Uncertainty-Aware Modeling of Hemolysis in Rotary Blood Pumps Using Reduced-Order Models and Bayesian Parameter Estimation

Unsicherheitsbewusste Modellierung der Hämolyse in rotatorischen Blutpumpen mittels Surrogatmodellen und Bayesscher Parameterschätzung

Von der Fakultät für Maschinenwesen der Rheinisch-Westfälischen Technischen Hochschule Aachen zur Erlangung des akademischen Grades eines Doktors der Ingenieurwissenschaften genehmigte Dissertation

vorgelegt von

Christopher Blum

Berichter: Univ.-Prof. Dr.-Ing. Ulrich Steinseifer
Univ.-Prof. Marek Behr, Ph. D.

Tag der mündlichen Prüfung: 02. Februar 2026

Acknowledgments

This dissertation was completed during my time at the Institute of Applied Medical Engineering in the Department of Cardiovascular Engineering. I am grateful for the stimulating scientific environment provided by the institute, which offered both the freedom to explore new ideas and the structure needed to pursue them in a systematic manner.
I would like to express my sincere gratitude to Prof. Ulrich Steinseifer for his valuable feedback, continuous support, and for fostering a team-oriented research environment that encouraged independent work. The combination of scientific freedom, trust, and constructive discussion made every day at work feel less like an obligation and more like the pursuit of a genuine interest. I am also deeply thankful to Dr. Michael Neidlin, who played a central role in shaping me as a scientist. He taught me how to think critically, how to approach scientific problems systematically, and how to act with integrity in research. Beyond providing guidance on specific projects, he invested substantial time and effort in my development, for which I am very grateful.
I thank Prof. Marek Behr for serving as the second referee of this dissertation and Prof. Jørgen Magnus for chairing the defense committee. I appreciate their time, interest in my work, and their valuable contributions to the evaluation process.
My sincere thanks also go to my colleagues at the Department of Cardiovascular Engineering. The many scientific discussions and shared experiences inside and outside the work environment fostered a supportive and motivating atmosphere. Being part of a team where rigorous science and collegiality went hand in hand greatly enriched my doctoral journey and made everyday work truly enjoyable. I am also particularly grateful for the excellent research environment. Access to well-equipped laboratories, high-performance computing clusters, and specialized software, together with this consistently collegial and supportive atmosphere, was essential to the successful completion of this work.
More broadly, this work builds on the insights and efforts of many before me. It benefits not only from decades of scientific progress, but also from the accelerating development of research tools that continue to transform how we explore, analyze, and communicate knowledge. The pace of innovation in recent years has been extraordinary, and it has played a substantial role in shaping the path and process of this project.
Finally, I would like to thank my family for their steady support, patience, and encouragement throughout this period, which provided a constant source of motivation, especially during challenging phases. I am also deeply grateful to my partner and to those closest to me for their understanding, encouragement, and constant support during this journey.

Abstract

In extracorporeal membrane oxygenation (ECMO), a life-support therapy used in patients with severe cardiac or respiratory failure, blood is circulated outside the body through an artificial circuit that includes a membrane oxygenator and a rotary blood pump (RBPs). While ECMO can be life-saving, it exposes blood to artificial surfaces and elevated mechanical stresses, particularly in RBPs, which operate under supraphysiological shear conditions and are a primary source of red blood cell (RBC) damage. A key complication is shear-induced hemolysis, the mechanical destruction of RBCs in high-shear blood flows. This damage releases hemoglobin and contributes to adverse events such as thrombosis, inflammation, and organ dysfunction. To better understand and reduce hemolysis, numerical models have become an essential tool in the design and evaluation of blood-contacting devices. However, current *in-silico* approaches face major limitations. They are computationally expensive, often lack validation against clinical data, and typically neglect biological and experimental variability. As a result, their ability to reliably predict hemolysis across realistic clinical conditions remains limited. The overarching aim of this dissertation is to develop a computationally efficient, uncertainty-aware *in-silico* framework that can predict hemolysis across the full clinical operating range of RBPs and explicitly account for key sources of uncertainty.

To make large-scale predictions feasible, a reduced-order model was developed based on non-intrusive polynomial chaos expansion. This approach enables rapid yet accurate simulation of flow conditions across the entire clinical operating space. Building on this, stress-based hemolysis predictions were validated to *in-vivo* data from 580 ECMO patients. While the results showed good agreement with cohort-level trends, they also revealed substantial inter- and intra-patient variability, pointing to the limits of purely mechanical models. To investigate the source of this variability, a controlled experimental study was conducted using a custom-built shearing device. Repeated hemolysis measurements under identical conditions showed unexpectedly high intra-donor variability, challenging assumptions underlying standard calibration practices and highlighting inherent uncertainty in experimental data. To address this, a Bayesian parameter estimation framework was introduced that incorporates this variability directly into model calibration. The resulting probabilistic model provides not only hemolysis predictions but also quantified uncertainty, increasing robustness, interpretability, and credibility in both clinical and regulatory contexts.

Together, these contributions help to establish a more comprehensive framework for hemolysis modeling by combining computational efficiency, clinical validation, and uncertainty quantification. Building on this foundation, future applications such as *in-silico* clinical trials or patient-specific modeling across diverse patient conditions may become more feasible.

Zusammenfassung

Bei der extrakorporalen Membranoxygenierung (ECMO), einer lebenserhaltenden Therapie für Patienten mit schwerem Herz- oder Lungenversagen, wird Blut außerhalb des Körpers durch einen künstlichen Kreislauf geleitet, der einen Membranoxygenator und eine rotierende Blutpumpe (RBP) umfasst. Obwohl ECMO Therapie lebensrettend sein kann, kommt das Blut mit künstlichen Oberflächen und hohen mechanischen Belastungen in Kontakt. Besonders in RBPs wirken extreme Scherkräfte, die eine Hauptursache für Schäden an roten Blutkörperchen (RBCs) darstellen. Eine zentrale Folge ist die scherinduzierte Hämolyse, also die mechanische Zerstörung von RBCs in Hochscherströmungen. Dabei wird Hämoglobin freigesetzt, was zu Komplikationen wie Thrombosen, Entzündungen und Organversagen führen kann. Um Hämolyse besser zu verstehen und zu reduzieren, kommen numerische Modelle zunehmend bei der Entwicklung blutführender Geräte zum Einsatz. Aktuelle *in-silico*-Ansätze stoßen jedoch an Grenzen: Sie sind rechenintensiv, oft nicht klinisch validiert und vernachlässigen biologische und experimentelle Schwankungen. Ihre Vorhersagekraft unter realistischen Bedingungen ist daher begrenzt. Ziel dieser Arbeit ist es, ein effizientes, unsicherheitsbewusstes *in-silico*-Modell zu entwickeln, das Hämolyse im gesamten klinischen Einsatzbereich von RBPs vorhersagen kann und zentrale Unsicherheiten berücksichtigt. Um großflächige Vorhersagen zu ermöglichen, wurde ein Surrogatmodell auf Basis nicht-intrusiver polynomieller Chaos-Expansion entwickelt. Dieser Ansatz erlaubt schnelle und präzise Simulationen im gesamten klinischen Betriebsbereich. Darauf aufbauend wurden scherspannungsbasierte Hämolysevorhersagen mit *in-vivo*-Daten von 580 ECMO-Patienten validiert. Die Ergebnisse stimmten gut mit Trends auf Kohortenebene überein, zeigten aber auch große inter- und intraindividuelle Unterschiede, was auf die Grenzen rein mechanischer Modelle hinweist. Um die Ursache zu untersuchen, wurde eine experimentelle Studie mit einem eigens entwickelten Scherversuchsstand durchgeführt. Wiederholte Messungen unter gleichen Bedingungen zeigten überraschend hohe Schwankungen innerhalb einzelner Spender, was gängige Kalibrierannahmen infrage stellt und auf eine grundlegende Unsicherheit in experimentellen Daten hinweist. Zur Berücksichtigung dieser Unsicherheit wurde ein Bayessches Parameterschätzverfahren eingeführt, das diese Schwankungen direkt in die Modellkalibrierung einbezieht. Das resultierende probabilistische Modell liefert Vorhersagen samt Unsicherheitsangabe und erhöht so Robustheit, Nachvollziehbarkeit und Glaubwürdigkeit in klinischem und regulatorischem Kontext. Diese Arbeit trägt durch die Kombination aus Recheneffizienz, klinischer Validierung und Unsicherheitsquantifizierung dazu bei, ein umfassenderes Konzept zur Hämolysemodellierung zu schaffen. Darauf aufbauend könnten künftige Anwendungen wie *in-silico*-Studien oder patientenspezifische Modellierungen realistischer werden.

Contents

1 Introduction

Shear-induced hemolysis refers to the mechanical destruction of red blood cells (RBCs) in blood flows, particularly relevant in blood-contacting medical devices such as rotary blood pumps (RBPs), where supraphysiological shear stresses prevail. The breakdown of RBCs releases hemoglobin into the plasma, leading to increased risks of thrombosis, oxidative damage, and inflammatory responses, ultimately resulting in severe complications and elevated mortality.

One clinical application of RBPs is in extracorporeal membrane oxygenation (ECMO) support, a life-saving therapy that provides essential respiratory and cardiac support for critically ill patients, often as a last resort to maintain heart and/or lung function. Despite its life-saving potential, the survival rate for adult patients undergoing veno-venous ECMO (VV ECMO) remains at approximately 60 % [1]. A major contributor to this limited survival rate is the occurrence of severe complications, such as bleeding and thrombosis, which affect nearly 40 % of patients [2] and are strongly associated with increased morbidity and mortality [3]. These adverse events have been closely linked to hemolysis [4], primarily caused by the mechanical stresses exerted by the RBPs integrated into ECMO circuits [5]. Minimizing such hemolysis through informed design choices is therefore a critical step toward improving the safety, effectiveness, and clinical outcomes of ECMO and other RBP-based therapies.

To address this challenge, *in-silico* modeling has become an invaluable tool for optimizing device design. By rapidly testing and refining configurations, these computational approaches reduce the need for extensive *in-vitro* or *in-vivo* experimentation. In this context, computational fluid dynamics (CFD) simulations are used to calculate blood flow inside devices, providing the foundation for various hemolysis modeling techniques. However, despite substantial advances in *in-silico* hemolysis modeling in recent years, no existing model has yet been able to accurately predict real-world hemolysis values.

One potential reason is that current numerical hemolysis models typically neglect biological variability inherent in blood, despite the fact that significant variability is observed at both microscopic scales (e.g., cellular properties) [6, 7] and macroscopic scales (e.g., experimental hemolysis measurements) [8]. Capturing these well known uncertainties is thus a key requirement for novel hemolysis models. Moreover, validation of hemolysis models has predominantly relied on *in-vitro* data, which can substantially diverge from clinical, *in-vivo* conditions, thereby limiting confidence in the predictive capacity of current models for device certification purposes. Additionally, RBPs are usually optimized for specific design operating points, yet clinical practice requires operation over wide ranges of flow rate (0–7 L/min) and pressure head conditions (0–400 mmHg) [9, 10]. While comprehensive evaluation across this entire operational range is essential, it remains computationally infeasible with current hemolysis

modeling approaches, as each operating point requires a separate, resource-intensive CFD simulation.
To overcome these shortcomings, the overarching aim of this dissertation is to develop a computationally efficient, uncertainty-aware *in-silico* hemolysis modeling approach. To achieve this, the following research questions will be addressed:

1. How can fluid variables across a wide range of RBP operating points be computed efficiently?

2. How accurately do *in-silico* hemolysis models capture clinical *in-vivo* hemolysis data?

3. How can uncertainties in experimental hemolysis data be quantified?

4. How can these experimental uncertainties be incorporated into uncertainty-aware *in-silico* hemolysis models?

Answering these questions will enable a more robust evaluation of current hemolysis models across diverse clinical operating conditions. It will also shed light on the extent of uncertainty in predictions, paving the way for uncertainty-aware numerical hemolysis models, which is of particular importance for the potential future use of numerical methods in the regulatory approval process of medical devices.
To address the outlined research questions and accomplish the overarching aim, this dissertation is structured into a series of interlinked chapters that build upon one another in a logical progression (*see* Figure 1.1).
Chapter 2 provides essential clinical and technical background, detailing the physiology of hemolysis, the application of ECMO therapy, and the fundamentals of numerical hemolysis modeling, laying the groundwork for understanding the clinical significance and technical challenges addressed throughout the dissertation.
Chapter 3 addresses the first major methodological challenge: the computational cost of evaluating RBP performance across their full clinical operating range. By introducing a reduced-order model (ROM) based on non-intrusive polynomial chaos expansion (NIPCE), this chapter lays the computational groundwork for all following numerical analyses. The ROM enables efficient and accurate flow predictions across a wide parameter space, thereby supporting both large-scale validation studies and uncertainty quantification.
Building directly on this, **Chapter 4** shifts the focus from computational efficiency to clinical relevance. It critically evaluates the predictive accuracy of stress-based hemolysis models, leveraging a large-scale dataset of *in-vivo* hemolysis measurements from 580 ECMO patients. This chapter demonstrates that, when paired with the ROM developed in Chapter 3, stress based numerical models can reflect real-world hemolysis trends, while also revealing considerable inter-patient variability that motivates a closer examination of experimental uncertainty.
Chapter 5 addresses this gap by systematically quantifying variability in hemolysis measurements using a custom-built, rheometer-based shearing device. Through extensive repeat experiments ($n = 100$), it uncovers substantial intra-donor variability that is often overlooked in conventional hemolysis studies, challenging assumptions about data reliability and revealing a key source of uncertainty in model calibration.
Chapter 6 builds on these insights by integrating experimental variability into a Bayesian

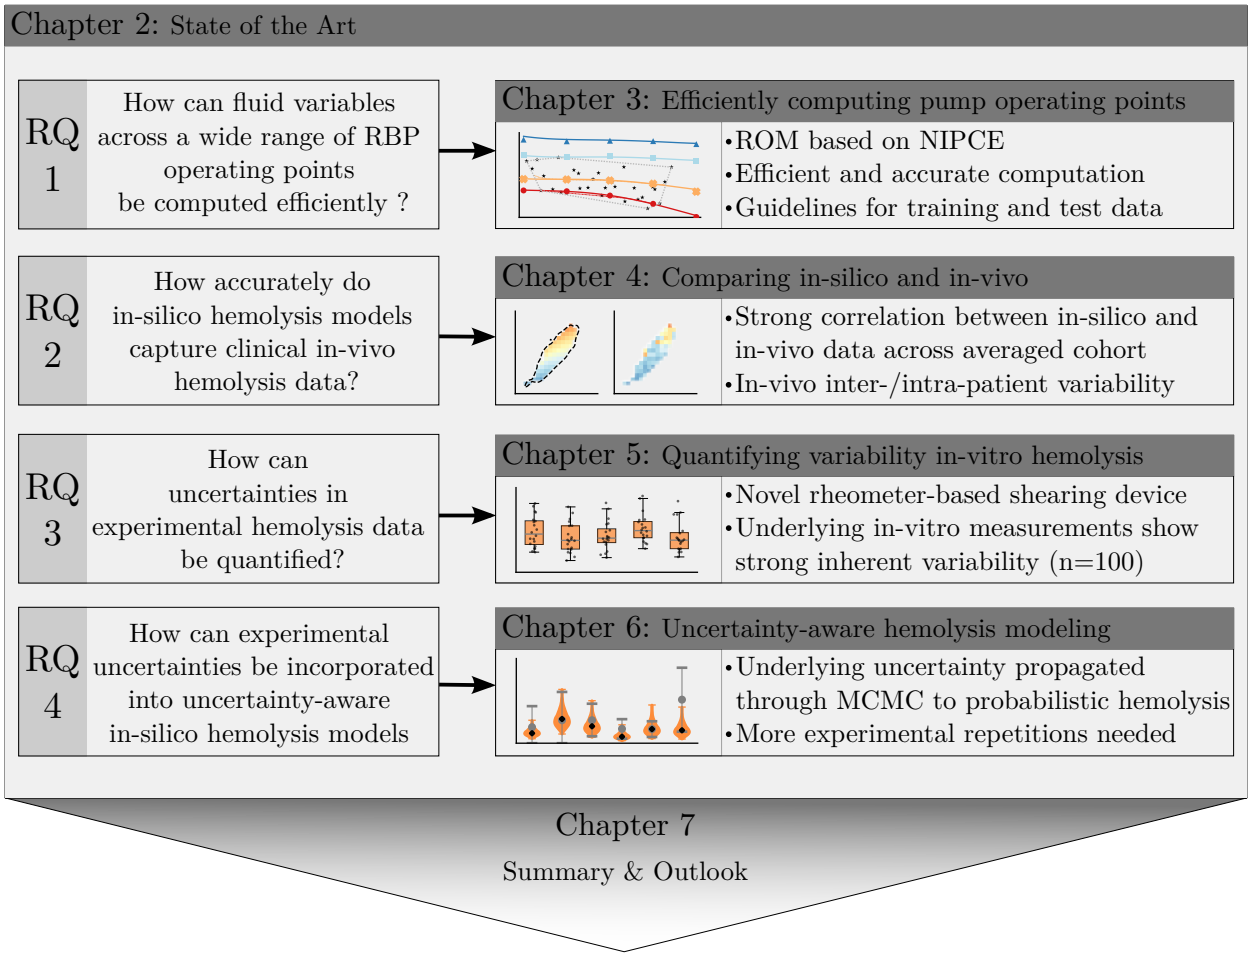

Figure 1.1: Graphical overview of the dissertation structure. Chapter 2 provides the clinical and technical background, framing the four central research questions (RQ1–RQ4) addressed in Chapters 3–6. Chapter 7 concludes with a summary and outlook.

inference framework, using Markov Chain Monte Carlo (MCMC) sampling to estimate full posterior distributions for model parameters. This transforms classical deterministic hemolysis models into uncertainty-aware tools, improving robustness, interpretability, and applicability in clinical and regulatory contexts.

Finally, **Chapter 7** summarizes the key findings of the dissertation, placing the developed methods and insights into a broader scientific and clinical context, and concludes by discussing implications and perspectives for future research directions.

2 State of the art

This chapter lays the foundational background required for the subsequent chapters by introducing key clinical (2.1), technical (2.2), and modeling (2.3) aspects. It covers the physiological basis of blood and hemolysis, the clinical use of ECMO therapy, the design principles of RBPs, and the fundamentals of numerical hemolysis modeling, establishing the clinical relevance and technical context for the challenges addressed in this dissertation.

2.1 Clinical background

Providing the clinical foundation for understanding hemolysis in the context of ECMO, this chapter begins with an overview of the most relevant properties of blood, focusing on physiological flow characteristics and factors contributing to blood damage (2.1.1). Building on this, the application of ECMO therapy is discussed in detail, including its components, monitoring strategies, and typical therapy stages (2.1.2).

2.1.1 Blood properties

Blood is a biological fluid with essential physiological roles, primarily delivering oxygen from the lungs to tissues throughout the body. Beyond oxygen transport, blood carries nutrients, removes metabolic wastes, facilitates immune responses, and acts as a medium for biochemical signaling [13–15]. To fulfill these diverse functions, blood consists of plasma, a protein-rich fluid primarily composed of water (89 %), and various cellular elements, including RBCs, leukocytes, and platelets [11].

Table 2.1: Average physiological composition of human blood, showing the relative proportions of plasma and cellular components, as well as the major constituents within each fraction [11, 12].

Components	**Percentage**	**Components**	**Percentage**
Plasma	55	Water	89
		Electrolytes	2
		Lipids	3
		Proteins	6
Cellular components	45	RBC	99
		Leukocytes	<1
		Platelets	<1

The approximate average volume composition of these major components is summarized in Table 2.1. However, this composition varies markedly between individuals due to factors such as sex, age, ethnicity, and environmental influences [12, 16], which in turn may influence susceptibility to hemolysis [17].

This composition of plasma and cellular components is rhythmically propelled by the heart through a hierarchical vascular network that spans several orders of magnitude in vessel diameter, ranging from large arteries measuring approximately 30 mm to capillaries less than 5 µm wide [18, 19]. Throughout this network, blood encounters a broad range of mechanical and rheological conditions. Shear stresses vary considerably, from around 1–3 Pa in large arteries such as the aorta and up to approximately 15 Pa in the microcirculation [20, 21]. Under

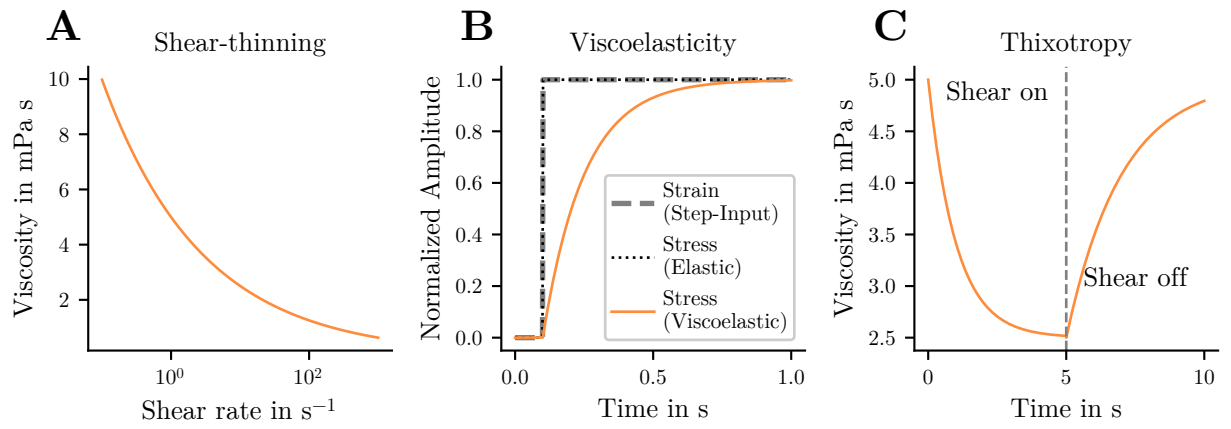

Figure 2.1: Key rheological behaviors of blood. Panel A illustrates shear-thinning behavior, where viscosity decreases with increasing shear rate due to the alignment and deformation of red blood cells. Panel B shows viscoelasticity, characterized by a time-dependent stress response following a step strain input, highlighting the difference between purely elastic and viscoelastic materials. Panel C depicts thixotropy, where viscosity progressively decreases under sustained shear stress and gradually recovers after the stress is removed.

these conditions, blood exhibits complex non-Newtonian behavior, including shear-thinning, viscoelasticity, and thixotropy [22].

