}} = 600\,\text{mm/s}$ (Re = 6,649) and with $U_{\text{max}} = 50\,\text{mm/s}$ (Re = 554) for turbulent and laminar conditions, respectively. In the tape, the tentative inlet velocity profile was parabolic with $U_{\text{max}} = 800\,\text{mm/s}$ (Re = 8,865) and $U_{\text{max}} = 100\,\text{mm/s}$ (Re = 1,108). Iterative minimisation of the discrepancy between $\mathbf{u}$ and $\mathbf{u}_{\text{obs}}$ was carried out to determine the optimal velocity profile for CFD simulations, verifying that it matches the parabolic profile used to generate the synthetic observations.

Progressively Demanding Tests. Next, the method was benchmarked through three progressively more demanding tests:

1. A 2D straight conduit in transient conditions (which represents a case of 2D VarDA also involving time and can thus be addressed using the term we define as the *2DVar+t* benchmark). This benchmark shared the same domain as the 2DVar

benchmark. However, both the auxiliary simulation for the generation of the experimental observations and the tape consisted of a sequence of two transient simulations: in the first simulation, velocity was initially equal to 0 mm/s everywhere in the domain, and at the inlet a parabolic velocity profile was imposed whose peak velocity increased linearly from zero to $\frac{U_{\max}}{2}$ over 0.3 s. In the second simulation, the velocity field computed by the first simulation was used as the IC, and the inlet velocity parabolic profile was scaled by the time-dependent function $f(t)$:

$$f(t) = \begin{cases} \frac{U_{\max}}{2} \cos\left(\frac{\pi}{T_s}\left(t - \frac{T_s}{2}\right)\right), & \text{if } t \le T_s, \\ \frac{U_{\max}}{2}, & \text{if } T_s < t \le T_d, \end{cases} \tag{7.4}$$

where $T_s = 300$ ms and $T_d = 540$ ms are the cardiac cycle's systolic and diastolic phases, respectively (Katz, 1977). Besides determining the optimal velocity profile for CFD simulations, spatial and temporal regularization terms, as in (7.3), were incorporated into the optimisation process and subjected to a sensitivity analysis.

2. A 3D straight conduit under steady-state conditions (which represents a case of 3D VarDA and can thus be addressed using the term we define as the *3DVar* benchmark). This benchmark evaluated the computational cost increase when transitioning from a 2D to a 3D problem. The fluid domain was a 3D cylinder with a radius of 30 mm and a length of 200 mm, consisting of 74,968 mesh elements. The IC and BCs, as well as the objective function, were identical to those in the 2DVar benchmark.

3. Patient-specific AAA geometry under steady-state conditions (which represents a case of 3DVar applied to a patient-specific AAA geometry and can thus be addressed using the term we define as the AAA benchmark). This benchmark aimed to test VarDA in a complex 3D domain using real experimental observations. 4D flow imaging data were acquired from an adult male with AAA using a 3T coronary magnetic resonance system (Ingenia, Philips Healthcare, Netherlands) at Linköping University Hospital. The 4D flow data were processed with in-house Python (Saitta et al., 2024), and coronary magnetic resonance angiography was performed for 3D AAA geometry segmentation. Two tests were carried out with laminar flow in the AAA. In the first test, observations consisted of 4D flow data acquired during early systole, corresponding to the third time frame (about 63 ms from the start of the cardiac cycle, with a 21-ms time step), while the tape was generated by CFD simulation fed by 4D flow-based inlet velocity profiles. The second test assessed the method's robustness to noise, using an inlet velocity profile scaled by 0.15 at peak systole to produce the tape's output. Noisy observations were generated by processing the tape's output according to the medium noise setting of Saitta et al. (2024). In addition to metrics mentioned, wall shear stresses (WSSs) from the final simulation were analysed using custom ParaView filters.

The associated codes can be found at https://github.com/saraparatico/proceedingsCodes/tree/main.

Results

7.2.5 Computational Costs

Numerical experiments were conducted on various setups: a workstation with 24 CPUs and 64 GB RAM for the 2DVar and 3DVar benchmarks and a high-performance computing system with 40 CPUs and 190 GB RAM for the 2DVar+t and AAA benchmarks. The 2DVar benchmark took 30 minutes, and the 3DVar benchmark took 6 hours on 12 CPUs; on the other hand, 6 hours were required for the 2DVar+t benchmark and 17 hours for the AAA benchmark.