Shear-thinning refers to the characteristic nonlinear, monotonically decreasing relationship of blood viscosity with increasing shear rate, as illustrated in Figure 2.1A. At high shear rates, typical of small vessels and capillaries, blood viscosity decreases substantially, facilitating efficient perfusion through microvessels. Conversely, at low shear rates common in large vessels, viscosity increases, which suppresses the transition to turbulence and prevents RBC sedimentation, thereby supporting stable and efficient circulation across the hierarchical vascular network. This behavior primarily results from the reversible aggregation and disaggregation of RBCs, coupled with their interactions with plasma proteins [23, 24]. At high shear rates, RBC aggregates break down, allowing cells to align and deform along the flow direction, reducing frictional forces. At low shear rates, increased RBC aggregation elevates viscosity and contributes to the overall shear-thinning behavior.

Viscoelasticity refers to the combined elastic and viscous behavior of blood, especially evident under pulsatile flow conditions typical of arterial circulation. Blood exhibits elastic recovery and delayed deformation, enabling it to buffer pressure fluctuations and smooth out

flow pulsations. This property contributes to energy-efficient blood transport and protects vascular structures from mechanical fatigue [25]. Figure 2.1B illustrates viscoelasticity by comparing an ideal elastic material and a viscoelastic material (blood) under a sudden step in strain (gray dashed line). The elastic material (dotted black line) responds instantaneously, matching the step input. In contrast, the viscoelastic material (solid orange line) exhibits a gradual, time-dependent increase in deformation, smoothing the step over a characteristic timescale of 0.1–0.5 seconds, which is typical for human blood [22]. This behavior arises primarily from the deformability of RBCs and their interactions with the surrounding plasma [26, 27]. The elastic response is largely driven by the highly deformable membrane of RBCs, while the viscous component is governed by plasma viscosity and cell-cell interactions. The balance between these behaviors influences wave propagation, impedance, and damping within the arterial system, particularly under pulsatile conditions.

Thixotropy refers to the time-dependent decrease in blood viscosity under sustained shear, even when the shear rate remains constant. This property allows blood to adjust its flow resistance in response to changing hemodynamic conditions, supporting flow stability during transitions such as exercise, rest, postural changes, or blood pressure fluctuations. Thixotropy is particularly important in the venous circulation and other low-shear environments, where RBC aggregation is more pronounced [28]. Figure 2.1C presents a simplified, isolated view of thixotropic behavior, illustrating the typical timescale of 1–5 seconds over which this process occurs in human blood [22]. This representation highlights the gradual breakdown and reformation of RBC aggregates under sustained or interrupted shear. Continuous shear leads to the gradual breakdown of RBC aggregates, decreasing viscosity over time. In the absence of shear, these aggregates gradually reform, restoring higher viscosity. It is important to note, however, that thixotropy and viscoelasticity represent distinct rheological mechanisms. Viscoelasticity arises from the rapid elastic and viscous response of RBCs and plasma, typically within 0.1–0.5 seconds, whereas thixotropy reflects slower changes related to the dynamic formation and breakup of RBC aggregates. Although both phenomena influence blood's time-dependent flow behavior, they act on separate timescales and are governed by different underlying processes.

Among the cellular components responsible for these properties, RBCs play the most dominant role. They are numerically and volumetrically predominant, typically constituting about 40–

Table 2.2: Geometric properties of human red blood cells (RBCs) with corresponding mean values and standard deviations (std) [29, 30].

Property of RBC	**Unit**	**Mean ± std**
Diameter	µm	7.65 ± 0.02
Minimum thickness	µm	1.44 ± 0.01
Maximum thickness	µm	2.84 ± 0.01
Surface area	μm^2	129.95 ± 0.40
Volume	μm^3	97.91 ± 0.41
Sphericity index	µm	0.792 ± 0.001
Membrane thickness	µm	4.2

45 % of blood volume, a proportion known as the hematocrit. Each milliliter of blood contains roughly 4–6 billion cells. With an average diameter of 7–8 µm and a thickness of approximately 2 µm, their characteristic biconcave disc shape increases the surface-area-to-volume ratio by approximately 44 % compared to a sphere of equal volume [31]. This unique geometry reflects the primary role of RBCs in oxygen transport, enabling efficient gas exchange in the lungs and peripheral tissues. Inside the cell, oxygen is carried by hemoglobin, a protein capable of reversibly binding oxygen molecules. To accommodate a high concentration of hemoglobin and reduce internal mechanical resistance, thereby enhancing cell deformability during microvascular circulation, mature RBCs lack a nucleus and other organelles. As a result of these adaptations, nearly all hemoglobin is contained within intact RBCs, where it is densely packed, leading to physiological whole blood concentrations of 12–13.5 g/dL. In contrast, the concentration of extracellular hemoglobin, known as plasma-free hemoglobin (pfHb), remains minimal under physiological conditions, typically around 0.02 g/dL [32]. An overview of typical RBC geometry parameters, including surface area, volume, and membrane thickness, is provided in Table 2.2. The exceptional mechanical felxibility of RBCs is enabled by a uniquely engineered membrane that enables them to endure the repeated mechanical stresses of circulation. This membrane is composed of a phospholipid bilayer, which provides basic

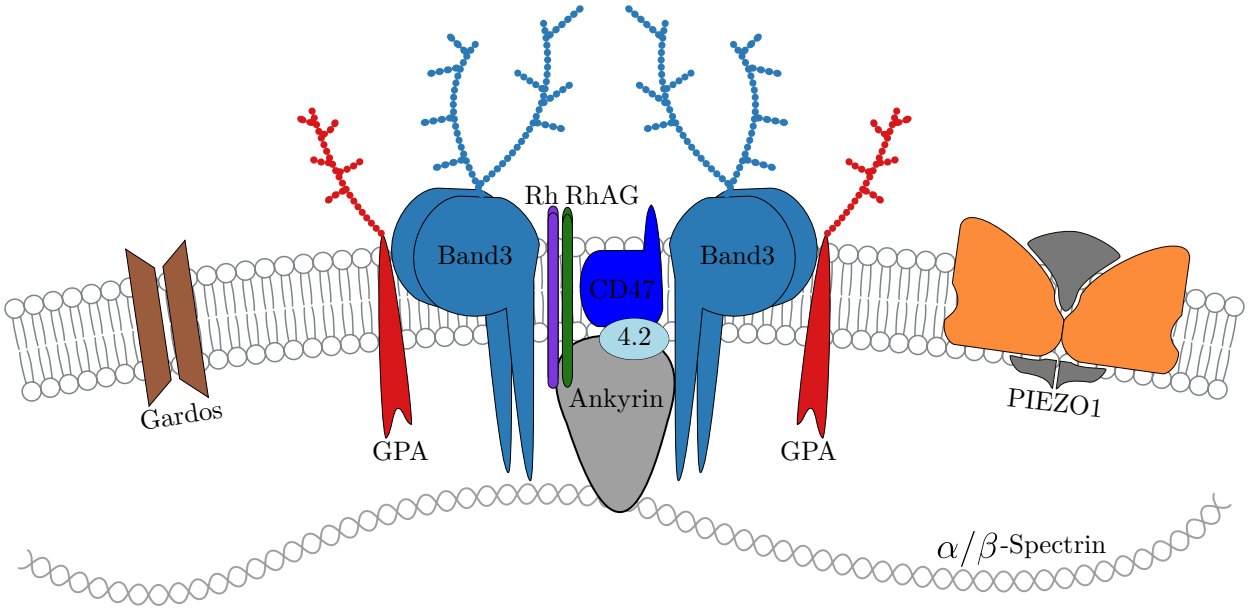

Figure 2.2: Schematic representation of selected key structural and functional components of the RBC membrane. The illustration highlights the lipid bilayer, a subset of transmembrane proteins, and the underlying cytoskeletal network of α- and β-spectrin. Shown membrane proteins include Band 3, glycophorin A (GPA), CD47 (integrin-associated protein), protein 4.2, and the Rh complex consisting of Rhesus polypeptide (Rh) and Rh-associated glycoprotein (RhAG), which are linked to the cytoskeleton via ankyrin and other adaptor proteins. Mechanosensitive and ion-regulatory channels such as PIEZO1 and the calcium-activated potassium (K^+) channel (Gardos channel) are depicted as important contributors to the RBC's mechanical responsiveness. Adapted from [33, 34], licensed under CC BY 4.0.

structural integrity and accommodates lateral motion of embedded proteins. Beneath the bilayer lies a highly organized cytoskeletal network, primarily made up of spectrin filaments arranged in a quasi-hexagonal lattice. This underlying scaffold is anchored to the bilayer

via linker proteins such as ankyrin and band 3, creating a dynamic composite that resists shear forces while allowing substantial deformation [35] (*see* Figure 2.2). This architecture allows RBCs to undergo reversible shape changes as they pass through narrow capillaries and splenic slits, often narrower than the cell itself, without compromising membrane integrity.
Beyond their mechanical design, RBCs exhibit a sophisticated level of physiological responsiveness. Despite lacking a nucleus and organelles, they are not simply inert carriers of hemoglobin but dynamic participants in blood flow regulation [13, 36]. Embedded within the membrane are various ion transporters and mechanosensitive ion channels that allow the cell to detect and respond to mechanical and osmotic changes (*see* Figure 2.2). Among these, the Piezo1 channel has emerged as a key mechanotransducer. Activated by membrane tension or shear stress, Piezo1 opens transiently to permit calcium influx, initiating a cascade of intracellular events. These include modulation of ion homeostasis, activation of downstream channels such as the Gardos channel (a Ca^{2+}-sensitive K^+ channel), resulting in changes in cell volume, and adjustments to membrane stiffness [37]. These responses are critical for maintaining RBC deformability, especially under conditions of mechanical load, and help prevent premature hemolysis in high-shear environments such as those encountered in narrow vessels or extracorporeal circuits.

Despite this remarkable deformability and capacity to adjust membrane stiffness, RBCs remain vulnerable to rupture when exposed to supraphysiological shear stress, whether in pathological vessel conditions or artificial circulation systems like RBPs [38]. Under such conditions, the mechanical integrity of the membrane fails, leading to hemolysis and the uncontrolled release of intracellular contents into the plasma. The resulting increase in pfHb has been associated with multiple adverse outcomes, including thrombosis, vascular dysfunction, and organ injury [39–41].
Importantly, hemoglobin is not the only bioactive molecule released during hemolysis. Another key molecule is adenosine triphosphate (ATP), which is normally stored within RBCs at high concentrations. In addition to its central role in cellular energy metabolism, ATP also functions as a potent extracellular signaling molecule. When released into the circulation, ATP interacts with receptors of endothelial and smooth muscle cells, modulating vascular tone [13]. This can alter blood pressure and local velocities, creating changes in flow conditions that may further elevate shear stress promoting additional RBC damage and sustained ATP release.
Beyond RBCs, supraphysiological shear stress can also damage other cellular and molecular components of blood. High shear rates are known to induce structural changes in high-molecular-weight multimers of von Willebrand factor (vWF), leading to the breakdown of vWF multimers and impairing their ability to mediate platelet adhesion resulting in acquired von Willebrand syndrome, which is a bleeding disorder frequently observed in patients with prolonged mechanical circulatory support such as in ventricular assist device therapy [42].
In addition to vWF, also platelets and leukocytes can be affected by unphysiological flow conditions. Platelet activation contributes to a prothrombotic state [43], while disturbed shear patterns in the vascular system can promote inappropriate leukocyte-endothelial interactions, potentially triggering inflammation and vascular injury [44].
The cumulative effect is a highly interconnected network of mechanical, biochemical, and immunological interactions, in which RBC destruction is not only a consequence of elevated

shear but also acts as a driver of additional vascular dysfunction, contributing to impaired vessel tone, inflammatory signaling, and hemostatic imbalance. In extreme cases of hemolysis, such as during ECMO support, this cycle can escalate into a widespread pathophysiological response resulting in hemolysis-induced thrombosis, bleeding, and organ injury [39–41]. The interplay between mechanical forces, cellular rupture, plasma protein degradation, and downstream signaling highlights the delicate balance RBCs and other blood components maintain within the circulatory system and how easily that balance can be disrupted under supraphysiological shear stress conditions.
Lastly, the complexity described at the cellular and molecular levels is further compounded by substantial biological variability in blood constitution [12, 16]. This variability may reflect differences in RBC aging, membrane characteristics, and ion transport, as well as broader factors such as genetic background, age, health status, and medication use. As a result, RBC mechanical properties like deformability and susceptibility to hemolysis can vary substantially both within and between individuals [17]. Such inter- and intra-donor variability is consistently observed in experimental settings and underscores the multifactorial nature of hemolytic responses.

While this overview does not capture the full breadth of blood's complexity, it highlights several of the most influential factors governing hemolysis. The RBC, and its potential for destruction, must be understood within the context of a highly dynamic and interconnected system. Mechanotransduction, membrane integrity, ion flux, protein degradation, vascular tone, and immunological signaling all interact across multiple orders of magnitude in both space and time. These processes range from nanometer-scale protein interactions and millisecond ion channel responses to systemic changes in vascular flow and thrombus formation occurring over hours or days. Taken together, they illustrate that RBC behavior and hemolysis are not isolated events but rather emergent outcomes of a finely balanced biological network that is highly sensitive to disruption under pathological mechanical stress.

- Blood shows complex flow behavior driven by RBCs, including shear-thinning, viscoelasticity, and thixotropy.
- RBCs are highly deformable but remain vulnerable to damage under supraphysiological shear stress.
- Hemolysis is a multifactorial process shaped by mechanical forces, cell properties, biochemical signaling, and biological variability.

2.1.2 Extracorporeal membrane oxygenation (ECMO)

ECMO is an advanced extracorporeal life-support technology utilized in critically ill patients whose heart and/or lungs are severely compromised, and conventional therapies have proven insufficient. ECMO therapy temporarily replaces or supports the native function of the heart and/or lungs, maintaining vital organ perfusion and oxygenation, thereby providing valuable time for recovery or further therapeutic interventions [45]. ECMO use has substantially increased over recent decades, with average annual case numbers reported in the registry of

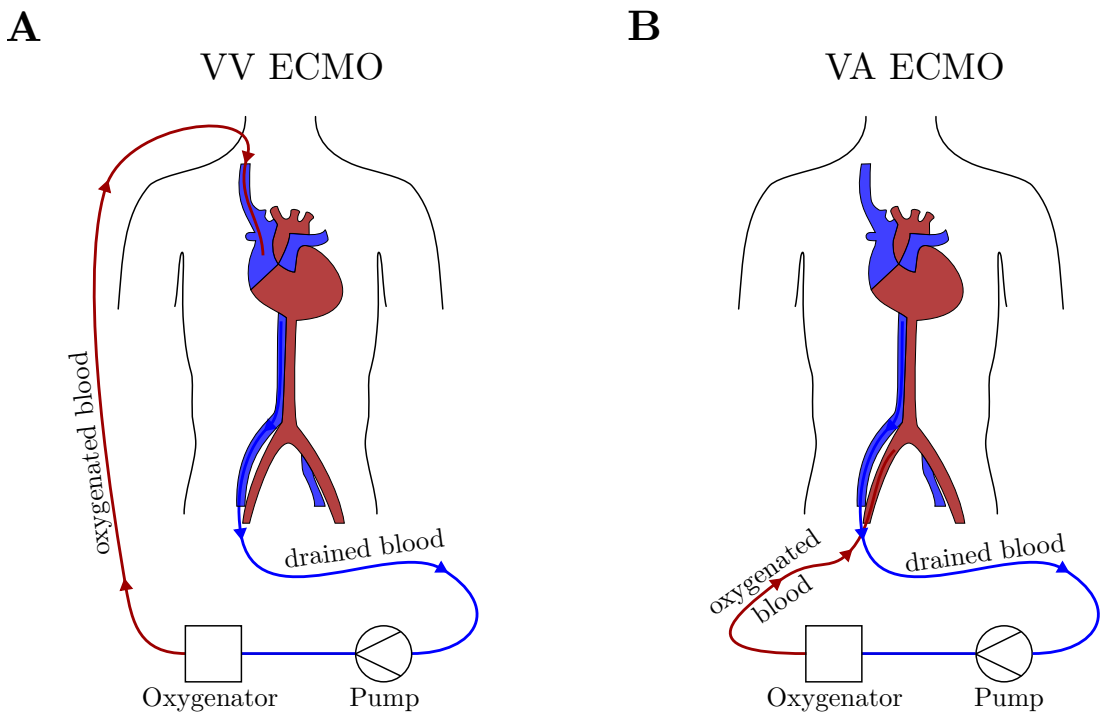

Figure 2.3: Schematic representation of veno-venous (VV) and veno-arterial (VA) ECMO configurations. Panel A shows VV ECMO, in which blood is typically drained from the femoral vein and returned via the internal jugular vein, providing respiratory support only. Panel B illustrates VA ECMO, where blood is drained from a vein and returned to an artery (e.g, from the femoral vein to the femoral artery) supporting both respiratory and circulatory function. The illustrations show common dual-cannula setups but cannulation strategies may vary based on clinical needs and patient anatomy. Adapted from [46], licensed under CC BY 4.0.

the Extracorporeal Life Support Organization (ELSO) reaching approximately 21,500 for adults and 5200 for pediatric and neonatal patients between 2020 and 2024 [1]. While ECMO in adult populations predominantly targets acute respiratory distress syndrome, cardiac arrest, or cardiogenic shock, pediatric indications often extend to conditions like congenital heart defects, neonatal respiratory distress, and acute respiratory failure due to infections or other congenital anomalies [47, 48]. Clinically, ECMO support is typically differentiated into two primary configurations based on the specific organs supported: veno-venous ECMO (VV ECMO) and veno-arterial ECMO (VA ECMO) (*see* Figure 2.3).
VV ECMO provides isolated respiratory support by removing deoxygenated blood from a large central vein, oxygenating it externally, and returning it to the venous circulation, thereby relying entirely on the patient's own cardiac function for systemic circulation.
Conversely, **VA ECMO** provides combined respiratory and circulatory support by withdrawing blood from the venous system, oxygenating it, and returning it directly to arterial circulation, thus partially or fully bypassing the heart.

A typical ECMO circuit comprises multiple components, including specialized cannulae inserted into major blood vessels, flexible tubing, a RBP that drives blood flow through

the extracorporeal circuit, a gas exchange membrane (oxygenator) including heat transfer and sweep gas mixer as well as multiple monitoring components like blood flow sensors, pressure sensors, bubble sensors and blood ports to sample blood probes for further analysis [49] (*see* Figure 2.4).

Cannulae form the critical interface between the patient's vascular system and the extracorporeal ECMO circuit. They are available in various lengths, diameters, and designs to suit individual patient anatomy and therapeutic requirements. Cannula selection depends on

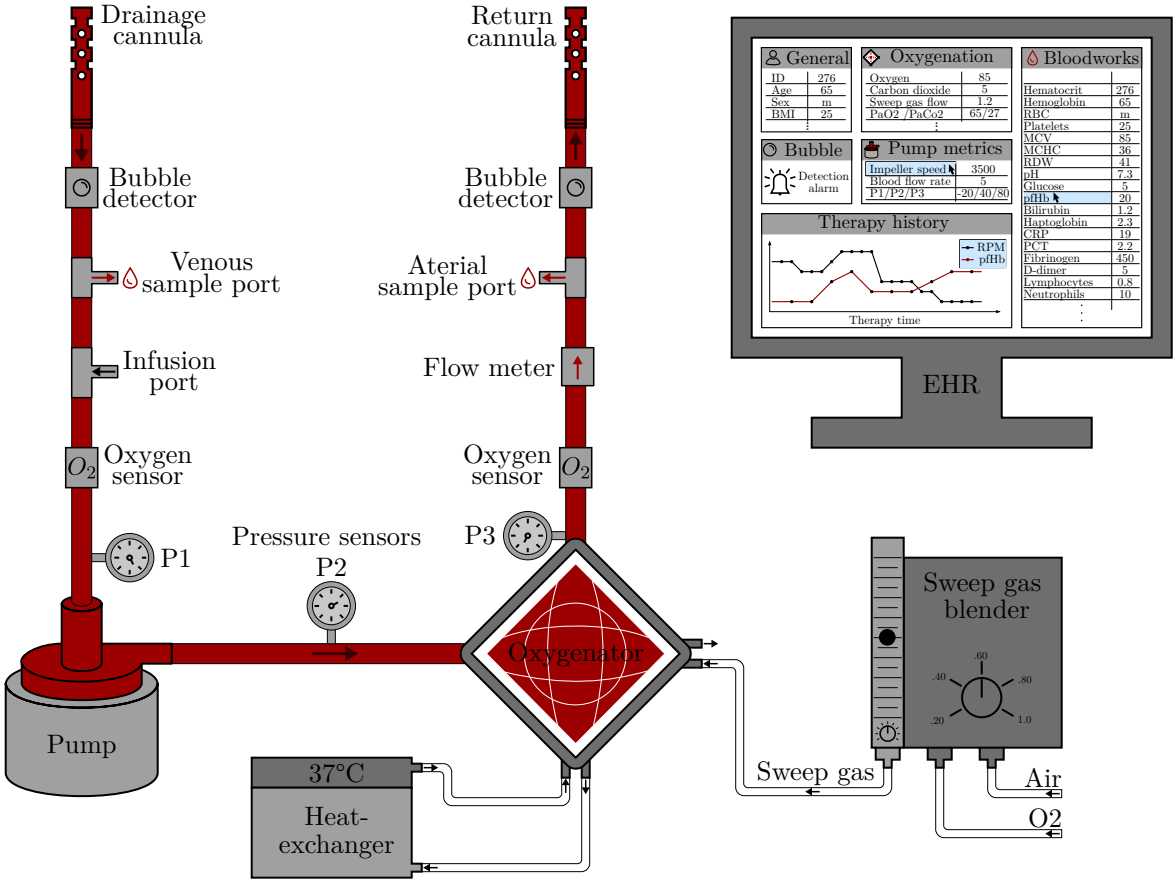

Figure 2.4: Schematic overview of an ECMO circuit with integrated monitoring and control components. The circuit includes sensors for pressure (P1–P3) and oxygenation, as well as ports for sampling (venous and arterial), infusion, and flow measurement. Bubble detectors enhance circuit safety. Sweep gas is delivered via a gas blender that regulates air and oxygen input to the oxygenator. Key parameters are continuously monitored and displayed on the electronic health record (EHR) interface.This schematic serves as an illustrative example; actual ECMO circuit configurations can vary depending on patient anatomy, clinical indication, equipment used, and institutional protocols.

several factors, including patient size, ECMO configuration, vessel diameter, and the required flow rate [50]. In VV ECMO, either a single dual-lumen cannula or two separate single-lumen cannulas may be used. In contrast, VA ECMO typically requires separate venous and arterial cannulas. Regardless of the configuration, correct positioning and secure fixation of all cannulas are critical to ensure safe and effective ECMO support. Malpositioning, particularly of the drainage cannula, can lead to critical suction events, where excessive negative pressure

causes partial vessel collapse and compromises circuit flow. Therefore, blood flow and pressure downstream of the drainage cannula are continuously monitored throughout ECMO support to detect and prevent such events.

The **RBP** forms the functional core of the extracorporeal circuit. Connected to the cannulae via flexible tubing, it drives blood flow by rotating an impeller, generating a pressure difference that enables circulation at the desired flow rate. Due to its rapidly rotating components, the pump is considered the primary source of blood damage within the ECMO circuit [5, 51]. During ECMO support, the pump's rotational speed, pressure build up, and flow rate are continuously monitored and adjusted to ensure adequate circulatory and/or respiratory support. A detailed description of RBPs and their operating principles is provided in Chapter 2.2.1.

After passing through the pump, blood enters the **oxygenator**, where gas exchange and temperature regulation take place. A controlled sweep gas mixture, typically composed of oxygen and either air or carbon dioxide, is used to regulate blood gas levels by adjusting the partial pressures of oxygen and carbon dioxide. Simultaneously, an integrated heat exchanger ensures that the circulating blood is maintained at desired physiological temperatures. To facilitate efficient gas exchange, oxygenators are constructed using dense arrays of hollow fiber membranes, typically with individual fiber diameters of approximately 300 µm [49]. These microporous polymer fibers provide a large surface area for diffusion, allowing oxygen to enter the blood while carbon dioxide is removed. The total gas exchange surface area varies by design and clinical use but typically ranges from 0.8–1.8 m^2 [52]. Although oxygenators are designed to be as biocompatible as possible, this large artificial surface may nonetheless activate inflammatory and coagulation pathways [53]. Combined with the relatively low flow velocities within the fiber bundle, the large surface area of the oxygenator makes it particularly susceptible to thrombus formation [54]. To detect this risk, the pressure drop across the oxygenator is continuously monitored, as a gradual increase during prolonged ECMO therapy can indicate thrombus formation, for which oxygenator replacement may be required.

In addition to the technical monitoring of pump speed, flow rate, and pressure differences across the pump, cannulae, and oxygenator, routine **blood sampling** plays a central role in assessing the patient's condition during ECMO support. These samples provide essential clinical information beyond what can be obtained from in-line sensor-based monitoring alone. In the context of hemolysis, pfHb serves as the primary biomarker. Given the high intracellular concentration of hemoglobin in RBCs and its negligible presence in plasma under normal conditions, even small increases in pfHb can serve as a highly sensitive and early indicator of RBC destruction. Clinical thresholds vary, but pfHb levels above 50 mg/dL are widely used to indicate hemolysis, while concentrations exceeding 100 mg/dL are generally regarded as severe and may necessitate urgent intervention [49]. To broaden the diagnostic perspective, secondary hemolysis markers such as lactate dehydrogenase (LDH) and haptoglobin are also routinely assessed. LDH is a general marker of cellular injury and may also increase in conditions unrelated to ECMO, such as liver dysfunction or systemic infection [55, 56]. Haptoglobin, which binds free hemoglobin in plasma, decreases in response to hemolysis, though its levels may already be depleted in critically ill patients. Because these secondary markers can be influenced by a range of factors, they must be interpreted in conjunction with pfHb and the overall clinical picture. In addition to hemolysis-related markers, other laboratory parameters such as those reflecting oxygenation status, inflammation, and coagulation are also monitored

to guide clinical decision-making and assess overall patient stability.
To integrate and visualize this complex dataset, most modern hospitals employ **electronic health record** (EHR) systems, which collect and store both laboratory values and technical monitoring data. These platforms help clinicians identify parameter trends over time, bringing together data from a wide variety of sensors and laboratory findings to form a unified picture of the patient's condition, enabling early detection of complications and supporting well-informed treatment decisions.

Following the description of individual ECMO components and monitoring systems, it is important to consider how these elements are applied throughout the course of a typical ECMO therapy. Clinical ECMO support can generally be divided into three overarching

Table 2.3: Stages of a typical ECMO treatment, outlining the key clinical steps during initiation, active support, and the weaning phase.

Initiation	**ECMO support**	**Weaning**
• Vascular access	• Monitoring	• Gradual support reduction
• Cannula placement	• Pain & Santicoagulation management	• Weaning trials
• Circuit connection	• Rehabilitation	• Decannulation
• Pump initiation		

phases: initiation, ECMO support, and weaning (*see* Table 2.3).
In the **initiation phase**, vascular access is first evaluated to determine the appropriate cannulation strategy. Based on patient anatomy and the planned ECMO configuration, suitable cannulae are selected [50]. Cannulation is then performed, and correct positioning is confirmed, typically using imaging guidance. Once vascular access is established, the extracorporeal circuit is connected. This circuit is usually pre-primed with an isotonic solution such as saline to ensure that it is air-free and ready for safe operation [57]. Depending on the priming volume relative to the patient's blood volume, a substantial hemodilution may occur. In cases where hemoglobin levels fall below critical thresholds, such as 7 g/dL [58], RBC transfusion may be required to maintain adequate oxygen delivery and physiological blood characteristics.
ECMO support begins once the circuit is running and the patient is physiologically stabilized. At this stage, the correct operating point must be established through continuous adjustment of pump speed, blood flow rate, and sweep gas settings to maintain adequate perfusion and gas exchange. While the primary objective of ECMO therapy is to stabilize physiological circulation and oxygenation, other clinical priorities must also be managed. These include pain and sedation management (balancing conscious and unconscious states) [59], anticoagulation and bleeding controls [60], and, if the patient is stable enough, the initiation of rehabilitation measures, which have been associated with shortened therapy time [61].
If cardiac or pulmonary recovery is observed, the therapy enters the **weaning phase**, during which ECMO support is gradually reduced. This is typically achieved by stepwise lowering of blood flow and oxygenator settings. Before final removal of support, weaning trials are conducted, in which ECMO assistance is temporarily minimized while closely monitoring the patient's ability to maintain stable gas exchange and hemodynamics [62, 63]. If the trials are

successful, the patient is decannulated, and the therapy concludes with close post-treatment monitoring and follow-up.

Overall, ECMO is a life-saving intervention for patients with severe cardiac and/or respiratory failure. However, it imposes substantial stress on blood due to exposure to supraphysiological shear forces and extensive artificial surfaces. These factors contribute to complex physiological responses, including hemolysis, inflammation, and impaired coagulation, which must be carefully managed in the context of an already critically ill patient. Effective ECMO support therefore requires continuous monitoring and an multifaceted approach that balances multiple therapeutic goals.

- ECMO provides temporary respiratory and/or circulatory support in critically ill patients when conventional therapies fail.
- The circuit exposes blood to artificial surfaces and supraphysiological shear stress, contributing to hemolysis, inflammation, and coagulation complications.
- Effective ECMO therapy requires careful component selection and continuous monitoring across initiation, support, and weaning phases.

2.2 Technical background

This section introduces the technical foundations relevant to hemolysis in ECMO systems. It covers the evolution and working principles of RBPs in ECMO application, with a focus on centrifugal pump design and related shear-induced challenges (2.2.1). It also outlines the regulatory landscape for these devices, including approval pathways and the increasing role of *in-silico* methods in device evaluation (2.2.2).

2.2.1 Rotary blood pumps (RBPs)

This dissertation focuses on the numerical modeling of hemolysis, using ECMO as a representative clinical application. As blood pump design and flow conditions directly influence hemolysis, the following section provides an overview of RPB technologies relevant to ECMO. This includes a general introduction to RBP types, their historical development in ECMO circuits, and the clinical rationale behind the current standard of centrifugal blood pumps before explaining their working principles and design challenges related to hemocompatibility.
RBPs are key components in various forms of mechanical circulatory support, including ventricular assist devices, total artificial hearts, and ECMO. These devices are employed to provide short- or long-term circulatory support in patients with acute or chronic cardiac and/or respiratory failure. RBPs function by imparting energy to blood using a rotating impeller, and are typically categorized into three types based on their flow path geometry: axial, centrifugal, and diagonal [64].
In **axial pumps**, the blood moves parallel to the impeller axis, driven forward by a spinning helical impeller. These pumps are compact and capable of generating high flow rates, but often at the cost of higher shear stress, which can increase hemolysis [64]. **Centrifugal pumps**, by contrast, use a rotating impeller to generate flow perpendicular to the rotor axis through radial acceleration and are generally associated with better hemocompatibility [65]. **Diagonal pumps** combine elements of both axial and centrifugal designs, offering a balanced configuration in terms of flow dynamics and hemocompatibility.
Although RBPs are now firmly established in modern ECMO systems, this has not always been the case. In the early decades of ECMO development, RBPs were not yet used. Instead, roller pumps, which were originally developed for cardiopulmonary bypass, served as the primary pumping technology in ECMO circuits [45]. These positive-displacement pumps generate flow by compressing flexible tubing in a peristaltic motion. While mechanically simple and effective, roller pumps are associated with several drawbacks, including high rates of hemolysis, risk of tubing rupture, and limited responsiveness to changes in venous drainage [45]. Because the flow rate is mechanically imposed by the rollers, a reduction in venous blood availability does not result in a corresponding reduction in pump output. This can lead to excessive negative pressure at the drainage cannula, increasing the risk of suction events and hemolysis [45]. Despite these limitations, roller pumps remain used in pediatric ECMO due to their ability to maintain accurate and stable flow at very low volumes, which centrifugal pumps may struggle with due to flow sensitivity and low-inertia behavior at reduced rotational speeds.
The transition from roller to centrifugal pumps in ECMO began in the early 2000s. By 2015,

over 95 % of all reported ECMO cases in the ELSO registry were utilizing centrifugal blood pumps [45]. This widespread clinical adoption is attributable to several factors: reduced mechanical trauma to blood, lower inlet suction pressures and compatibility with modern magnetically driven systems that eliminate the need for dynamic shaft seals. Therefore, centrifugal pumps are now considered the standard of care in most adult ECMO systems, offering improved safety, reliability, and long-term performance [66, 67].

Fundamentally, a centrifugal pump is designed to raise the energy level of a fluid at a given flow rate, most commonly measured in terms of hydraulic head [68]. The head is defined as the energy per unit weight of the fluid and is typically expressed in meters. It represents the height to which a pump can raise a fluid and is directly related to the pressure generated. In this sense, the fundamental task of the centrifugal pump is to convert mechanical energy from a rotating shaft into hydraulic energy in the form of increased pressure head. In medical pump applications, however, pressure is more commonly expressed in millimeters of mercury (mmHg), consistent with clinical practice. This can be converted to meters of fluid column using the hydrostatic relation, $h = \frac{p}{\rho g}$, where ρ is the fluid density and g is the gravitational acceleration. For blood with an approximate density of $\rho = 1060\,\mathrm{kg/m^3}$, 1 mmHg corresponds to approximately 0.0128 m of hydraulic head.
To achieve the increase in hydraulic head, fluid enters the centrifugal pump axially through the eye of the impeller, where the flow is guided into the rotating impeller blades (*see* Figure 2.5A). Due to the no-slip condition at the blade surfaces, the fluid adheres to the impeller motion, gaining angular momentum as it is driven outward in a radial direction. This process increases both the tangential and radial components of the fluid velocity, thereby raising its kinetic energy. Upon exiting the impeller, the high-velocity fluid enters a volute or diffuser, where the progressively expanding cross-sectional area decelerates the flow and converts kinetic energy into static pressure. The energy transfer within the impeller can be idealized using Euler's turbomachinery equation (equation (2.1)) [68]:

$$\Delta h = \frac{1}{g}\left(u_2 V_{u2} - u_1 V_{u1}\right) \tag{2.1}$$

where Δh is the theoretical head gain, u is the blade tip speed (with $u = R \cdot \omega$), and V_u is the tangential component of the absolute velocity of the fluid. The subscripts 1 and 2 refer to the inlet and outlet of the impeller, respectively. This equation describes the ideal energy transfer based solely on the change in angular momentum and assumes inviscid, incompressible flow without losses. In real systems, the actual head generated is lower due to hydraulic losses such as wall friction, flow separation, and turbulence. Moreover, Euler's equation does not account for secondary flows, slip at the impeller outlet, or mechanical losses in bearings and seals [68].
The kinematic interaction between the impeller and fluid is best illustrated using velocity triangles at both the inlet and outlet of the impeller (*see* Figure 2.5B). Each triangle represents the vectorial relationship between three velocities: the *absolute velocity* of the fluid ($\vec{V}$), the *relative velocity* of the fluid with respect to the rotating blades ($\vec{W}$), and the *blade speed* (u). The shape of the velocity triangle at the inlet is influenced by the incoming flow direction and impeller blade angle, while the triangle at the outlet is governed by the impeller geometry and flow deflection. The tangential components of the absolute velocity, V_{u1} and V_{u2}, determine

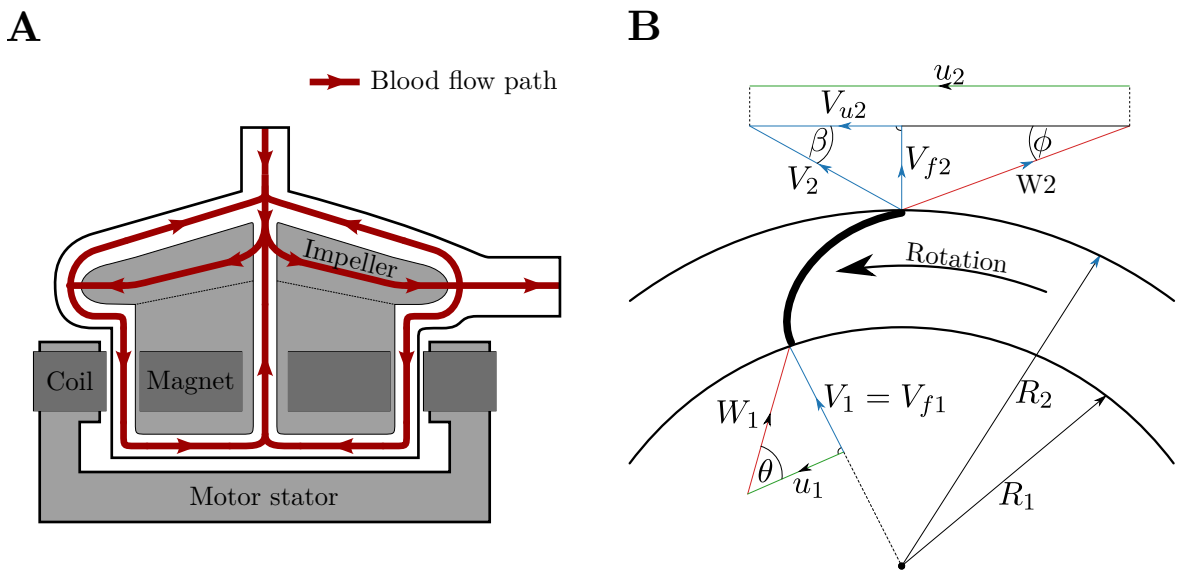

Figure 2.5: Panel A shows a cross-sectional view of a centrifugal blood pump, highlighting the blood flow path through the housing and around the rotating impeller, which is magnetically driven. Panel B illustrates the velocity triangles at the impeller inlet and outlet, representing the vectorial relationship between the absolute velocity of the fluid ($\vec{V}$, blue), the relative velocity with respect to the rotating blades ($\vec{W}$, red), and the blade speed (u, green).

the angular momentum at the inlet and outlet, respectively. In Figure 2.5B, the inlet velocity triangle is simplified by assuming $V_{u1} = 0$, meaning the incoming flow is aligned with the impeller blade angle and carries no tangential velocity component. This simplification aids illustration but must be considered when interpreting the theoretical energy transfer. According to Euler's equation, it is the difference $u_2V_{u2} - u_1V_{u1}$ that sets the theoretical head gain. Depending on the specific blade design (e.g., radial, backward-, or forward-curved) and operating conditions, both V_{u1} and V_{u2} may contribute significantly to the energy transfer. Therefore, the velocity triangles provide a visual and analytical tool to assess how effectively the impeller transfers energy to the fluid.

Once the fluid enters the volute, the increase in cross-sectional area causes the flow velocity to drop. According to Bernoulli's principle, this reduction in velocity results in an increase in static pressure, finalizing the conversion from kinetic energy to pressure.

The overall performance of the pump is characterized by the relationship between pressure head and volumetric flow rate, typically depicted in the pump characteristic curve (*see* Figure 2.6A). This curve generally shows a decreasing trend in pressure head with increasing flow rate, due to rising hydraulic losses and diminished energy transfer per unit mass. The hydraulic efficiency curve (*see* Figure 2.6B) typically exhibits a single peak at the best efficiency point, indicating the operating condition where energy conversion losses are minimized. This characteristic reflects the fact that centrifugal pumps are generally optimized for a single specific design point. When operating away from this optimum, whether at higher or lower flow rates, the pump is subjected to increased recirculation, flow separation, and incidence losses, all of which lead to a marked decline in efficiency.

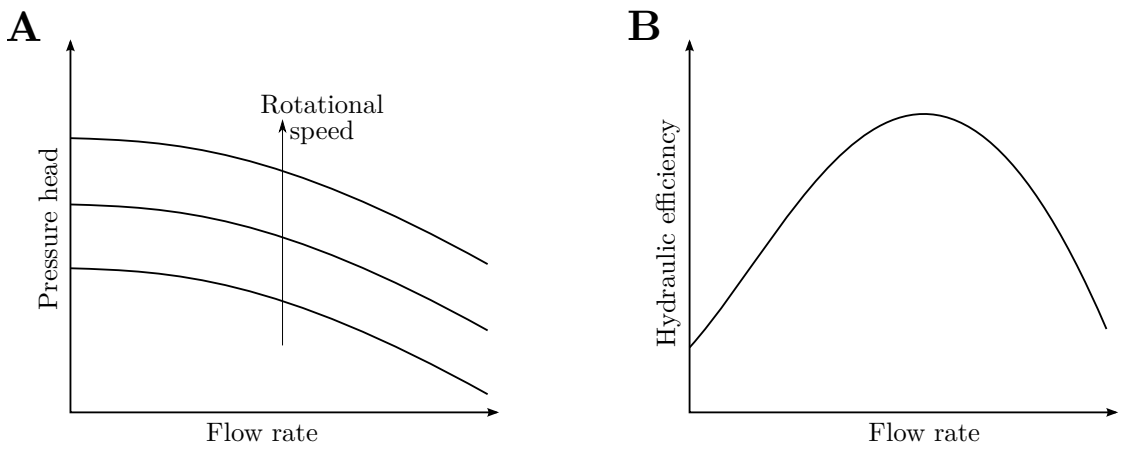

Figure 2.6: Schematic representation of typical pump performance characteristics. Panel A shows a set of head-flowrate (HQ) curves, which describe the relationship between pressure head and flow rate at different constant rotational speed lines. Panel B displays a schematic hydraulic efficiency curve as a function of flow rate with maximum efficiency achieved at an optimal flow rate.

The presence of a distinct best efficiency point is typical of conventional turbomachinery design, such as in water supply pumps or industrial cooling systems, where the pump is expected to operate steadily at or near a fixed design flow rate and pressure head. In those settings, operating near the best efficiency point ensures minimal energy loss and therefore lower operational costs or improved overall system efficiency. However, in ECMO therapy, this becomes a practical challenge, since the RBP is operated across a wide range of flow rates and pressure heads rather than at single operating point as clinical operating point requirements during ECMO treatment often change substantially over time, with flow rates ranging from 0–7 L/min [9, 10]. An example previously discussed in Chapter 2.1.2 is the weaning phase, during which the flow rate is progressively reduced, forcing the pump to operate under low-flow, off-design conditions. Such operation places the pump far from its best efficiency point, which can significantly increase blood damage due to local variations in shear stress and turbulence [69].

The need to operate across a wide range of flow rates is not the only factor that sets RBPs apart from traditional turbomachinery. While their fundamental working principle is rooted in classical turbomachinery theory, the design of blood pumps diverges substantially due to the unique physiological requirements of blood-handling systems. In conventional industrial pumps, minimizing hydraulic losses is typically achieved by employing narrow tolerances and small gap sizes, which reduce internal leakage and enhance efficiency [68]. However, applying these strategies to blood pumps would lead to high local shear stress peaks, which are known to cause hemolysis [70]. To mitigate this risk, blood pumps are deliberately designed with larger gap regions that help to reduce shear stress gradients and promote smoother, less traumatic flow conditions (*see* Figure 2.5A).

Another key distinction lies in the design of bearing systems. Traditional turbomachinery commonly uses mechanical contact bearings [68], which are efficient and robust for handling industrial fluids but unsuitable for blood due to frictional heating and surface wear. While several early RBP designs incorporated miniaturized ball bearings, which were favored for

their simplicity and low manufacturing cost, direct mechanical contact with blood led to increased risks of hemolysis and thrombus formation, limiting their suitability for long-term clinical use [64]. Despite these drawbacks, ball bearings remain a practical option in short-term or cost-effective disposable devices, where design simplicity and low production costs are prioritized. Over time, bearing technologies in blood pumps evolved toward either hydraulic (fluid) bearings or fully magnetically levitated rotors. Hydraulic bearings provide a balance between mechanical stability and low shear exposure, while magnetic bearings enable completely contact-free operation, eliminating friction and surface wear altogether.

A closely related design difference is the approach to shaft sealing. Conventional turbomachines often rely on mechanically sealed shafts to transmit torque from the motor to the impeller [68], which introduces potential leakage paths and contamination risks, which are both unacceptable in blood-contacting devices [64]. To address this, RBPs frequently integrate magnets directly into the impeller and achieve rotation through an external magnetic drive (*see* Figure 2.5A). This configuration allows for a completely sealed flow path without the need for dynamic shaft seals, thereby improving hemocompatibility and reducing the risk of contamination.

However, despite their advantages, centrifugal blood pumps are not without limitations. One major concern is the risk of hemolysis, which increases in regions of high shear stress, particularly around narrow impeller clearances or during off-design operation, such as at low flow rates during weaning phases. In contrast, regions of low shear within the pump housing or at circuit junctions can promote platelet activation and thrombus formation. These opposing risks, hemolysis under high shear and thrombosis under low shear, highlight a fundamental design challenge and underscore the need to balance hydraulic performance with hemocompatibility.