7.2.6 Preliminary Tests

In high–Reynolds number tests, IPCS optimisation reduced the RMSE from 142.60 mm/s to 6.70 mm/s, while the coupled scheme faced convergence issues. Under low-Reynolds number conditions, the IPCS proved to be five times faster than the coupled scheme and achieved a final RMSE of 1.76 mm/s, compared to 4.22 mm/s for the coupled scheme.

7.2.7 Progressively Demanding Tests

2DVar+t Benchmark. When a zero-velocity field was imposed as an IC, the post-optimisation velocity field showed inconsistencies with respect to the observations. In particular, a high-velocity region just downstream of the inlet section was obtained, while low velocity values were computed in the rest of the domain. Moreover, these tests did not yield improvements from changing α, and increasing β further worsened the performance (Fig. 7.2).

When the initial velocity and pressure fields were set equal to those obtained from the previous post-optimisation simulation, the results showed a more homogeneous flow that better matched parabolic characteristics and had lower $J + R$ values.

3DVar Benchmark. VarDA was performed with spatial regularization terms set to $\alpha = 10^{-2}$, 10^{1}, and 10^{3}. The lowest value of $J+R$ was achieved with $\alpha = 10^{-2}$, but it did not correspond to the lowest RMSE. The velocity field exhibited a peak near the inlet, suggesting continuity loss. The lowest RMSE was achieved for $\alpha = 10^{1}$, with

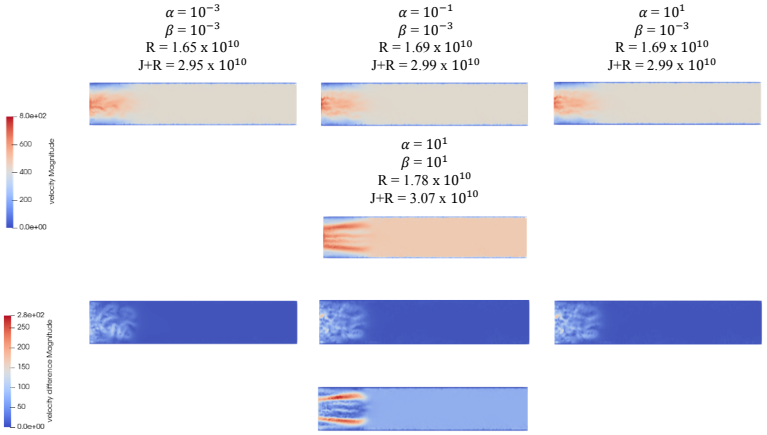

Fig. 7.2: Sensitivity analysis of the regularization parameters at the systolic peak. First row: The 2DVar+t velocity magnitude for different α values (10^{-3}, 10^{-1}, 10^{1}), with $\beta = 10^{-3}$. Second row: A test with $\beta = 10^{1}$ to assess time regularization. Third and fourth rows: The velocity difference between the reference and 3DVar results for each α and β.

the post-optimisation velocity field more accurately reflecting the observed data. Increasing α to 10^3 resulted in significant deviations from the observations, with unexpected velocity behaviours and higher values of $J + R$ and the RMSE. This suggests that lower α values improve the RMSE, while higher α values lead to smoother solutions but can introduce inaccuracies when too large.

AAA Benchmark. In first tests, the RMSE improved from 59.3 mm/s to 55.1 mm/s, indicating better alignment with observations (Fig. 7.3a).

Generating a tape took about 25 minutes, while optimisation required 17 hours with 80 Gb of memory. The WSS distributions from the tape's output and the 3DVar predictions were consistent, identifying regions with high shear stress (Fig. 7.3b). WSS distributions from the tape's output and the 3DVar predictions were consistent in terms of the location of high-WSS regions. Moreover, while enforcing consistency with 4D flow-based velocity measurements, the 3DVar method yielded a regular and realistic WSS distribution. This is a major difference compared to the WSS distribution estimated directly from 4D flow data, which was unrealistic owing to their poor spatial resolution and the impact of noise in the near-wall region. In tests with noisy observations the 3DVar benchmark effectively reconstructed the velocity field, slightly reducing the RMSE from 36.8 mm/s to 36.4 mm/s while maintaining WSS predictions consistent with CFD results, particularly at the iliac bifurcation, which