While continuous research into impeller geometry, surface optimization, and flow dynamics seeks to further improve the safety and performance of RBPs [70, 71], it is important to recognize the complex nature of blood as a working fluid in addressing persistent challenges such as hemolysis, thrombosis, and bleeding. An increased focus on blood-informed design, which integrates biological responses such as shear-induced damage, protein deposition, and thrombus formation into engineering decisions, offers a promising path to mitigate these risks [72].

- Centrifugal pumps have become the standard in ECMO due to superior hemocompatibility, lower shear stress, and contact free, magnetically driven designs.
- Unlike industrial pumps, blood pumps require larger gaps and specialized bearing and shaft solutions to minimize shear-induced damage and thrombosis risk.
- Hemolysis and thrombosis remain key challenges, especially under off-design conditions, emphasizing the need for blood-informed pump design.

2.2.2 Device regulation and computational modeling

RBPs operate in a highly regulated environment. As life-sustaining Class III medical devices, they are subject to rigorous oversight to ensure safety, clinical efficacy, and long-term performance. Regulatory authorities define the clinical and technical evidence required for

market approval and ongoing monitoring, thereby exerting strong influence over the design, validation, and clinical application of these technologies. During ECMO support, RBPs replace essential physiological functions, and any failure may lead to immediate, life-threatening consequences. As such, regulation emphasizes verifying safety and reliability under clinical conditions, with particular focus on risks like thromboembolism, hemolysis, infection, and mechanical failure.

At the international level, the International Medical Device Regulators Forum plays a central role in aligning regulatory approaches across countries by promoting harmonized terminology, risk classification systems, and safety and performance principles [73]. Standards developed by the International Organization for Standardization (ISO), such as ISO 13485 for quality management systems [74] and ISO 14971 for risk management [75], are widely adopted or referenced by national regulators. Despite ongoing harmonization efforts, regulation of medical devices remains primarily governed at the national or regional level [76]. Among the various frameworks, those of the United States and the European Union hold particular prominence. Owing to their market size and regulatory maturity, these systems exert substantial global influence and are frequently used as reference models by other jurisdictions especially regarding device classification, nomenclature, approval processes, and post-market obligations [77].

In the United States, the U.S. Food and Drug Administration (FDA) evaluates RBPs through the Premarket Approval process. In the EU, Regulation (EU) 2017/745 (MDR) governs oversight, with Notified Bodies conducting conformity assessments leading to CE marking. Both systems require comprehensive technical documentation, risk management, and clinical evidence, with specific requirements based on device classification [78].

While regulatory texts do not provide a detailed, step-by-step procedure for developing a device, they do define the types of evidence and performance criteria that must be addressed throughout the development process. Manufacturers are guided by binding regulations, such as 21 CFR 814 for FDA submissions or Annexes II and XIV of the EU MDR.

To meet these regulatory expectations, regulatory bodies require that all relevant risks to be systematically identified, mitigated, and validated using appropriate methods. As illustrated in Table 2.4, the typical development process begins with device design and risk analysis within a certified quality management system. This is followed by *in-vitro* testing to evaluate key performance parameters, including hydraulic characteristics, structural durability, and hemocompatibility, among other critical metrics. Hemolysis is usually assessed according to ASTM F1841 [79]. Preclinical *in-vivo* studies in animal models help evaluate physiological compatibility, thrombogenicity, and device integration. Depending on the device's novelty and risk profile, clinical investigations may be required, ranging from early feasibility studies to confirmatory pivotal trials. Following market approval, manufacturers must maintain robust post-market surveillance. This includes adverse event reporting and, under the MDR, implementation of post-market clinical follow-up plans and periodic safety update reports.

The overall development and approval process is lengthy and resource-intensive. While estimates vary, the time from concept to market approval for Class III devices is often cited as three to seven years [80]. Recognizing the high costs, long timelines, and risks of late-stage failure, both the FDA and the European regulatory system have increasingly supported the use of computational modeling and simulation in regulatory submissions [81, 82]. These *in-silico* approaches offer the potential to accelerate device evaluation, reduce reliance on

Table 2.4: Typical evidence across regulatory phases for rotary blood pumps (RBPs). The asterisk (*) indicates that all entries in this column are examples only; actual submissions may include additional device-specific evidence.

Typical phase	Purpose	Illustrative evidence / tests*
Design & risk management	Concept sound, risks identified	Risk analysis (ISO 14971), traceability of design inputs and outputs
Bench (in-vitro)	Verify specs in lab	Functional / hydraulic / durability and hemocompatibility tests
Animal (pre-clinical)	Show safety & performance in-vivo	Representative animal model, blood-compatibility observations
Clinical	Confirm safety & effectiveness in humans	Feasibility and pivotal trials, satisfying clinical endpoint criteria
Post-market	Monitor real-world performance	Post-market monitoring & complaints, real-world evidence & registries

animal and human testing, decrease risk of late failure, and improve predictive insight into device safety and performance [82, 83].

To support appropriate use of computational models, the FDA endorses a credibility assessment framework based on the ASME V&V 40 standard [84]. This standard outlines a risk-informed approach in which the context of use (COU) defines the level of effort needed for verification, validation, and uncertainty quantification (VVUQ). The greater the regulatory impact and the potential consequences of incorrect model predictions, the more rigorous the VVUQ process must be to demonstrate that the model is sufficiently reliable for its intended purpose.

Verification ensures that the model's numerical implementation is correct. This includes software quality assurance, numerical code verification against analytical solutions, and estimation of discretization and solver error. In the context of RBPs, this might involve verifying that shear stress fields and scalar transport equations are solved consistently and are not substantially affected by mesh resolution or solver settings.

Validation determines whether the model accurately represents the physical system under the conditions specified in the COU. This requires justification of the model form (e.g., shear stress-based hemolysis models), selection of input parameters, and boundary condition relevance. Experimental comparisons should account for test equivalence and parameter consistency. Common validation approaches include comparing CFD predictions to particle image velocimetry measurements for flow fields or to *in-vitro* hemolysis tests for accuracy of hemolysis predictions. The strength of evidence should scale with model risk, from qualitative trend alignment at low risk to quantitative agreement at high risk.

Uncertainty Quantification evaluates how model input variability and knowledge gaps affect model outputs. In hemolysis models, this includes uncertainties in flow conditions, blood properties, and model coefficients. Sensitivity analysis and uncertainty propagation

help assess confidence in predicted hemolysis levels. The required effort depends on model risk, ranging from basic estimates to detailed uncertainty bounds at high risk.
This scaling of VVUQ effort is formalized through the concept of **model risk**, which provides a structured way to assess how much credibility evidence is needed for a given application. Currently, computational models such as hemolysis models are primarily used during the design phase of RBPs to support iterative optimization. While they play a valuable role in development, regulatory approval still relies predominantly on *in-vitro* and *in-vivo* evidence. To expand the role of modeling in regulatory submissions, the V&V 40 framework defines credibility requirements proportionate to the model's risk. As illustrated in Figure 2.7, model

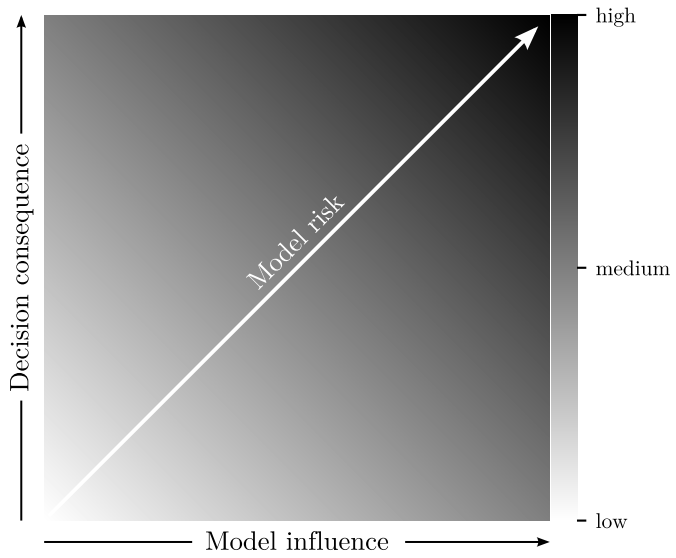

Figure 2.7: Conceptual illustration of model risk based on the ASME V&V 40 standard [84]. Model risk increases with both Model influence, which refers to the extent to which a model informs a decision, and Decision consequence, which describes the potential impact if the decision is incorrect. The diagonal gradient visualizes increasing model risk and helps determine the required rigor of verification and validation activities.

risk is defined by two independent factors: model influence and decision consequence.
Model influence reflects the extent to which a regulatory decision depends on model output, while decision consequence refers to the clinical impact of an incorrect decision. Higher values of either increase the overall model risk. For instance, if a hemolysis model is used to replace physical testing in a Class III device submission, model risk is considered high. Conversely, if the model is used only to inform test conditions, the risk is substantially lower [84].
A practical example of applying this framework to hemolysis modeling in RBPs has been published by Morrison *et al.* [85], illustrating how model credibility depends on the specific COU. Two scenarios were considered: one with low model influence, where CFD was used to select worst-case conditions for subsequent physical testing; and one with high model influence, where hemolysis predictions informed safety assessments across the full operating range with limited experimental data. The study demonstrated how higher model risk required more rigorous verification and validation, including mesh convergence, experimental uncertainty quantification, and multi-point validation of hemolysis predictions. This case illustrates

how model credibility is not static but shifts with the model's role in the decision-making process. It also shows that, even for high-risk Class III devices, computational modeling can contribute meaningfully to regulatory submissions when credibility is thoroughly assessed and documented.

To fully apply this risk-based framework to hemolysis modeling, all relevant sources of uncertainty must be addressed in a systematic and transparent manner. While the importance of uncertainty quantification (UQ) is well recognized in computational science and starting to be applied in biomedical engineering across disciplines such as cardiovascular mechanics [86], imaging [87], and pharmacology [88], its implementation in numerical hemolysis modeling remains largely absent. Existing models typically operate deterministically, providing single-point estimates without accounting for physiological variability or uncertainty in model parameters. As a result, predictions are often reported without confidence intervals or formal error bounds, limiting their interpretability in risk-sensitive applications such as regulatory decision-making.

Uncertainty in computational modeling is typically classified into two main types: aleatoric and epistemic [89].

Aleatoric uncertainty arises from intrinsic variability in the physical system and cannot be reduced through further model refinement. In hemolysis modeling, this includes physiological variability in blood properties such as RBC deformability and susceptibility to mechanical stress, which can vary both within and between individuals [17].

Epistemic uncertainty, in contrast, stems from incomplete knowledge or limitations in the model itself, including uncertain parameter values, assumptions in boundary conditions, and limited validation data. Unlike aleatoric uncertainty, it can be reduced through improved data, more representative experiments, or refinements in model formulation.

Despite the relevance of both types, current hemolysis models used in the context of RBPs generally neglect uncertainty quantification. This represents a major barrier to the credibility and regulatory acceptance of simulation-based predictions. Addressing uncertainty systematically would not only enhance transparency but also align numerical hemolysis modeling with the rigor expected for robust, risk-informed decision-making in regulatory settings.

To summarize, hemolysis modeling supports the development of RBPs but currently plays a limited role in regulatory decision-making. Its acceptance depends on the ability to demonstrate credibility through thorough verification and validation, scaled to the model's impact on regulatory outcomes. A key barrier remains the inadequate treatment of uncertainty, including both physiological variability and model-related knowledge gaps. Addressing these systematically is essential to improve the reliability of predictions and their value in high-risk regulatory contexts.

- RBPs require strict, risk-based evaluation with evidence from all development stages.
- Computational models are increasingly considered when supported by verification, validation, and uncertainty analysis.
- Hemolysis models aid design but must better reflect physiological variability and model uncertainty to inform regulatory decisions.

2.3 Modeling background

To address the clinical challenges posed by hemolysis in blood-contacting medical devices, reliable and efficient numerical modeling techniques are essential. This section introduces key methodologies used throughout this dissertation, beginning with a review of state-of-the-art approaches to numerical hemolysis modeling, including their capabilities and limitations. Subsequently, a ROM approach, designed specifically to overcome the computational challenges inherent in evaluating wide range of operating conditions of RBPs, will be introduced. Finally, Bayesian parameter estimation techniques are outlined, focusing on their ability to capture experimental uncertainties in parameter distributions, ultimately leading to more robust and clinically meaningful modeling outcomes.

2.3.1 Hemolysis modeling

As highlighted previously in Chapter 2.1.1, blood is a complex biological fluid influenced by numerous interrelated factors across a wide range of spatial and temporal scales. Capturing the full extent of this complexity within a single model is not feasible. Even when narrowing the focus exclusively to shear-induced hemolysis, the number of potential influencing factors and their respective interaction scales remains extensive.
To address this, the field of numerical shear-induced hemolysis modeling is generally structured across multiple scales, each offering distinct insights (*see* Figure 2.8). At the microscopic

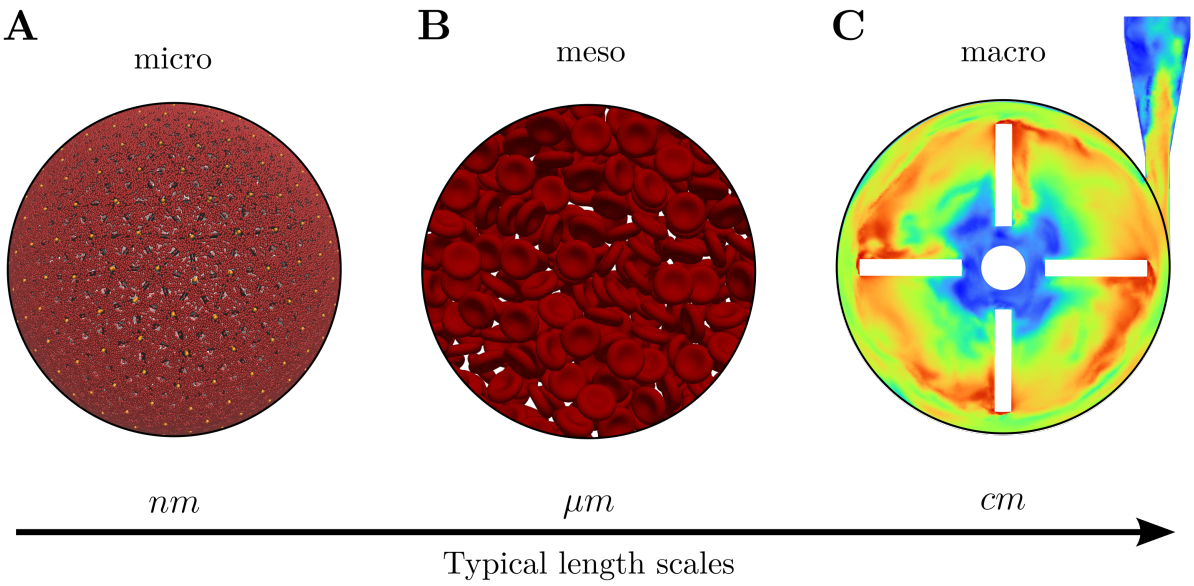

Figure 2.8: Conceptual representation of hemolysis modeling approaches across different length scales. Panel A (micro-scale, nanometer range) shows a coarse-grained model of a RBC membrane, representing individual membrane molecules and the underlying spectrin network (image generated using OpenRBC [90]). Panel B (meso-scale, micrometer range) illustrates an ensemble of deformable RBCs interacting within a flow field, modeled using the HemoCell framework [91]. Panel C (macro-scale, centimeter range) presents a cross-sectional view of blood velocity distribution in the FDA benchmark RBP [92], obtained from a continuum-based CFD simulation, serving as a basis for macro-scale hemolysis models.

scale (nanometers to micrometers), detailed coarse-grained models simulate the response of individual RBC membranes to mechanical stresses by resolving membrane structure, deformation mechanics, and the organization and interactions of major membrane proteins [93–96]. At the mesoscale (micrometers), models typically simulate ensembles of up to thousands of interacting, deformable RBCs, enabling the investigation of collective cell dynamics, cell-cell collisions, aggregation, and flow-induced deformation [97, 98]. At the macroscale (millimeters to centimeters), models describe bulk blood behavior relevant to clinical settings and large medical devices such as RBPs [99–101]. Additionally, multiscale models [102, 103] attempt to bridge two or more of these scales, though their practical accuracy and predictive reliability remain limited due to high computational costs and inherent modeling simplifications.
For the prediction of hemolysis in blood-contacting medical devices, such as RBPs, which typically span from centimeters to decimeters, microscopic and mesoscale simulations are computationally infeasible. Consequently, current numerical approaches for device-induced hemolysis predominantly rely on macroscale formulations.

Although such models do not explicitly resolve detailed cell mechanics, such as membrane deformation, rupture, or internal stresses, they indirectly account for their cumulative effects by fitting model parameters to experimental outcomes that reflect these underlying biological processes [99]. As a result, macroscale models can approximate the net impact of complex multiscale damage mechanisms without explicitly simulating them, enabling practical and computationally efficient predictions for large-scale device simulations.
In such experiments that inform macroscopic hemolysis models, blood samples are exposed to well-defined shear stress levels over specific exposure durations, typically using Couette geometries or capillary flow devices [104, 105]. A schematic of a Couette geometry is shown in Figure 2.9A, where the inner and/or outer cylinder is rotated to generate a well defined shear stress in the blood-filled gap (highlighted in red). The degree of hemolysis is quantified using the Hemolysis Index (HI), which measures the concentration of plasma-free hemoglobin released due to RBC lysis, normalized by hematocrit and total hemoglobin concentration:

$$HI = \frac{\Delta\text{pfHb}(100 - \text{HCT})}{Hb} \tag{2.2}$$

Here, ΔpfHb denotes the increase in plasma-free hemoglobin from before to after shear exposure, Hb is the total hemoglobin concentration, and HCT is the hematocrit. These experiments are performed at various operating points, typically covering shear stress levels of several hundred pascals and exposure durations in the range of a few seconds. A schematic example of the resulting experimental data is shown in Figure 2.9B. The schematic illustrates considerable variability across the tested operating points. Such variation is characteristic of biological systems and could stem from the inherent heterogeneity of blood samples as well as the sensitivity of the measurement techniques. To increase confidence in the observed trends and reduce the influence of outliers, measurements are typically performed in triplicate [104, 106, 107]. While this does not eliminate the underlying scatter, it enables a more reliable estimation of the mean response, which is particularly relevant when fitting empirical models. To mathematically capture the average empirical relationship between shear stress, exposure time, and hemolysis, a power-law equation is commonly employed:

$$HI = C\,\tau^{\beta}\,t^{\alpha} \tag{2.3}$$

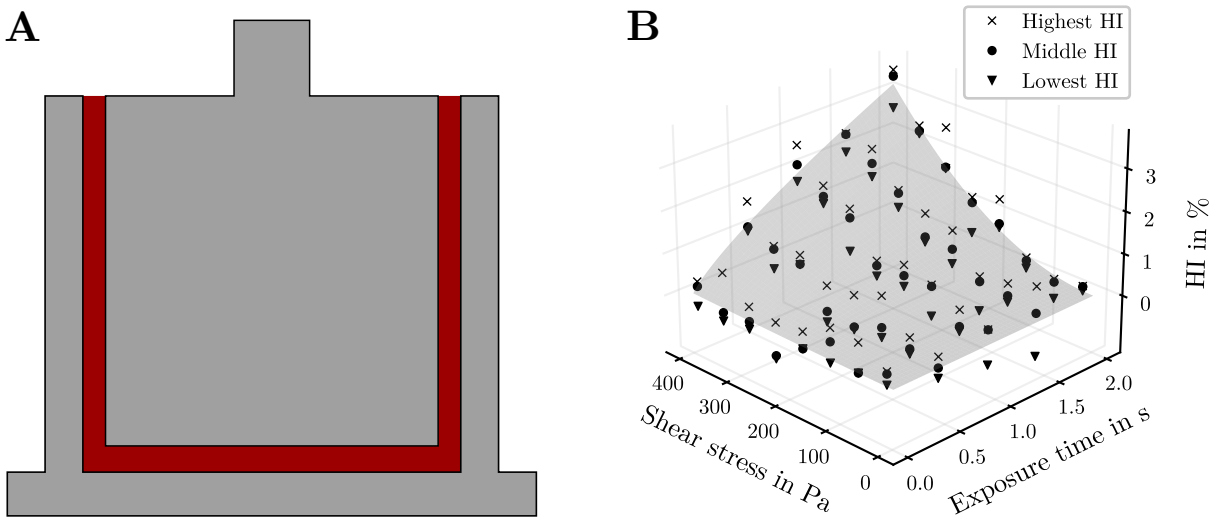

Figure 2.9: Panel A shows a schematic of a Couette geometry commonly used in hemolysis testing, where blood (depicted in red) is sheared between concentric rotating cylinders. Panel B presents representative experimental results, illustrating hemolysis index (HI) as a function of shear stress and exposure time. The spread between data groups (highest, middle, lowest HI) highlights the biological variability in hemolysis susceptibility across different blood samples and the grey surface represents averaged fitting of equation (2.3).

In this equation, τ denotes the applied shear stress, t is the exposure time, and C, α, and β are empirically determined fitting constants. The fitting process results in a single best-fit surface within the parameter space, which is represented in grey in Figure 2.9B. While this approach enables a compact and interpretable representation of the hemolysis response, it also introduces a major limitation: the fitted surface averages out the biological variability and measurement uncertainty, thereby neglecting the inherent stochastic nature of hemolysis. Despite this drawback, power-law-based hemolysis models remain widely used, particularly for comparative purposes where relative differences between conditions or devices are of interest. Numerous research groups have attempted to identify optimal fitting parameters through various experimental setups, applying different fluid shear stress scenarios to blood from different species [104, 106–108]. However, no universal set of parameters has yet been identified that accurately predicts absolute hemolysis values in real-world RBP applications. A comparison of published parameter sets reveals substantial discrepancies, underscoring the difficulty of generalizing hemolysis predictions across systems [109].

To overcome this issue, recent studies have adopted advanced optimization techniques to refine model parameters for specific devices. Craven *et al.* [110] developed a method to calibrate the empirical hemolysis model coefficients (C, α, β) for a specific device and blood species. They ran multiple CFD simulations of the device, each with different coefficient values, and built a surrogate model using Kriging to interpolate hemolysis outcomes. By comparing these predictions to experimental measurements, they identified the coefficient sets that best match the observed data. Demonstrated on laminar flow in a capillary tube, their method revealed that traditional coefficients from uniform shear experiments can result in absolute errors exceeding 2000 %, highlighting the benefits for device-specific calibration.

Similarly, Torner *et al.* [111] employed a multi-objective particle swarm optimization algorithm

to calibrate power-law coefficients across three benchmark test cases (capillary tube, FDA nozzle, FDA pump [92, 112, 113]) and twelve turbulent flow conditions. Using both relative error and correlation as fitness criteria, the algorithm identified Pareto-optimal parameter sets that improved consistency between numerical and experimental hemolysis values. However, while the optimized models showed better alignment than traditional coefficient sets, significant discrepancies in other geometries remained.
Despite these advancements, optimization-based approaches remain inherently limited in their generalizability and cannot yet guarantee quantitatively accurate predictions across diverse flow regimes. Nevertheless, they represent meaningful progress toward bridging the gap between *in-vitro* calibration and *in-silico* hemolysis modeling in complex medical devices.

Beyond fine-tuning the power-law parameters to improve predictive accuracy, other studies have focused on refining the implementation of the functional relationship between shear stress, exposure time, and hemolysis (*see* equation (2.3)) within computational models.
A key step in this direction is recognizing the frame of reference implicit in the original experimental setup, which underpins how damage is accumulated over time. Although the power-law model originates from experiments conducted on stationary blood volumes, such as in Couette viscometers, its formulation is inherently Lagrangian in nature [114]. This is because it describes the cumulative damage experienced by a fluid element as it moves through a shear field over time. Hemolysis depends not only on the local shear stress magnitude but also on the exposure duration along a fluid particle's trajectory. The original Couette experimental setup implicitly assumes that all fluid particles are subjected to the same stress history, effectively making the formulation equivalent to tracking a material volume, which is the defining characteristic of a Lagrangian description. This distinction is illustrated in

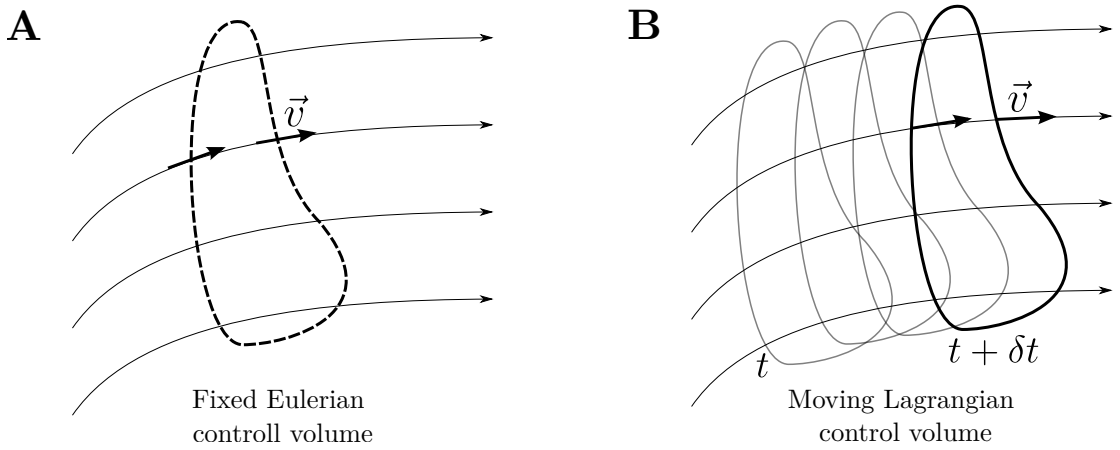

Figure 2.10: Comparison of Eulerian and Lagrangian reference frames. Panel A illustrates the Eulerian perspective, where fluid motion is observed through a fixed spatial control volume. Panel B shows the Lagrangian perspective, where individual fluid elements are followed as they move through space and time.

Figure 2.10. Panel A shows the Eulerian perspective, where a control volume is fixed in space and the motion of fluid through it is observed. In contrast, Panel B depicts the Lagrangian viewpoint, which follows a fluid element as it moves through the domain. Mathematically,

these two perspectives are related by:

$$\frac{D}{Dt} = \frac{\partial}{\partial t} + \vec{v} \cdot \nabla \tag{2.4}$$

where $\frac{D}{Dt}$ is the material (Lagrangian) derivative, $\frac{\partial}{\partial t}$ is the local (Eulerian) time derivative, and $\vec{v} \cdot \nabla$ accounts for the convective transport of a quantity due to fluid motion.
RBP simulations typically rely on CFD to resolve the flow field, including velocity, pressure, and shear stress, within an Eulerian framework. In this approach, these flow quantities are computed at fixed points in space on a mesh. This allows for efficient numerical solving and is particularly well suited for the complex geometries encountered in RBP designs. The fixed spatial reference of the Eulerian approach ensures full domain coverage and aligns with mature CFD solvers that include turbulence models, wall functions, and robust discretization schemes. Lagrangian CFD solvers, which track individual fluid particles throughout the domain, do exist in the form of mesh-free particle methods such as smoothed particle hydrodynamics [115]. However, these methods are not commonly used for device-scale simulations due to several limitations, including high computational cost, difficulties in resolving boundary layers, and challenges in maintaining uniform domain coverage, especially in complex or turbulent flows. As a result, the Eulerian formulation continues to be the standard in both commercial and research CFD tools.
Since most CFD solvers are based on an Eulerian framework, implementing hemolysis models within the same formulation is consistent and straightforward. For this the classical power-law formulation (*see* equation (2.3)), which describes blood damage as a function of shear stress and exposure time, can be reformulated as a scalar transport equation and solved alongside the momentum and continuity equations of classical CFD. For this the power law gets linearized over time:

$$\frac{\partial \widehat{HI}}{\partial t} + \vec{u} \cdot \nabla \widehat{HI} = C^{\frac{1}{\alpha}} \cdot \tau^{\frac{\beta}{\alpha}}, \quad \text{with} \quad \widehat{HI} = HI^{\frac{1}{\alpha}} = C^{\frac{1}{\alpha}} \cdot \tau^{\frac{\beta}{\alpha}} \cdot t \tag{2.5}$$

This makes Eulerian hemolysis models easy to integrate and computationally efficient, which explains their widespread use. However, transforming the original Lagrangian power-law formulation into an Eulerian transport equation involves two important approximations. First, exposure time is treated as a local variable tied to fixed spatial points, which is a reasonable simplification only in idealized, uniaxial flows with constant velocity along streamlines. In more complex or accelerating flows, this can lead to noticeable discrepancies. Faghih *et al.* [116] demonstrated that such deviations may reach up to 257 % in different geometries when compared to Lagrangian estimates. Second, the reformulation typically applies the exposure time exponent outside an integral, an operation that assumes either constant stress or unity exponents. As shown by Faghih and Sharp [117], this can affect accuracy in flow regions where stress and exposure time vary substantially across space. Despite these drawbacks, Eulerian hemolysis models remain popular due to their ease of implementation, low computational cost, and ability to provide domain-wide predictions. These qualities make them attractive for design studies that require a spatial overview of hemolysis risk in the entire device.
In contrast, Lagrangian hemolysis models follow fluid particles along their trajectories and

accumulate damage based on the full history of shear stress experienced over time. This path-based perspective is often seen as more closely aligned with the underlying physical process of hemolysis, which depends on the magnitude, direction and duration of mechanical stress. In this sense, Lagrangian frameworks offer a more direct representation of the cumulative nature of blood damage. Over time, a variety of extensions to Lagrangian hemolysis models have been proposed, many of which incorporate biological insights derived from *in-vitro* studies. Among these are approaches that account for stress history effects [118], viscoelastic properties of RBC membranes [101], and the potential influence of prior sublethal damage [119]. These examples reflect broader efforts to develop more detailed and physiologically motivated representations of hemolysis beyond the classical formulations. However, Lagrangian models also face practical challenges. They are typically more computationally demanding than Eulerian formulations, and their ability to capture hemolysis-relevant regions depends on the density and distribution of tracked particle paths. In complex flow fields, this may lead to incomplete spatial coverage, particularly in regions where flow features are highly variable or transient. A comparative study by Yu *et al.* [120] evaluated both Eulerian and Lagrangian hemolysis models and found that neither was able to predict absolute hemolysis values with high accuracy. The predicted results from various models spanned several orders of magnitude when compared to experimental data, underscoring the fact that, despite advances in modeling strategies, a quantitatively reliable prediction of hemolysis remains a challenging goal.

However, beyond parameter calibration and the choice of reference frame, the predictive performance of a macroscopic hemolysis model also depends on how forces acting on RBCs are modeled. To address this, a range of modeling strategies have been developed, each offering a different interpretation of how mechanical forces contribute to flow-induced damage. These approaches are typically categorized into stress-based, strain-based, and dissipation-based formulations:

Stress-based hemolysis models are the most widely used because they are conceptually simple and easy to integrate in commercial CFD solvers. They assume that hemolysis correlates with a single scalar shear stress (SSS), derived from the viscous stress tensor, and that damage monotonically accumulates over time according to the power-law formulation in equation (2.3). A central numerical challenge is how to define SSS in three-dimensional, time-varying flows. The empirical coefficients in the power-law model were calibrated under steady, simple-shear Couette conditions, so any generalized scalar should reduce to the applied shear stress in that limit. Faghih and Sharp [121] showed that several commonly used definitions fail this consistency check and can yield different predictions in complex geometries. A common approach to address this challenge is to compute SSS from the norm of the strain-rate tensor T_{ij}, leading to an equivalent stress definition:

$$SSS = \mu \sqrt{2\, T_{ij} T_{ij}} \tag{2.6}$$

where μ is the dynamic viscosity. Because stress-based formulations are computationally efficient and historically well documented, they remain the predominant option for device-level studies, despite limited sensitivity to stress mode (shear versus extension) and possible load–unload effects, in which cyclic or intermittent stress exposures may result in cumulative damage not well predicted by models based solely on instantaneous shear stress values. These

limitations motivate strain-based models, which aim to link hemolysis more directly to RBC deformation mechanics.

Strain-based hemolysis models offer a complementary alternative to stress-based formulations by focusing on the deformation history of RBCs, rather than their instantaneous exposure to mechanical stress. These models are motivated by the viscoelastic nature of the RBC membrane, which responds not only to the magnitude of applied forces but also to their duration and directionality. Rather than relying on single scalar stress values, strain-based approaches attempt to quantify the evolving shape distortion of RBCs and relate this directly to hemolysis. Early implementations typically followed a Lagrangian framework, where the deformation of individual RBCs was computed along their flow trajectories [122, 123]. While physically detailed, this method is computationally intensive and difficult to scale to simulations of entire devices involving millions of particles. To address these limitations, Dirkes *et al.* [124] recently proposed a fully Eulerian strain-based model, known as the Tank-Treading Morphology model. Their approach builds on the ellipsoidal deformation framework introduced by Arora *et al.* [122], in which RBCs are modeled as neutrally buoyant ellipsoids described by a morphology tensor S that evolves according to local fluid strain and vorticity. By transforming the model into an Eulerian formulation, Dirkes *et al.* [124] enabled domain-wide computation without the need for particle tracking. The model solves for the eigenvalues λ_i of S, which correspond to the squared semi-axis lengths of the ellipsoid, and assumes a tank-treading regime in which the cell's orientation reaches equilibrium through a balance of strain and rotational effects. This reduces computational cost while maintaining a strong link to underlying cell mechanics.

$$G_{\text{eff}} = \frac{2Df_1}{(1-D^2)f_2} \tag{2.7}$$

The output of the model is an effective shear rate, G_{eff}, which is calculated from the instantaneous cell distortion D, which in turn is derived from the eigenvalues λ_i of the morphology tensor. The constants f_1 and f_2 are empirically determined and calibrate the relationship between mechanical distortion and effective shear. The resulting G_{eff} is then used as input to the conventional empirical power-law model for hemolysis prediction.

Within the broader context of hemolysis modeling, strain-based approaches aim to complement existing formulations by incorporating additional mechanical detail at the cellular level. By explicitly capturing cell deformation, they naturally account for different stress modes, such as shear and extension, as well as time-dependent effects associated with the viscoelastic behavior of RBCs. Notably, recent experimental findings by Lommel *et al.* [125] indicate that extremely short exposure times do not produce measurable hemolysis, highlighting the importance of incorporating time-dependent damage accumulation into such models. The Eulerian formulation introduced by Dirkes *et al.* [124] represents a substantial step forward, making strain-based modeling computationally feasible for complex three-dimensional flow fields, including those in RBPs. This opens up the possibility of incorporating more detailed cell mechanics into device-scale simulations. At present, the method relies on solving coupled transport equations for the morphology tensor. Since this capability is not yet standard in commercial CFD solvers, it may pose a barrier to broader adoption. Furthermore, comprehensive experimental validation in full-scale devices remains an important next step. Although not employed in this dissertation, strain-based models hold considerable promise for improving

the biophysical accuracy of hemolysis predictions and may play an important role in future generations of blood damage modeling frameworks.

Dissipation-based hemolysis models constitute a third class of mechanical force formulations, linking RBC damage to the local energy-dissipation rate ε rather than to scalar stress or strain. The underlying hypothesis is that hemolysis correlates directly with the mechanical energy imparted to RBCs, which is ultimately dissipated through viscous friction. This approach is particularly appealing for turbulent and transitional flows, where traditional stress-based models may overestimate damage due to interpreting Reynolds stresses as direct mechanical loads on cells. Wu *et al.* [126] introduced an energy-dissipation-based power law,

$$\bar{h} = C \cdot \varepsilon^{\alpha} \tag{2.8}$$

where $\bar{h}$ is the hemolsyis rate, ε is the local energy dissipation rate per unit volume, and C, α are empirically fitted constants. The hemolysis index is then computed by multiplying $\bar{h}$ by the exposure time. Evaluated across a capillary tube [112], FDA nozzle [113], and FDA centrifugal pump [92], the model demonstrated reduced prediction error and improved alignment with experimental data, particularly in regions of high Reynolds-number where viscous dissipation is relatively low despite high velocity fluctuations. Faghih and Sharp [127] further examined the theoretical basis for dissipation-based modeling by analyzing membrane tension in flows with matched ε. Their results showed that dissipation-based modeling improved consistency between laminar and turbulent shear flows, although extensional flows continued to generate higher tension levels. This suggests that while energy dissipation captures important features of RBC loading, it may not fully represent the anisotropic and nonlinear deformation responses in all flow types.

Overall, dissipation-based models provide a physically motivated and computationally accessible alternative to stress-based formulations, particularly in turbulent flow regimes. Their integration into eulerian transport frameworks makes them compatible with standard CFD solvers. However, further studies in flows dominated by extension or strong anisotropy are necessary to evaluate their broader applicability and predictive reliability.

Taken together, stress-based, strain-based, and dissipation-based hemolysis models represent three distinct strategies for describing the mechanical forces that contribute to RBC damage. Each approach emphasizes different physical aspects of the flow: stress-based models offer simplicity and efficiency, strain-based models aim to incorporate cell-level deformation mechanics, and dissipation-based models seek to capture the mechanical energy lost due to viscous and turbulent deformation processes, which is postulated to be a key driver of hemolysis in turbulent flow environments. While each formulation has demonstrated strengths in specific applications, none fully captures the complexity of hemolysis across all flow conditions. As a result, the choice of model should be guided by the flow regime, the desired level of physical fidelity, and the available computational resources. Continued efforts to integrate experimental findings, biophysical insights, and numerical strategies remain essential for improving the predictive capabilities of hemolysis models in complex medical devices.

Future directions will likely continue along two complementary paths. On the one hand, existing improvements such as accounting for the viscoelastic behavior of RBC membranes and introducing deformation-based metrics are expected to advance further. These developments aim to represent the mechanical aspects of cell damage more realistically, with a

focus on accurately describing the forces acting on the membrane and the resulting rupture mechanics, while keeping computational demands within practical limits. On the other hand, a complementary path involves incorporating aspects of biological variability by expanding the mechanistic scope of hemolysis models beyond purely fluid force induced mechanical failure. This could include integrating additional biological markers or pathways associated with RBC damage, such as oxidative stress, metabolic alterations, or changes in membrane protein composition. Such extensions could improve the accuracy and physiological relevance of hemolysis models and help capture the multifactorial and heterogeneous nature of blood damage observed under clinical conditions.

To summarize, hemolysis modeling enables practical predictions of blood damage in medical devices by simplifying complex multiscale processes into macroscale formulations. These models, typically calibrated against *in-vitro* data, represent averaged responses and do not capture individual variability or the full stochastic nature of hemolysis. While stress-, strain-, and dissipation-based approaches offer different perspectives on RBC damage, all current models involve trade-offs between physical fidelity, generalizability, and computational cost. The selection of an appropriate model should therefore reflect the specific flow conditions, application goals, and required balance between detail and efficiency.

- Hemolysis models for blood-contacting devices rely on empirical macroscale formulations due to the computational cost of detailed simulations.
- Current models are fitted to averaged *in-vitro* data and lack reliability for predicting absolute hemolysis.
- Key directions for improvement include incorporating additional flow-related mechanisms and developing methods that reflect physiological variability more accurately.

2.3.2 Reduced order modeling

As discussed in Chapter 2.2.2, there is growing interest in using validated computational models to support regulatory decision-making. Frameworks such as the ASME V&V 40 [84] provide structured guidance for integrating modeling and simulation into the medical device approval process. Within this context, ROMs have emerged as a promising approach to support or complement current standard of *in-vitro* and *in-vivo* approaches, provided that their applicability, limitations, and uncertainties are well understood.
As outlined in Chapter 2.3.1, high-fidelity CFD simulations are widely used in hemolysis modeling to capture detailed flow structures and quantify shear stresses. These simulations form the basis for predicting shear-induced blood damage in RBPs. However, the computational cost associated with high-fidelity CFD remains a major constraint, particularly when simulations need to be repeated across multiple operating conditions, design variants, or parameter studies. This limits the practical use of CFD in iterative design processes, uncertainty quantification or regulatory applications.
Reduced-order modeling addresses this challenge by approximating the behavior of high-fidelity simulations using simplified, lower-dimensional representations. This modeling approach aims

to preserve the dominant physical features of the modeled system while substantially reducing computational effort. As a result, ROMs can accelerate simulation workflows and facilitate integration into broader design and decision-making processes. Their ability to produce fast, interpretable, and parameterizable outputs makes them particularly attractive in the development and analysis of medical devices such as RBPs [128]. In the context of CFD, techniques for ROM can generally be classified into two categories: intrusive and non-intrusive approaches [129].

Intrusive methods operate by modifying the governing equations of the CFD solver. These methods extract a low-dimensional subspace from full-order solutions and project the governing equations onto that subspace. While they can yield highly accurate and interpretable models, they require access to the CFD source code and extensive modifications, which makes them difficult to implement in commercial software environments. Under substantial effort, Song and Heuveline [130] demonstrated such an approach by developing an in-house finite element solver tailored to RBP simulations. Their work combined a variational multiscale turbulence model with an intrusive stochastic Galerkin method to quantify the effect of uncertain input parameters on blood flow and hemolysis risk.

Non-intrusive methods, in contrast, treat the CFD solver as a black box and construct a ROM using input and output data from selected simulation runs. These ROMs are trained to replicate the system response across the input space and can be evaluated at a negligible computational cost. This makes non-intrusive ROMs well suited for practical engineering applications where modifying the underlying solver is infeasible.

The general non-intrusive concept is illustrated in Figure 2.11. The approach begins with the definition of a set of input variables $X = [x_1, x_2, \ldots, x_n]$. In the context of RBP modeling these could be for example flow rate and rotational speed (panel A). Rather than evaluating the full CFD model $f(X)$ at every possible combination of inputs, which would be computationally prohibitive, the high-fidelity simulation is run only at selected sample points (grey dots), chosen to cover the input space efficiently.

In a second step at each of these sampled input combinations, the full-order model is used to compute the corresponding outputs $Y = f(X) = [y_1, y_2, \ldots, y_d]$, where the output might represent, for example, the predicted hemolysis index (panel B). The results reflect how the physical system responds to variations in the input parameters, and each simulation adds one data point to the growing dataset of input–output pairs. These data form the training basis for constructing a ROM, depicted as $\hat{f}(X)$, that approximates the true function $f(X)$ across the full range of inputs.

Finally, Panel C illustrates the resulting continuous system response surface $\hat{f}(X)$, which represents the ROM's prediction of system behavior based on the limited training data. Unlike the full CFD model, which must solve complex governing equations at each point, the surrogate $\hat{f}(X)$ allows rapid evaluations at new parameter values, enabling applications such as sensitivity analysis, optimization, or uncertainty quantification. The black crosses in the plot denote additional test points that are evaluated using the full CFD model to assess the surrogate's accuracy. These are not part of the training data and serve to quantify how well $\hat{f}(X)$ generalizes beyond the sampled inputs. The complexity of the surrogate model and the number of required training samples depend on several factors, including the dimensionality of the input space, the nonlinearity of the system response, and the desired accuracy [131]. Most ROM techniques are limited to predictions within the training domain and do not

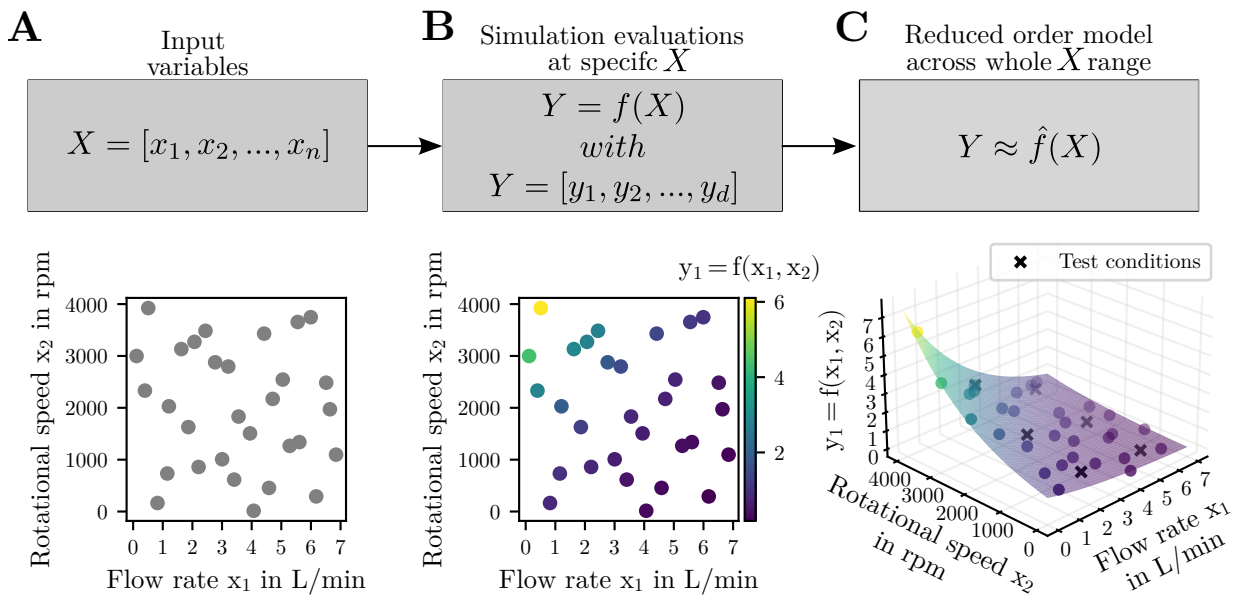

Figure 2.11: Illustration of the general non-intrusive ROM approach. Panel A defines the input variable space $X = [x_1, x_2, \ldots, x_n]$, where each point represents a combination of operating parameters. In case of a RBP this could be flow rate and rotational speed. Panel B shows selected high-fidelity simulation evaluations $Y = f(X)$ at discrete input points (colored by output value), where the output $Y = [y_1, y_2, \ldots, y_d]$ may represent system responses like hemolysis index. Panel C visualizes the resulting ROM model $\hat{f}(X)$, which approximates the full simulation response across the input space and can be efficiently evaluated. The black crosses indicate independent test points used to evaluate the ROM's predictive accuracy beyond the training data.

perform well outside of this range. Therefore, careful planning of the sampling strategy and the definition of the input parameter space are essential.
In addition, the choice of mathematical structure influences the performance and applicability of the ROM, as different formulations result in distinct modeling capabilities. Accordingly, non-intrusive ROMs can be classified into several main categories based on their underlying mathematical formulation:

- Parametric regression models approximate the response surface using a linear combination of basis functions, which may be predefined or selected adaptively. Models like Multivariate Adaptive Regression Splines [132] use piecewise-defined functions to capture local trends. These methods are efficient and easy to implement but may struggle with highly nonlinear behaviors.