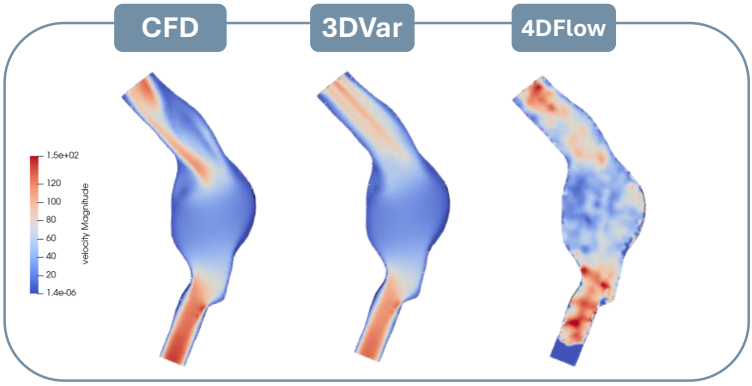

(a) AAA benchmark velocity magnitude maps

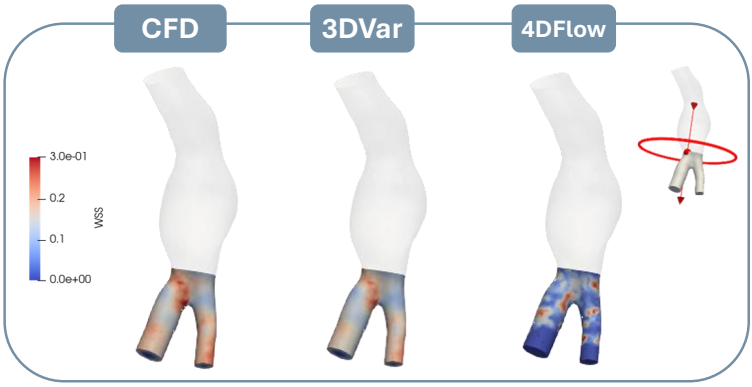

(b) AAA benchmark WSS maps

Fig. 7.3: (a) Velocity and (b) WSS fields obtained on the AAA computed by CFD (left), derived directly from 4D flow MRI (right), and computed by data assimilation (centre).

is where the abdominal aorta splits into two smaller arteries, carrying blood to the pelvis and legs.

7.3 Discussion

7.3.1 From 2DVar+t to 3DVar

The 2DVar+t benchmark reveals challenges due to inertial effects and short simulation durations, causing reconstruction defects from incomplete flow development. Extending the simulation time for optimisation is impractical due to high computational costs. A potential solution includes proper initialization of CFD simulations and implementing a checkpointing method to reduce computational costs by using only the last cardiac cycle for gradient calculations. Moreover, a key difference between the 2DVar+t and 3DVar benchmarks is the flow field's response to regularization. In the 2DVar+t case, increasing α has minimal effect due to dominant time-dependent effects, reducing the impact of spatial regularization. Additionally, increasing β deteriorates the results, a challenge that does not arise in the 3DVar case, emphasizing the difficulty of balancing temporal and spatial regularization in dynamic flows. Conversely, in the 3DVar case, moderate α values (10^1) significantly improve the velocity field, reducing inlet peaks and lowering the RMSE. However, excessive regularization ($\alpha = 10^3$) causes unrealistic velocity patterns.

7.3.2 AAA Benchmark

The AAA benchmark effectively reconstructs the velocity field in the AAA geometry, maintaining high consistency with data obtained from 4D flow imaging. It identifies regions of high shear stress despite challenges due to the lower resolution of 4D flow data near boundaries. The method remains robust to noise.

7.4 Conclusions

This study applies VarDA to estimate optimal inlet BCs for CFD using noisy 4D flow MRI data, minimising mismatches with in vivo velocity measurements. The method yields a resolved, noise-free velocity field and has been validated on 2D, 3D, and patient-specific AAA geometries, demonstrating potential for personalised haemodynamic simulations. The IPCS framework enhances efficiency and accuracy in transient flow analyses. Despite challenges in transient cases, this work lays the groundwork for future VarDA advancements with clinical implications.