- Spectral methods, including Polynomial Chaos Expansions [133], represent the output as a weighted sum of orthogonal polynomials defined over the input space. These models allow analytical computation of statistical quantities such as variance and sensitivity indices, but they become increasingly complex in high-dimensional spaces.

- Kernel-based models such as Gaussian Process Regression [134], Radial Basis Functions [135], and Support Vector Regression [136] use similarity kernels to interpolate

the training data. They are capable of representing complex functions with relatively few data points and include built-in measures of uncertainty, but their training cost grows significantly with the number of samples.

- Neural network-based models, such as multilayer perceptrons [137] and autoencoders [138], are highly flexible and can capture nonlinear relationships in large datasets. However, they require substantial training data and often lack interpretability, which may limit their applicability in regulated environments like RBPs.
- Finally, hybrid and multi-fidelity approaches combine information from models of varying resolution or incorporate gradient information of model outputs with respect to inputs to enhance accuracy. These methods can achieve high accuracy with relatively few high-fidelity evaluations, assuming sufficient correlation between model fidelities.

In the context of RBP simulations, ROMs must meet a number of constraints. These include computational efficiency, robustness under limited training data, and interpretability for purposes such as regulatory documentation. Spectral approaches, particularly Polynomial Chaos Expansions in the non-intrusive formulation (NIPCE), align well with these requirements. NIPCE is compatible with black-box CFD solvers, allows analytical computation of statistical measures, and performs well with relatively few training samples. Cheng *et al.* [139] benchmarked several surrogate methods and identified NIPCE as one of the most accurate and computationally efficient approaches for global sensitivity analysis in smooth, moderately low-dimensional problems. Its offline cost is limited to a regression fit of the polynomial coefficients (i.e., training the ROM using high-fidelity CFD results), while online evaluation (i.e., predicting outputs for new input values) is reduced to a simple weighted sum, enabling rapid prediction at negligible cost. For these reasons, the present work adopts NIPCE as the ROM framework, with its detailed implementation and validation discussed in Chapter 3.2.3. This fast evaluation capability makes ROMs particularly valuable for applications that require repeated model queries, such as sensitivity studies, optimization, or uncertainty quantification.

To summarize, reduced-order modeling offers an efficient and scalable approach for hemolysis modeling in RBPs. It enables rapid evaluation across a wide range of flow conditions and parameter settings, supporting tasks such as sensitivity analysis, uncertainty quantification, and model calibration. These capabilities make ROMs a valuable approach for extending and accelerating detailed hemolysis evaluations in complex device geometries.

- ROMs speed up hemolysis predictions by approximating high-fidelity CFD simulations across parameter spaces.
- ROMs enable sensitivity studies, optimization, or uncertainty quantification
- NIPCE is ideal for medical devices, combining solver compatibility with accuracy, efficiency, and interpretability.

2.3.3 Bayesian parameter estimation

As discussed in Chapter 2.1.1 and illustrated in Figure 2.9B, shear stress-induced hemolysis is influenced by a variety of factors, and current macroscopic hemolysis models are based on averaged fittings of experimental data, which fail to capture the underlying biological variability. Despite this averaging simplification, such models have proven to be useful and reliable tools for predicting relative differences between designs, thereby informing the development of RBPs and contributing to improved hemodynamic performance. However, comparison with experimental hemolysis data remains challenging, as measured hemolysis levels often exhibit substantial variability [92, 140] that current deterministic macroscopic models cannot capture.

This observation highlights a fundamental limitation of deterministic models: they produce single point estimates that cannot capture the inherent biological and experimental variability in hemolysis measurements. As a result, the reliability and interpretability of such models are limited, particularly when comparing results across different experimental conditions or donors. Probabilistic modeling frameworks, such as Bayesian inference, address this limitation by treating model parameters as uncertain quantities and updating them systematically as new data become available. Instead of yielding fixed values, Bayesian methods represent model parameters and predictions as full probability distributions, allowing variability and uncertainty to be explicitly quantified and propagated through the model [141]. In the

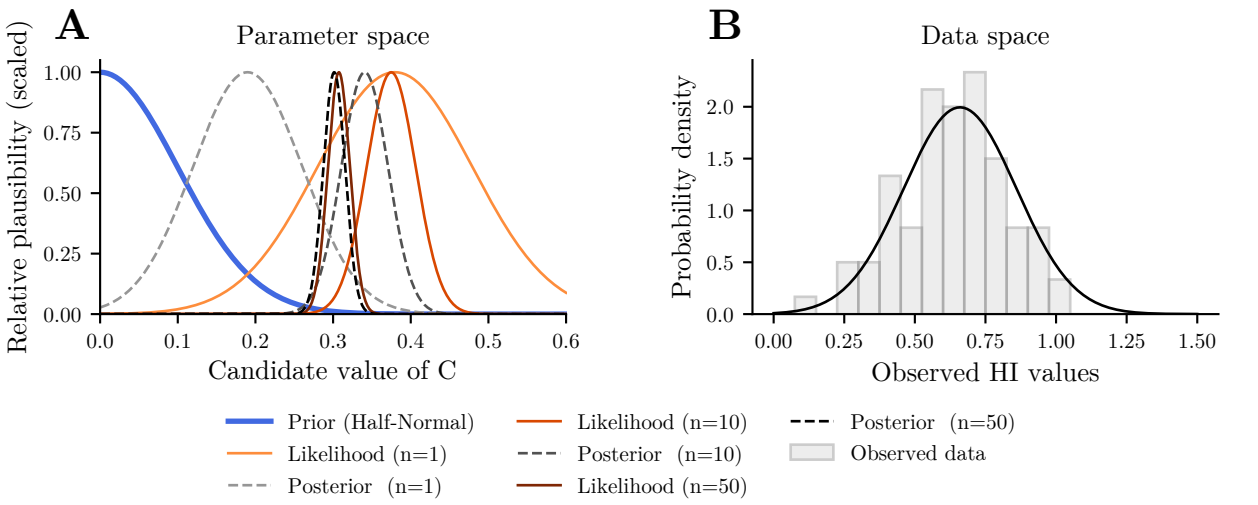

Figure 2.12: Illustration of Bayesian inference for parameter estimation in hemolysis modeling. Panel A shows the evolution of the posterior distribution for a model parameter C as the number of hemolysis data points increases. The prior (blue) represents initial beliefs, while likelihood functions (orange) and resulting posteriors (grey to black dashed lines) reflect how data updates those beliefs. As the sample size n increases, the posterior distribution becomes narrower, indicating reduced uncertainty about the parameter. Panel B illustrates the observed distribution of hemolysis index (HI) values in the data space (grey bars).

context of hemolysis modeling (*see* Chapter 2.3.1), this can be illustrated using the power law formulation (*see* Equation 2.3), which describes hemolysis as a function of shear stress and exposure time. This section focuses on an illustrative example of estimating the fitting parameter C, which scales the hemolysis response, while keeping the parameters α and β

and shear stess and exposure time fixed for illustrative purposes.
The conceptual workflow of Bayesian parameter estimation, with it's three main components prior, likelihood and posterior, is visualized in Figure 2.12, where the goal is to estimate the unknown parameter C based on experimental observations of HI values. These observations vary due to biological and measurement-related noise and form the basis for updating prior beliefs about plausible values of C.
The **prior** distribution encodes domain knowledge or assumptions about the parameter C before any data are observed. For example, in the context of the power law model, it is known that C must be strictly positive, as negative values would imply physically meaningless negative hemolysis. Such constraints can be formalized through the prior by assigning higher probability density to plausible values and near-zero density to implausible regions (e.g., using a half-normal distribution). This initial assumption is represented in Figure 2.12A as a broad blue curve covering the parameter space.
The **likelihood** function quantifies how well different candidate values of the parameter C explain the observed data. Specifically, it represents the probability of obtaining the observed hemolysis values given a particular parameter value. Since the data exhibit biological and experimental variability, the likelihood function reflects how consistent each candidate value of C is with that observed variability. In Figure 2.12A, the orange curves show the likelihood functions for increasing sample sizes ($n = 1, 10, 50$). As more data points are included, the likelihood becomes increasingly peaked around the value of C that most plausibly explains the data.
The **posterior** distribution results from combining the prior and the likelihood using Bayes' theorem:

$$\text{Posterior}(C \mid \text{data}) = \frac{\text{Likelihood}(\text{data} \mid C) \times \text{Prior}(C)}{\text{Evidence}(\text{data})} \tag{2.9}$$

In this expression, the vertical bar "|" denotes a conditional relationship. For example, $\text{Posterior}(C \mid \text{data})$ represents the updated probability distribution of the parameter C after incorporating the observed data. The numerator combines the prior belief about C with the likelihood of observing the data given a particular value of C. The denominator, referred to as the Evidence, serves as a normalization constant to ensure that the posterior distribution integrates to one [142]. In Figure 2.12A, the grey dashed curves show the resulting posterior distributions for increasing sample sizes. As more data are considered, the posterior becomes narrower and moves toward the region of the parameter space favored by the data. While the true value C is not explicitly marked in the figure, the convergence of the posterior reflects the process of learning from data, in which uncertainty about the parameter is progressively reduced as more information becomes available.
It is important to note that the likelihood and posterior distributions shown in Panel A address a different question than the data distribution shown in Panel B. While the posterior characterizes the remaining uncertainty about the parameter C after incorporating the available data, the data distribution in Panel B illustrates the variability of the hemolysis measurements themselves. This distinction is essential because, even if the true value of C were known, the hemolysis measurements would still vary due to biological and measurement noise. While incorporating more data reduces uncertainty about the parameter C, as reflected by the narrowing posterior, it does not eliminate the intrinsic variability of the measurements. That

is, although the estimate of the parameter becomes more precise, the spread of the underlying measurement distribution remains governed by inherent biological and experimental variability. In practice, however, it is important to recognize that the number of experimental data points available for hemolysis modeling is typically limited. Published datasets often contain a fixed number of measurements, and no new data can be collected retrospectively. Consequently, the posterior distribution derived from such datasets reflects the full extent of parameter uncertainty that can be inferred from the available data. This constraint must be kept in mind when interpreting the results of Bayesian inference in this context.

While the conceptual framework of Bayesian parameter estimation offers a powerful tool for incorporating uncertainty, the practical computation of posterior distributions is often challenging. In practice, computing the posterior distribution analytically is only possible for a small class of simple models where so-called conjugate priors can be applied. Conjugate priors are specific choices of prior distributions that, when combined with the likelihood, result in a posterior distribution of the same mathematical form. While this property greatly simplifies calculations, it is limited to very specific cases and assumptions, which are not applicable to complex real-world problems [142]. In macroscopic hemolysis modeling, the nonlinear nature of the models, the presence of measurement noise, and the multidimensional parameter spaces make analytical solutions intractable. Therefore, numerical methods are required to approximate the posterior distribution and to compute quantities of interest.
Several computational strategies have been developed to enable Bayesian inference in cases where analytical solutions are not feasible. Depending on the characteristics of the problem and the available computational resources, different algorithms can be employed to sample from or approximate the posterior distribution. Among these, MCMC methods are widely used and serve as the foundation for many modern Bayesian inference techniques.
To introduce the basic idea of MCMC, the Metropolis-Hastings algorithm [143] is frequently used as a foundational example. It provides an intuitive understanding of how sampling from a posterior distribution can be achieved when direct analytical computation is not feasible. MCMC methods rely on the concept of a Markov chain, which is a sequence of random values where the probability of moving to the next value depends only on the current value, not on the full history of the chain. This property makes Markov chains well suited for exploring complex parameter spaces, because the algorithm does not need to store or analyze the entire sample history in order to generate new values. When carefully designed, a Markov chain can be constructed such that its long-run distribution matches the posterior distribution of interest. The core idea of Metropolis-Hastings is to construct such a chain by generating a sequence of candidate parameter values and deciding, at each step, whether to accept or reject the proposed value based on how well it agrees with the observed data and prior information. To illustrate the process in the context of hemolysis modeling, a simplified example is provided for estimating the parameter C in the power law equation. A prior distribution for C is assumed, and experimental hemolysis data are used to define the likelihood. The steps of the Metropolis-Hastings algorithm are shown in the pseudocode below:

Algorithm 1 Metropolis-Hastings algorithm for estimating the parameter C

Choose an initial value C_{current}
Set step size σ
for $i = 1$ to N **do**
    Propose $C_{\text{new}} \leftarrow C_{\text{current}} + \varepsilon$, where $\varepsilon \sim \mathcal{N}(0, \sigma^2)$
    Compute acceptance ratio:

$$r = \frac{\text{Likelihood}(C_{\text{new}}) \cdot \text{Prior}(C_{\text{new}})}{\text{Likelihood}(C_{\text{current}}) \cdot \text{Prior}(C_{\text{current}})}$$

    Draw $u \sim \text{Uniform}(0, 1)$
    if $u < r$ **then**
        $C_{\text{current}} \leftarrow C_{\text{new}}$
    end if
    Store C_{current} as sample i
end for

Each stage of the algorithm serves a distinct purpose:

- **Initialization:** A reasonable starting value for the parameter C is chosen, typically based on domain knowledge or centered within the prior distribution.
- **Proposal mechanism:** A new candidate value C_{new} is proposed by adding Gaussian noise $\varepsilon \sim \mathcal{N}(0, \sigma^2)$ to the current value. The step size σ controls how far the new value can deviate. Smaller values result in slow but stable exploration; larger values enable faster movement but may increase rejection rates.
- **Acceptance ratio computation:** The acceptance ratio r is calculated as the ratio of the unnormalized posterior probabilities of the new and current parameter values. It reflects how well the proposed value explains the data and how consistent it is with the prior.
- **Acceptance decision:** A random number $u \sim \text{Uniform}(0, 1)$ is drawn. If $u < r$, the proposed value is accepted; otherwise, the current value is retained. This probabilistic decision process prevents the chain from getting stuck and ensures proper exploration.
- **Storage:** The current value of C is recorded as a sample from the posterior. Repeating this process over many iterations produces a sequence of values that approximate the target posterior distribution.

Over many iterations, the collected values form a sample from the posterior distribution of C. After discarding an initial burn-in period, the remaining samples can be used to estimate posterior means, medians, credible intervals, and other quantities of interest.
This basic MCMC procedure illustrates the core concepts of numerical Bayesian parameter estimation and serves as a foundation for understanding more advanced sampling algorithms. In this dissertation, Bayesian inference is performed using the No-U-Turn Sampler (NUTS) [144], an extension of Hamiltonian Monte Carlo that automatically adjusts step sizes and trajectory lengths during sampling. Due to its improved efficiency in high-dimensional spaces

and its ability to eliminate manual tuning, it is particularly well suited for the parameter estimation task addressed in Chapter 6. A detailed explanation of the NUTS algorithm is provided in the corresponding methodological section 6.2. Nevertheless, to briefly highlight

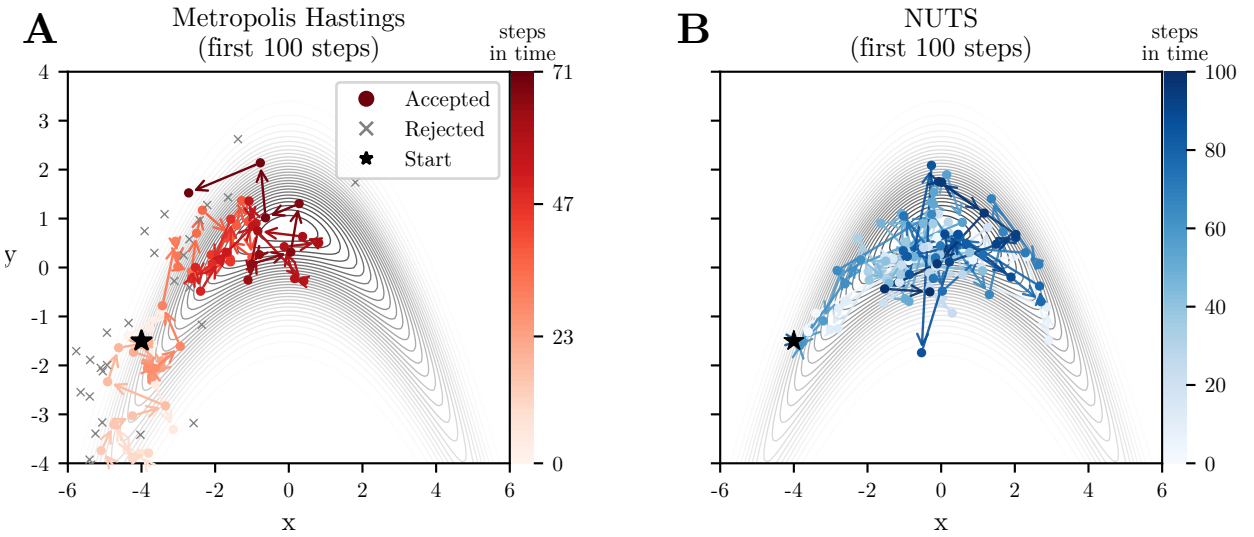

Figure 2.13: Visualization of the first 100 sampling steps for two MCMC algorithms applied to a two-dimensional posterior distribution with curved geometry. Panel A shows the Metropolis-Hastings sampler, with accepted samples (colored dots), rejected proposals (grey crosses), and the starting point (black star). The high rejection rate and limited spread of accepted samples indicate poor coverage of the target distribution. Panel B displays the No-U-Turn Sampler (NUTS), showing accepted samples and the starting point, with improved coverage of the distribution. Contour lines represent the target distribution, while colorbars (white to red/blue) indicates sampling progression over time.

the practical differences in sampling behavior, Figure 2.13 compares the performance of the Metropolis-Hastings algorithm and the NUTS on a two-dimensional illustrative example with highly correlated parameters x and y. This correlation is represented by a two-dimensional distribution with a characteristic banana-shaped curvature (grey contour-lines). In Panel A, Metropolis-Hastings is initialized at the black star. The colored path represents the sequence of accepted proposals, while grey crosses denote rejected ones. The algorithm explores the posterior by proposing random steps from a symmetric proposal distribution. Moves toward regions of higher probability are more likely to be accepted, while moves away from them are often rejected. As a result, the chain gradually builds up more density in high-probability regions. However, when the posterior has strong curvature or elongated geometry, as in this case, exploration becomes inefficient. A small step size must be used to maintain a reasonable acceptance rate, leading to slow movement through the parameter space and a high number of rejections. This is evident from the cluster of rejected points and the limited coverage of the posterior after 100 proposals.

In contrast, the Panel B shows the performance of the NUTS algorithm on the same posterior, starting from the same initial point. In this approach, the parameter values of the model are treated as the position of a particle moving through a potential energy landscape defined by the posterior distribution. The intuition behind this method is that, instead of proposing random steps, NUTS simulates the movement of a particle under physical laws that cause it to follow smooth and energy-conserving trajectories. These trajectories allow the sampler to

move efficiently through regions of high probability, particularly when the posterior exhibits strong correlations between parameters.
Building on this physical intuition, NUTS introduces an adaptive mechanism that determines the appropriate trajectory length automatically during sampling. It uses the so-called "no U-turn" criterion, which halts the trajectory once it begins to reverse direction toward its starting point. This prevents redundant exploration and enhances computational efficiency. One point along the valid portion of the trajectory is then selected as the next state in the Markov chain. By drawing from an entire simulated trajectory rather than proposing a single random step, NUTS naturally incorporates exploration into each iteration. This procedure enables the sampler to travel across complex posterior landscapes in a single iteration, resulting in more efficient exploration and faster convergence compared to traditional random-walk approaches.

To summarize, Bayesian parameter estimation offers a systematic approach to quantify uncertainty in hemolysis model parameters by integrating prior knowledge with experimental data. Instead of relying on single point estimates, it produces full probability distributions that capture the range and likelihood of plausible parameter values. This enables a more transparent treatment of biological variability and measurement noise, enhancing the model's interpretability and reliability.

- Biological and experimental variability in hemolysis data limits the reliability of deterministic models and motivates probabilistic approaches.
- Bayesian parameter estimation quantifies uncertainty by treating model parameters as probability distributions informed by experimental data.
- Advanced sampling algorithms like NUTS enable efficient estimation of posterior distributions, even in strongly correlating and high-dimensional parameter spaces.

3 Efficiently computing a wide range of pump operating points

The general findings of this chapter are based on the article "Systematic analysis of non-intrusive polynomial chaos expansion to determine rotary blood pump performance over the entire operating range" published in *Computers in Biology and Medicine* [145].

3.1 Introduction

RBPs used in ECMO support must operate across a wide range of clinically relevant conditions, including flow rates from 0 to 7 L/min and pressures up to 400 mmHg [9, 10]. These requirements reflect the diverse therapeutic goals and patient-specific needs encountered in clinical practice, as further detailed in Chapter 2.1.2. Ensuring safe and effective RBP operation across this broad spectrum is critical. This applies not only to clinical outcomes but also to regulatory assessment, as outlined in Chapter 2.2.2. As Class III medical devices, RBPs must demonstrate consistent safety and functionality across the full intended operating range. Typically, this involves the evaluation of multiple representative operating conditions to ensure robust performance and hemocompatibility. This is commonly achieved through extensive *in-vitro* and *in-vivo* testing, which can be both time-consuming and costly. As a result, *in-silico* methods are increasingly being explored as a promising complement, offering potential to reduce experimental burden while maintaining rigorous performance evaluation [81, 83].
To support such efforts, the FDA endorses the ASME V&V 40 standard [84], which provides a risk-informed framework for VVUQ. The level of rigor required in VVUQ efforts depends on the regulatory impact of the computational model and the potential consequences of incorrect predictions. As shown in Chapter 2.3.1, assessing RBP performance, including hydraulic characteristics and shear-induced blood damage, generally relies on high-fidelity CFD simulations that resolve detailed flow structures. However, each new operating point requires a separate simulation, leading to substantial computational costs. This limitation poses a major obstacle to a comprehensive evaluation across the full operating point range of RBPs and motivates the first research question of the dissertation, which is addressed in the present chapter:

1. *How can fluid variables across a wide range of RBP operating points be computed efficiently?*

Addressing this question is essential to enable efficient performance evaluation and to support hemocompatibility modeling under realistic and clinically relevant conditions.
To overcome this challenge, various ROM techniques have been explored across scientific and engineering disciplines, aiming to reduce the number of required simulations. Notably, Cheng *et al.* [139] compared several ROM approaches and identified NIPCE as one of the most accurate and computationally efficient methods. Their comprehensive evaluation provides a well-established basis for selecting NIPCE in the present chapter, allowing the focus to shift from comparing ROM techniques to applying the most promising method to the specific context of efficiently evaluating RBPs across the full clinical operating point range in ECMO therapy. While NIPCE has seen frequent use in general turbomachinery applications [146–149], its application to RBPs, and particularly the modeling of the entire operating point range, remains relatively unexplored. For instance, Song and Heuveline [130] employed intrusive polynomial chaos expansion to adjust RBP parameters, but did so within a limited operating point range. Similarly, Mohammadi *et al.* [150] applied NIPCE to the FDA benchmark RBP geometry [92] without exploring parameter variation across the full operating point range. These prior studies also lack systematic test simulations, leaving open questions regarding the influence of individual parameters on model performance. Furthermore, the impact of modeling parameter choices on the performance of NIPCE-based ROMs has not yet been systematically investigated in the context of RBPs.
To summarize, this approach has the potential to enable efficient computation of RBP performance across a wide clinical range, which is increasingly relevant given the growing role of *in-silico* evidence in the regulatory process and the required use of advanced numerical methods, such as uncertainty quantification, as outlined in ASME V&V 40 [84]. Based on insights from the literature identifying NIPCE as a high-performing ROM technique, the present chapter systematically investigates its application to RBPs, with a particular focus on the influence of NIPCE modeling parameters. The results are validated against comprehensive test data, and practical guidelines are derived to support reliable performance evaluation and hemocompatibility modeling.

3.2 Materials & Methods

Expanding on the ROM fundamentals discussed in Chapter 2.3.2, this chapter introduces a structured methodology for training, applying, and validating a NIPCE-based ROM aimed at efficiently predicting RBP performance over a broad operating range.

3.2.1 Geometries, mesh and simulation parameters

To test the ROM method over the entire operating range of a RBP, the FDA benchmark geometry for RBPs was chosen [92]. This RBP geometry is freely available, and extensive validation data from several groups exist in the literature [8, 151]. Validating the CFD model with experimental data reduces uncertainty in the model, which is essential for training the ROM. This validation of the CFD model was carried out in a previous investigation by Gross-Hardt *et al.* [152]. For details of the mesh and simulation parameters, reference is made to that work. In summary, the geometry was meshed using ANSYS 2021 R1 Meshing tool

(ANSYS Inc., Canonsburg, USA) with unstructured tetrahedral elements and prism layers. This resulted in a total mesh element count of 9.3 million cells, corresponding to the "middle" mesh described in [152]. ANSYS CFX was used as the CFD solver and the turbulence model selected was k-ω shear stress transport.
To reduce the simulation effort required for training data generation, steady-state simulations were performed, in contrast to the transient approach adopted by Gross-Hardt *et al.* [152]. Comparison with validation data [8] for the pump's pressure head indicated that this simplification was justified, given the minor differences from unsteady simulations and validation data. The pressure head across the RBP was determined by the area-averaged pressure difference between pump inlet and outlet. The axial force was defined as the force acting on the impeller in the direction of the rotational axis. Hydraulic efficiency was calculated using the following equation:

$$\eta = \frac{\dot{m}\,\Delta p}{\rho\,P_{\text{shaft}}} \tag{3.1}$$

with $\dot{m}$ denoting the inlet mass flow rate, Δp the pressure head over the pump, ρ the fluid density, and P_{shaft} the shaft power, which is influenced by rotational speed and impeller torque.
As outlined in Chapter 2.3.2 by the overview Figure 2.11, the underlying data used to train the ROM is provided by the CFD model. The rotational speed and flow rate of the CFD simulations were treated as stochastic variables to enable coverage of the entire operating point range of the RBP. The term "stochastic variable" is commonly used in NIPCE literature to describe input variables that follow a defined probability distribution. In the present chapter, a uniform probability distribution was assumed for both variables, as this ensures equal exploration of the parameter space and promotes full coverage of the operating range. The results of the CFD simulations defined the quantities of interest (QoI), which included the pressure head, axial force, hydraulic efficiency, and fluid velocity. In principle, additional QoI could be considered for training the NIPCE model. For instance, variables derived from numerical hemolysis models may also be of interest when assessing blood damage potential. However, these were excluded in the current chapter, as the definition of hemolysis metrics involves division by flow rate. As the flow rate approaches zero, these expressions tend toward infinity, making them unsuitable for inclusion in an analysis spanning the full theoretical operating range. This limitation will be addressed in Chapter 4, where hemolysis-relevant QoI under clinically realistic conditions are analyzed using a ROM.

3.2.2 Training and test data

Creating a ROM across the operating range of a RBP required the definition of the parameter space and the probability distribution of the stochastic variables. The interval [0–5000 rpm] was defined for the rotational speed and the interval [0–7 L/min] for the flow rate. To train the ROM method on this operating range, training data were required, i.e., CFD simulations solved at specific operating points within the defined parameter space. In order to distribute these points as evenly as possible, the Latin hypercube sampling algorithm was applied [153]. This algorithm uses a quadratic grid over the parameter space and the probability distribution of the stochastic variables to select representative operating points that evenly cover the entire domain. The number of training data points required for the NIPCE method was initially

determined to be 30, based on the recommendation of Hosder *et al.* [131]. This number results from the polynomial order used ($p = 4$) and the number of stochastic variables ($d = 2$). Further details are provided in the following subchapter 3.2.3. In addition to the training data points, test data points were generated at flow rates of 0, 2, 4, 6 and 8 L/min, combined with rotational speeds of 1000, 3000, 5000 and 6000 rpm. These 20 simulations served to assess the predictive accuracy of the NIPCE method both within the training range and under extrapolation conditions, such as 6000 rpm and 8 L/min. As they were not used during the training process, the test data provide an independent basis for validation. An overview of the distribution of training and test data points is shown in Figure 3.1. To investigate the effect

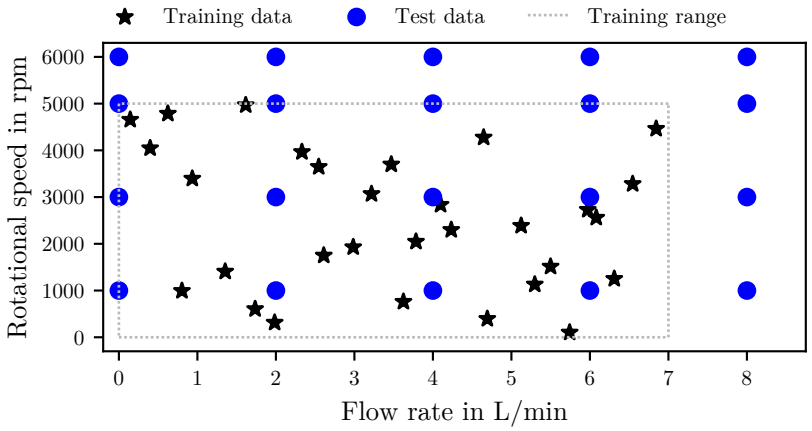

Figure 3.1: Distribution of training and test data points across the operating space. The training range spans rotational speeds from 0–5000 rpm and flow rates from 0–7 L/min, indicated by the dotted grey box. Training data (30 points) are marked as black stars, and test data (20 points) as blue dots.

of the number of training data points on the predictive power of the NIPCE method, subsets were generated from the original set of 30 training data points. Subsets consisting of 6, 12 and 20 data points, in addition to the full 30-point set, were compared. When creating these subsets, the minimum distance between points was maximized to ensure that each subset covered the parameter space as uniformly as possible. For this purpose, the training data were scaled between zero and one in both rotational speed and flow rate, and the distances were computed using the Euclidean distance. An algorithm calculated the Euclidean distances for all possible subsets of 6 points and selected the subset with the greatest minimum distance between points. After the first 6-point subset was identified, only subsets containing the already selected points were considered for the 12- and 20-point sets. This procedure ensured that each larger subset built upon the previous one while maintaining maximal spacing.

3.2.3 Non-intrusive polynomial chaos expansion (NIPCE)

There are several approaches available to explore the impact of varying operating conditions on system performance. A widely used method is the Monte Carlo approach [154], in which random samples are drawn from the input parameter space and individually simulated. While conceptually simple, this method requires a large number of high-fidelity CFD runs to produce statistically meaningful results, making it impractical for comprehensive analyses

involving expensive simulations. To reduce this computational burden, various reduced-order modeling techniques have been proposed. Among these, NIPCE has demonstrated particularly strong performance in a comparative study [139]. NIPCE enables efficient prediction over a parameter space based on a limited number of CFD simulations, offering a substantial reduction in computational effort. In NIPCE, the CFD solver is treated as a black-box model. The unknown quantity of interest, Y, is represented by a function f of uncertain input variables $X = [x_1, x_2, \ldots, x_d]$ (*see* equation (3.2)).

$$Y = f(X) \tag{3.2}$$

For the scenario considered in this chapter, Y would be, for example, the pressure head over the pump with x_1 being the rotational speed and x_2 the flow rate. Since a single evaluation of this black box equation involves solving a CFD model, it is desirable to require as little evaluations of this equation as possible. Therefore, a ROM for this equation is created below. Assuming that the input variables are independent of each other, Y can also be expressed by orthogonal polynomials as follows:

$$Y \approx \sum_{k=1}^{P} a_k \, \psi_k(X) \tag{3.3}$$

Where $\psi_k(X)$ is a multivariate orthogonal polynomial consisting of a linear combination of univariate orthogonal polynomials $\psi_{k_1}(x_1) \, \psi_{k_2}(x_2)$. The approximate-equal symbol is used because the polynomial chaos expansion is truncated at a finite order, denoted by p. The total number of terms in the expansion, P, depends on both the polynomial order p and the number of input variables d, as defined in equation (3.4):

$$P = \frac{(p+d)!}{p! \, d!} \tag{3.4}$$

The impact of these polynomials on the sum in equation (3.3) is determined by their coefficients, denoted as a_k. The polynomial basis can be freely chosen by the user, while an optimal polynomial basis exists depending on the probability distribution of the input parameters called Wiener-Askey scheme [155]. In the case of uniformly distributed stochastic variables, Legendre polynomials provide the appropriate orthogonal basis and were therefore selected for this chapter. To find the unknown coefficients a_k in equation (3.3), a common approach is the point collocation method, which is based on linear regression. Assuming that there are N evaluations of equation (3.3), a linear system of equations can be derived as shown in equation (3.5).

$$\begin{bmatrix} \psi_1(X_1) & \cdots & \psi_p(X_1) \\ \psi_1(X_2) & \cdots & \psi_p(X_2) \\ \vdots & \ddots & \vdots \\ \psi_1(X_N) & \cdots & \psi_p(X_N) \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \\ \vdots \\ a_p \end{bmatrix} = \begin{bmatrix} f(X_1) \\ f(X_2) \\ \vdots \\ f(X_N) \end{bmatrix} \tag{3.5}$$

Each entry on the right-hand side ($f(X_i)$) represents an evaluation of a CFD simulation with the stochastic input variables X_i. The required coefficients are determined by solving the

least square optimization:

$$a = \underset{a \in \mathbb{R}^p}{\operatorname{argmin}} \sum_{i=1}^{n} \Big(f(X_i) - a^T \psi(X_i) \Big)^2 \tag{3.6}$$

To determine the necessary number of training data points, a rough guideline suggesting an oversampling parameter of ≥ 2 was adopted based on the study by Hosder *et al.* [131]. The number of training data points depends on the number of stochastic variables and the polynomial order used. Based on these recommendations, the number of required training data points was determined for various combinations of p and d and is presented in Table 3.1. With

Table 3.1: Number of training data points with respect to polynomial order p and number of stochastic variables d according to Hosder *et al.* [131].

	p							
d	1	2	3	4	5	6	7	8
1	4	6	8	10	12	14	16	18
2	6	12	20	30	42	56	72	90
3	8	20	40	70	112	168	240	330
4	10	30	70	140	252	420	660	990

increasing polynomial order and increasing number of stochastic variables, the simulation effort for generating the training data points increases enormously. In this chapter, the assumption of using 30 training data points is tested by comparing the resulting predictions of the NIPCE method across varying polynomial orders and different numbers of training data points. For this the programming language Python was used to perform all necessary tasks for building the NIPCE model, including the creation of the polynomial basis, computation of the expansion coefficients using the point collocation method, and generation of training data via the Latin hypercube algorithm, all implemented with the open-source package *chaospy* [156]. The subsequent evaluation of the NIPCE method and the figure creation was also carried out using Python and the common packages numpy [157], pandas [158] , scipy [159] and matplotlib [160]. For easy reproducibility of the results presented in this chapter, both the raw data and the script to create the NIPCE-based ROM can be found in the corresponding publication [145].

3.2.4 Systematic investigation

To ensure a structured and transparent analysis, the present chapter follows a systematic approach consisting of three key stages: creation of the ROM, evaluation of its predictive performance, and sensitivity analysis with respect to model parameters. This methodology is illustrated schematically in Figure 3.2.
In the first step, the NIPCE-based ROM is created. As described in the previous sections, the polynomial chaos expansion is trained using 30 CFD simulations as training points and a polynomial order of 4, following the formulation proposed by Hosder *et al.* [131].
Subsequently, the predictive accuracy of the trained model is evaluated. For this purpose, a pressure plane in the rotational speed-flow rate space is generated from the CFD training

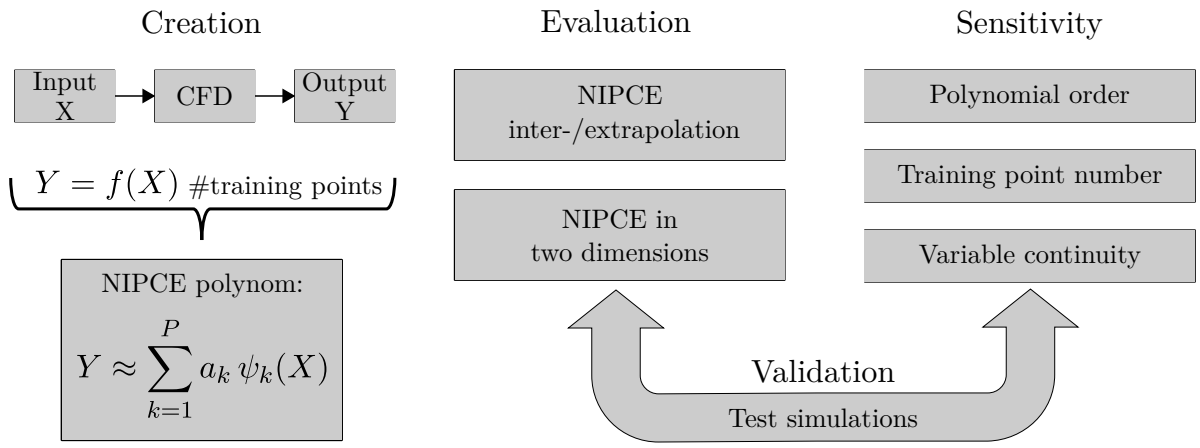

Figure 3.2: Schematic representation of the systematic investigation approach. First, the reduced-order model (ROM) is created using CFD-generated training data and a non-intrusive polynomial chaos expansion (NIPCE). Second, the predictive accuracy of the model is evaluated through inter-/extrapolation and comparison of two-dimensional flow fields. Third, a sensitivity analysis is conducted by varying the polynomial order, the number of training points, and the continuity of the input variables. All evaluations are validated against reference test simulations.

data using both linear and cubic interpolation. These interpolated surfaces are compared to the NIPCE prediction to assess agreement not only at the training points but across the entire rotational speed–flow rate domain. To further evaluate extrapolation capabilities, HQ-curves (*see* Figure 2.6) are used to quantify prediction accuracy both inside and outside the training range. Additionally, the model's ability to reproduce spatial flow information is assessed by comparing two-dimensional velocity distributions from a test CFD simulation and the NIPCE prediction at an operating point of 4 L/min and 3000 rpm.
Finally, the sensitivity of the NIPCE prediction to changes in model parameters is examined. This includes variation of the polynomial order (1, 2, 4 and 8) while keeping the number of training points constant at 30, as well as variation in the number of training points (6, 12, 20 and 30) at a fixed polynomial order of 4. In a final step, sensitivity to the continuity of the underlying data is investigated by comparing NIPCE predictions for a smooth QoI with those for a discontinuous QoI. Test simulations are used as a reference for all evaluations.

3.3 Results

This chapter presents the results of applying the NIPCE method to model RBP performance across the full theoretical operating range. The predictive capabilities of the method are evaluated and compared to conventional approaches, followed by a systematic investigation of factors influencing model accuracy.

3.3.1 NIPCE in comparison to interpolation

Figure 3.3 displays the pressure head prediction results of the NIPCE method in comparison to classical interpolation approaches. Panel A shows the location and pressure head of the

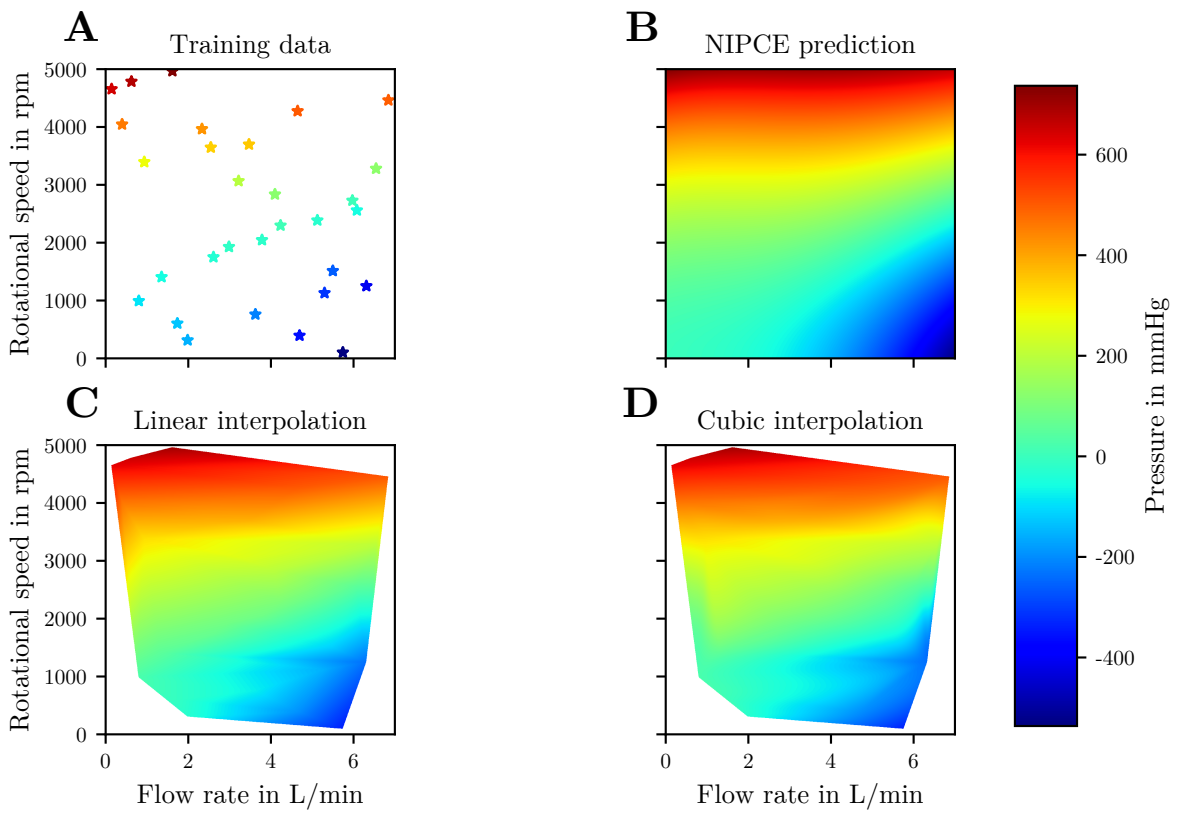

Figure 3.3: Distribution of the pressure head variable within the training range. Panel A shows the training data as stars colored by pressure head. Panel B presents the non-intrusive polynomial chaos expansion (NIPCE) prediction of pressure head over the whole training domain. Panels C and D display linear and cubic interpolation, respectively, based on the same training data. All interpolations are visualized within the convex hull of the training points.

underlying training data points in the rotational speed–flow rate space, represented as colored stars. Based on these training points, the quasi-continuous pressure distribution across the full operating range is computed using the NIPCE method, as illustrated in Panel B. To achieve this quasi-continuous representation, the method was evaluated on a 1000×1000 grid, yielding pressure head values at one million distinct operating points. A comparison to linear and cubic interpolation of the training data is shown in Panels C and D, respectively. It should be noted that interpolation is only possible within the convex hull of all training data points.