References

Alnæs MS, Blechta J, Hake JE, Johansson A, Kehlet B, Logg A, Richardson C, Ring J, Rognes ME, Wells GN (2015) The FEniCS Project Version 1.5. Archive of Numerical Software 3, doi:10.11588/ans.2015.100.20553

Bappoo N, Syed M, Khinsoe G, Kelsey L, Forsythe R, Powell J, Hoskins P, McBride O, Norman P, Jansen S, Newby D, Doyle B (2021) Low Shear Stress at Baseline Predicts Expansion and Aneurysm-Related Events in Patients With Abdominal Aortic Aneurysm. Circulation: Cardiovascular Imaging 14(12):1112–1121, doi:10.1161/CIRCIMAGING.121.013160

D'Elia M, Veneziani A (2013) Uncertainty quantification for data assimilation in a steady incompressible Navier-Stokes problem. ESAIM: Mathematical Modelling and Numerical Analysis 47(4):1037–1057, doi:10.1051/m2an/2012056

D'Elia M, Perego M, Veneziani A (2012) A variational data assimilation procedure for the incompressible Navier-Stokes equations in hemodynamics. Journal of Scientific Computing 52:340–359, doi:10.1007/s10915-011-9547-6

Dokken JS, Johansson A, Massing A, Funke SW (2020) A multimesh finite element method for the Navier-Stokes equations based on projection methods. Computer Methods in Applied Mechanics and Engineering 368:113129, doi:10.1016/j.cma.2020.113129

Dyverfeldt P, Bissell M, Barker AJea (2015) 4D flow cardiovascular magnetic resonance consensus statement. Journal of Cardiovascular Magnetic Resonance 17(1):72, doi:10.1186/s12968-015-0174-5

Figueroa CA, Vignon-Clementel IE, Jansen KE, Hughes TJ, Taylor CA (2006) A coupled momentum method for modeling blood flow in three-dimensional deformable arteries. Computer Methods in Applied Mechanics and Engineering 195(41):5685–5706, doi:10.1016/j.cma.2005.11.011

Funke SW, Nordaas M, Øyvind Evju, Alnæs MS, Mardal KA (2019) Variational data assimilation for transient blood flow simulations: Cerebral aneurysms as an illustrative example. International Journal for Numerical Methods in Biomedical Engineering 35(1):e3152, doi:10.1002/cnm.3152

Goda K (1979) A multistep technique with implicit difference schemes for calculating two- or three-dimensional cavity flows. Journal of Computational Physics 30(1):76–95, doi:10.1016/0021-9991(79)90088-3

Guzzardi D, Barker A, Van Ooij P, Malaisrie S, Puthumana J, Belke D, Mewhort H, Svystonyuk D, Kang S, Verma S, Collins J, Carr J, Bonow R, Markl M, Thomas J, McCarthy P, Fedak P (2015) Valve-related hemodynamics mediate human bicuspid aortopathy: Insights from wall shear stress mapping. Journal of the American College of Cardiology 66(8):892–900, doi:10.1016/j.jacc.2015.06.1310

Katz AM (1977) Physiology of the Heart. Lippincott Williams & Wilkins

Kheyfets V, Rios L, Smith T, Schroeder T, Mueller J, Murali S, Lasorda D, Zikos A, Spotti J, ReillyJr J, Finol E (2015) Patient-specific computational modeling of blood flow in the pulmonary arterial circulation. Computer Methods and Programs in Biomedicine 120(2):88–101, doi:10.1016/j.cmpb.2015.04.005

Koltukluoğlu TS, Blanco PJ (2018) Boundary control in computational haemodynamics. Journal of Fluid Mechanics 847:329–364, doi:10.1017/jfm.2018.329

Liu DC, Nocedal J (1989) On the limited-memory BFGS method for large scale optimization. Mathematical Programming 45:503–528, doi:10.1007/BF01589116

Nocedal J, Wright SJ (2006) Numerical Optimization, 2nd edn. Springer, doi:10.1007/978-0-387-40065-5

Saitta S, Carioni M, Mukherjee S, Schönlieb CB, Redaelli A (2024) Implicit neural representations for unsupervised super-resolution and denoising of 4D flow MRI. Computer Methods and Programs in Biomedicine 246:108057, doi:10.1016/J.CMPB.2024.108057

Stokes GG (2009) On the Theories of the Internal Friction of Fluids in Motion, and of the Equilibrium and Motion of Elastic Solids, Cambridge University Press, pp 75–129. Cambridge Library Collection - Mathematics, doi:10.1017/CBO9780511702242.005

Tiago J, Guerra T, Sequeira A (2017) A velocity tracking approach for the data assimilation problem in blood flow simulations. International Journal for Numerical Methods in Biomedical Engineering 33(10):e2856, doi:10.1002/cnm.2856