Using the data on rotational speed, flow rate, and pressure head, a standard pressure–flow rate (H-Q) curve diagram is generated, as shown in Figure 3.4. The underlying training data points are marked with black stars and their convex hull in the pressure-flow rate plane is drawn as dotted grey line. In addition, the prediction of the NIPCE method for the lines of constant rotational speed is shown by solid lines. The lines of constant rotational speed from the cubic interpolation are shown with dashed lines. To estimate the quality of the prediction, the test simulation data for the corresponding speeds are also drawn using circles, crosses, squares and triangles. The lines of constant rotational speed of the interpolation are only possible in the convex hull and deviate noticeably from the test data, especially in the lower

speed range. The NIPCE prediction, on the other hand, matches the test data well and can also provide good agreement beyond the training data range (6000 rpm & 8 L/min). In the upper right corner the mean average error (MAE) in mmHg between NIPCE prediction and test data indicates the quality of the prediction both inside the modeled range of 0–5000 rpm and 0–7 L/min, as well as outside of it.

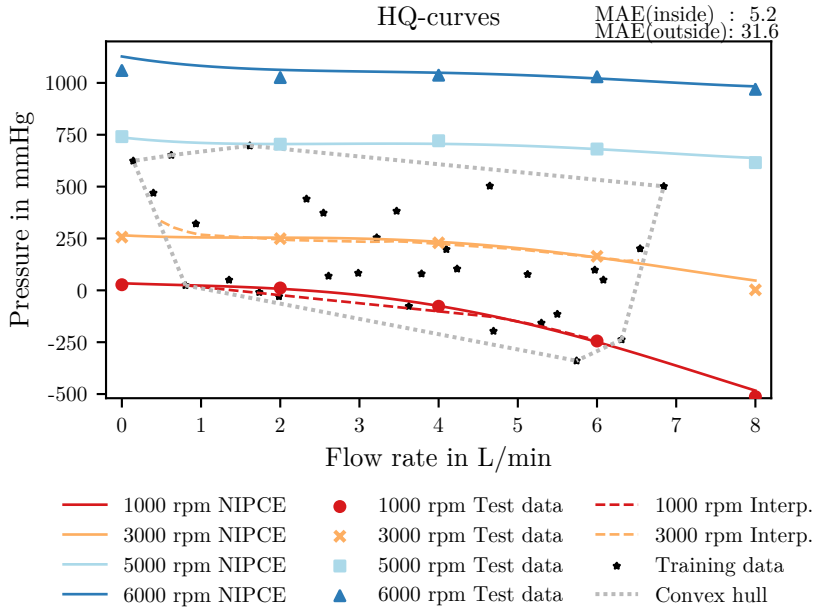

Figure 3.4: Pressure-flow rate curves from cubically interpolated training data inside the convex hull (dashed) and the prediction of the pressure-flow rate curves of the non-intrusive polynomial chaos expansion (NIPCE) method inside and outside the training area (solid). Test data for the respective curves are shown with circles, crosses, squares and triangles in the corresponding color and the mean absolute error (MAE) of NIPCE prediction to test data inside and outside the training range is shown in the upper right corner of the graph.

3.3.2 NIPCE in two dimensions

In addition to the one-dimensional consideration of the black box system from equation (3.1), Y is not limited to a global variable such as the pressure head considered so far. In principle, it is also possible to make a two-dimensional prediction if the training data is also available in two dimensions. Thus, it is possible to export a plane for each training simulation and construct a polynomial expansion for every individual mesh point using the entire training dataset. For this two-dimensional case, the NIPCE method was trained using the complete set of available operating points. As an example of its predictive capabilities, Figure 3.5 illustrates the results for one representative operating point at 3000 rpm and 4 L/min. Panel A shows the fluid velocity predicted by NIPCE at this condition, while Panel B presents the corresponding CFD simulation result. Panel C displays the pointwise difference between the two fields. Overall, the flow velocity in the plane is predicted with high accuracy, with a MAE of 0.1 m/s across all mesh elements. The observed deviations, particularly in the diffuser region, are consistent across all operating points and reflect a general pattern in the prediction behavior.

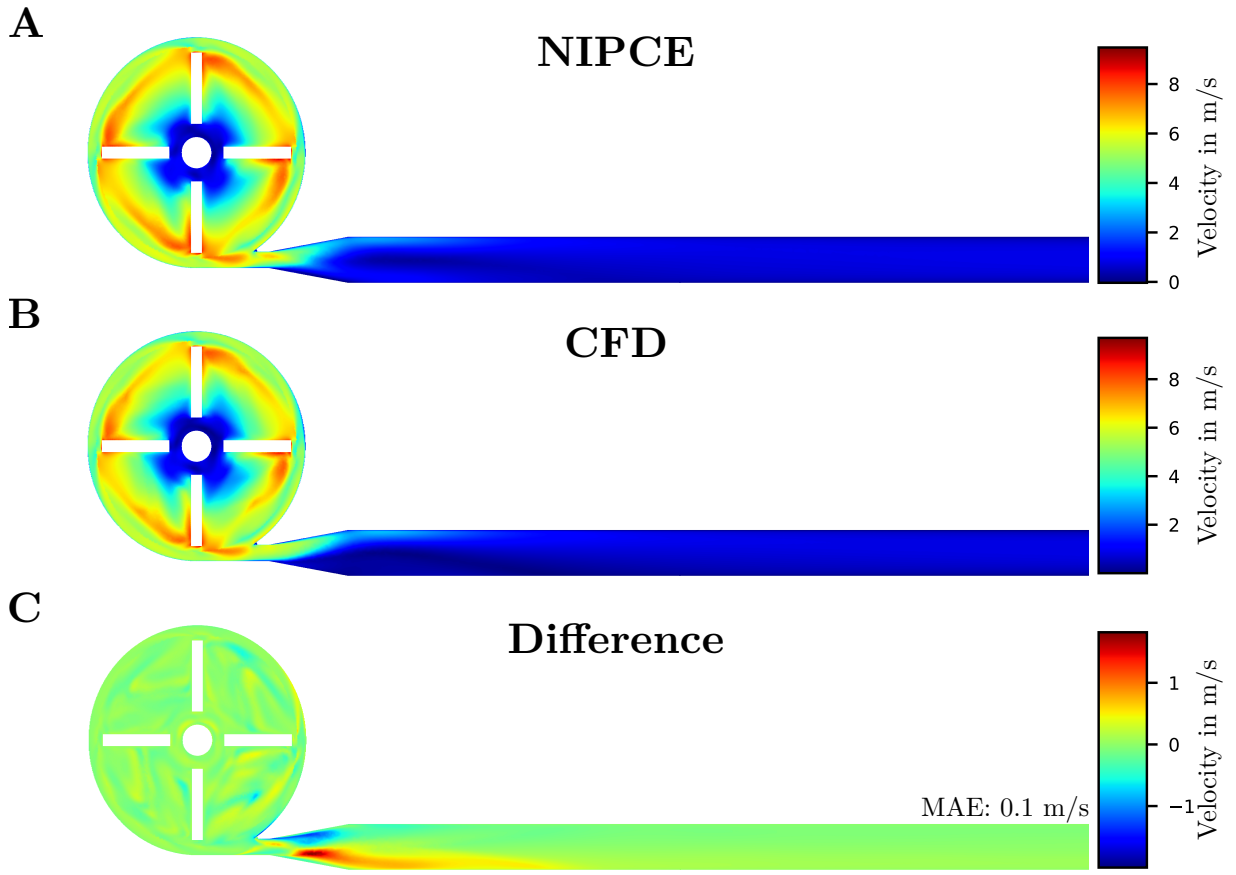

Figure 3.5: Two-dimensional NIPCE prediction of the velocity field in Panel A, shown on the plane $Z = 0.0065\,\mathrm{m}$ of the FDA RBP benchmark geometry at an operating point of 3000 rpm and 4 L/min. Panel B shows the corresponding velocity field from the CFD test simulation. Panel C presents the difference between NIPCE and CFD results, with the mean absolute error (MAE) computed over all mesh elements in the plane.

3.3.3 Sensitivity analysis

The sensitivity analysis of the NIPCE method is structured into three parts, each focusing on a key factor influencing predictive performance: polynomial order, number of training data points, and continuity of the training variable.

Polynomial order

Building on the previously introduced pressure–flow rate curves shown in Figure 3.4, Figure 3.6 illustrates the influence of increasing the polynomial order while keeping the number of training data points constant at 30. Panels A to D show the NIPCE predictions for polynomial orders 1, 2, 4 and 8, respectively. Panel A represents a linear approximation, which can only capture simple relationships. As the polynomial order increases, the model becomes capable of representing increasingly complex relationships. However, beyond a certain polynomial degree, prediction accuracy deteriorates. Panel D demonstrates that 30 training data points are insufficient to reliably train a NIPCE model of polynomial order 8.

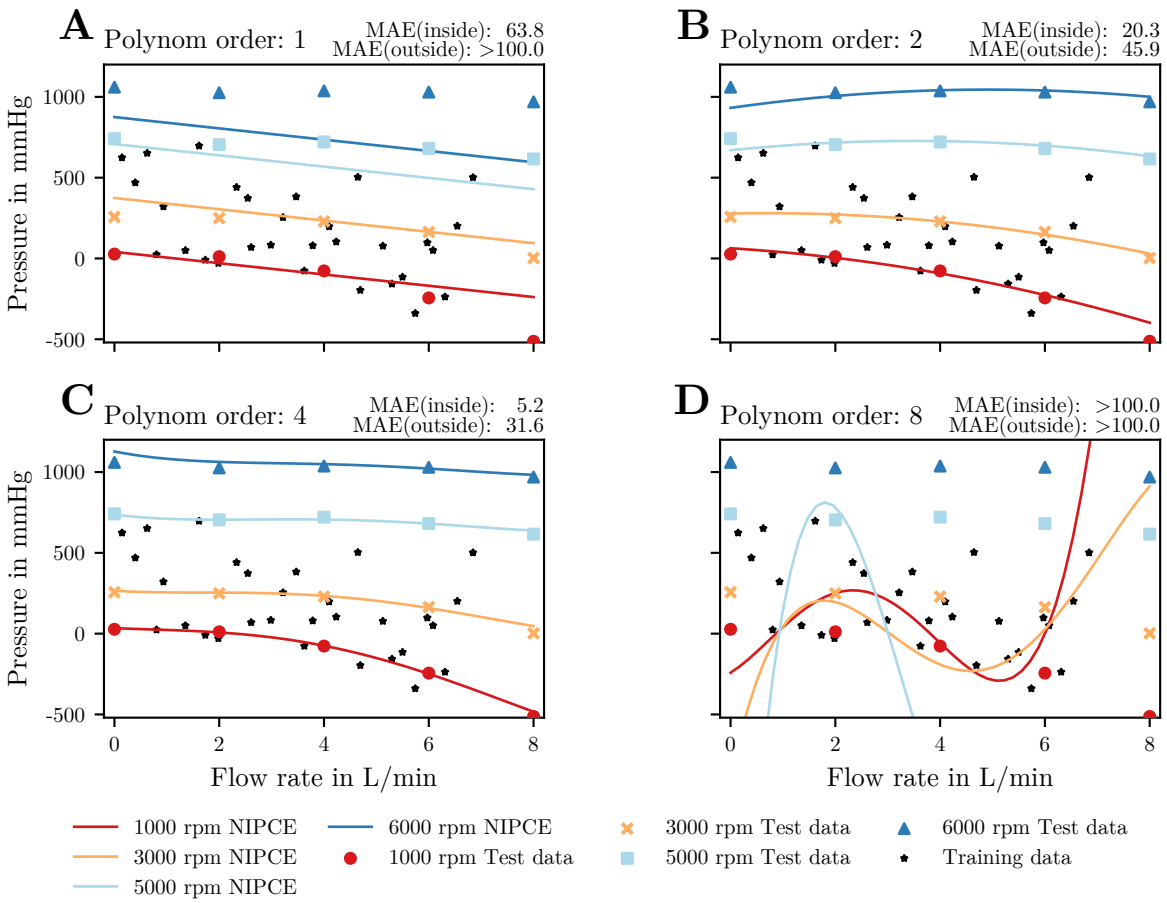

Figure 3.6: Influence of polynomial order on the predictive performance of the non-intrusive polynomial chaos expansion (NIPCE) method, using a constant number of 30 training points. Pressure–flow curves are shown over a range of 0–8 L/min and 0–6000 rpm. Panels A–D display the results for polynomial orders 1, 2, 4 and 8, respectively.

Training point number

Not only the polynomial order, but also the number of training data points has a great influence on the predictive power of the NIPCE method. Figure 3.7 is analogous to Figure 3.6, but here the number of training data points increases from 6 to 30 across Panels A to D, while the polynomial order is kept constant at 4. Panels A and B show that with only 6 or 10 training data points, the NIPCE method lacks sufficient predictive accuracy. However, with 20 training data points (Panel C), the pressure head can already be predicted reasonably well within the training range. As shown in Panel D, further increasing the number of training data points to 30 does not lead to a substantial improvement in predictive performance.

Variable continuity

In addition to the polynomial order and the number of training data points, the training variable itself also influence the predictive power of the NIPCE method. Figure 3.8 illustrates this effect using two examples. Panel A shows the distribution and values of the training

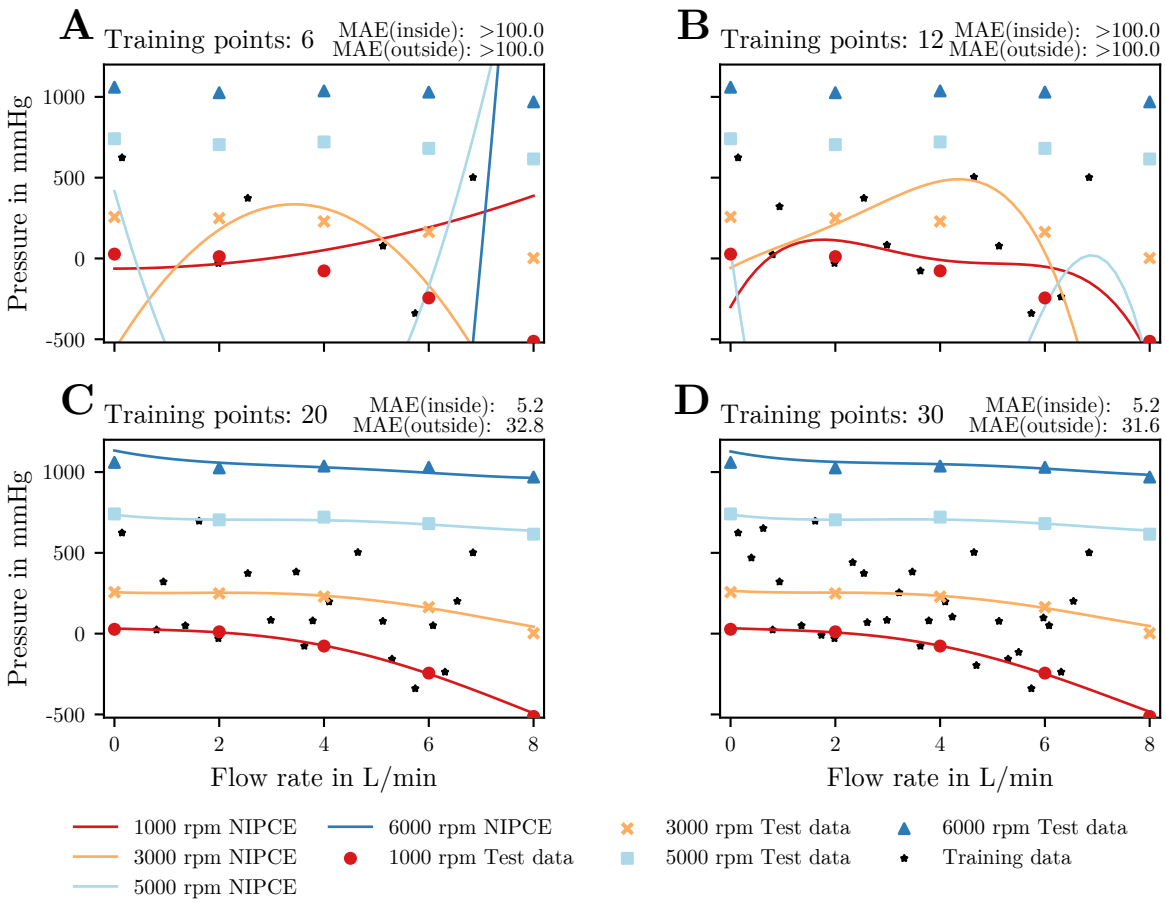

Figure 3.7: Influence of polynomial order on the predictive performance of the non-intrusive polynomial chaos expansion (NIPCE) method, using a constant polynomial order of 4. Pressure–flow curves are shown over a range of 0–8 L/min and 0–6000 rpm. Panels A–D display the results for number of training points 6, 12, 20 and 30, respectively.

points for the variable *axial force* using gray shaded stars. The training data span one order of magnitude and exhibit a continuous progression from low to high flow rates and rotational speeds. In this case, the NIPCE prediction agrees closely with the test data, resulting in a low MAE of 0.9 N across all test points. In contrast, Panel B displays the training data for *hydraulic efficiency*, which includes a noticeable outlier around 5.7 L/min. This discontinuity in the training distribution leads the NIPCE method to produce unrealistic predictions over the entire operating range, resulting in a substantial discrepancy from the test data.

3.4 Discussion

The present chapter addressed the central question of how fluid variables for a wide range of RBP operating points can be computed efficiently. For this, the NIPCE method was applied to model RBP performance across an operating point range of 0–8 L/min and 0–6000 rpm. The focus was placed not only on the feasibility of this approach but also on a systematic investigation of key modeling parameters. Namely, the polynomial order, the number of

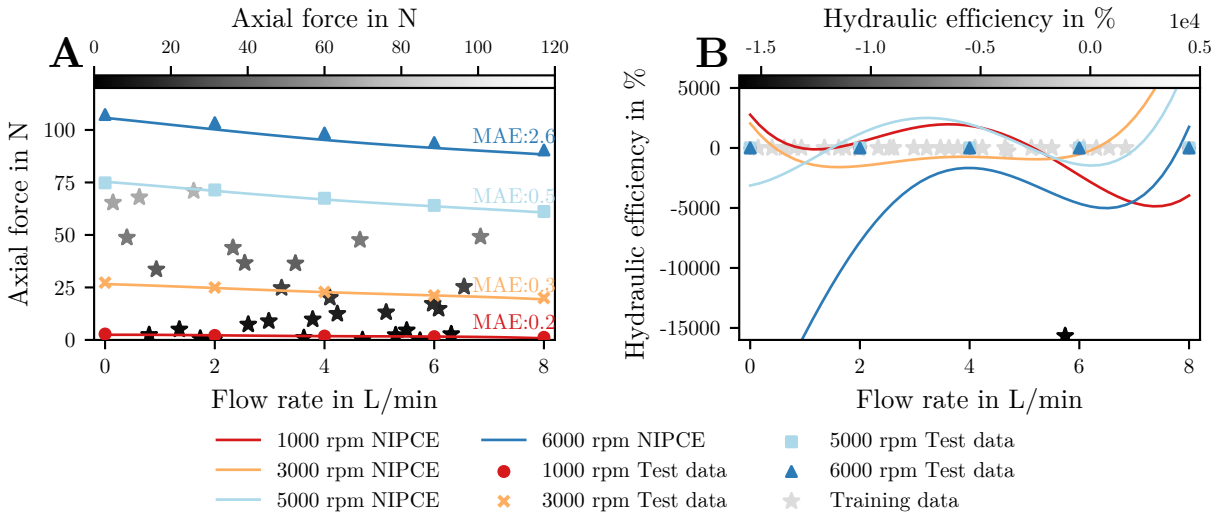

Figure 3.8: Influence of training data continuity on the predictive performance of the non-intrusive polynomial chaos expansion (NIPCE) method. Panel A shows the axial force–flow rate relationship for continuously distributed training data over a range of 0–8 L/min and 0–6000 rpm. The NIPCE predictions are compared against test data at selected constant rotational speeds lines. Training data points are shaded from gray to white, indicating the magnitude of the axial force. Mean absolute errors (MAE) for each prediction are annotated. Panel B illustrates the NIPCE prediction for the discontinuously distributed hydraulic efficiency variable (*see* equation (3.1)) along constant rotational speed lines, which are shown alongside the corresponding test data and training data points.

training data points, and the continuity of the training data. The results demonstrate that, when appropriately configured, NIPCE is capable of providing reliable predictions for quantities such as pressure head, velocity, and axial force within the training range of 0–7 L/min and 0–5000 rpm. Using a polynomial order of 4 and at least 20 well-distributed training points, accurate approximations of key performance variables were achieved. This confirms the potential of NIPCE to substantially reduce computational cost, while maintaining predictive accuracy within the desired range.

Previous applications of NIPCE in the context of turbomachinery have often lacked systematic validation or focused on narrow parameter variations [146–150]. In contrast, the present work expands upon earlier findings [130, 150] by demonstrating the method's applicability across a broad operating range relevant to ECMO therapy, rather then small variations of parameters (±10 %). The importance of rigorous model validation is highlighted through detailed comparisons with test data and through the evaluation of model sensitivity to parameter choices.

A comparison between the NIPCE method and simple interpolation of the same training data revealed that NIPCE provides more robust predictions both within and beyond the training domain. In contrast, interpolation between discretely recorded data points (*see* HQ-curves in Figure 3.4) is restricted to the convex hull of the training data and does not reach the predictive accuracy achieved by NIPCE, even within this range. The accuracy of NIPCE predictions in both interpolation and extrapolation settings is strongly dependent on the choice of polynomial order and the number of training data points. An excessively high polynomial order relative to the number of training points leads to overfitting and unstable

behavior across the entire operating range. This is due to the structure of equation (3.5), where the number of unknown coefficients a_k increases with the number of polynomial terms. Higher polynomial orders offer the potential for improved accuracy, but only if a sufficient number of training samples is available to reliably determine these coefficients through least-squares optimization. Conversely, a polynomial order that is too low may fail to capture the underlying nonlinearities of the CFD solution over the entire operating point range. This effect becomes particularly apparent when using a first-order polynomial, which inherently limits the model to linear relationships.

Another key factor influencing predictive performance is the smoothness of the training data. The NIPCE method assumes continuous functional behavior and is therefore sensitive to discontinuities or abrupt changes in the input–output relationships. As illustrated in Figure 3.8B, even a single outlier in the training data can propagate substantial errors throughout the entire prediction space. In contrast, Figure 3.8A demonstrates that a smooth and continuous variable, such as *axial force*, leads to stable and accurate predictions. This underscores the importance of ensuring continuity in the training dataset, especially when using moderate polynomial orders and a limited number of training points. To ensure reliable application of the method, both the consistency of training data and the validation of NIPCE predictions against independent test simulations must be verified. This validation step is essential to establish confidence in the model's predictive capability. This constraint is particularly relevant for QoI such as hemolysis-specific quantities or hydraulic efficiency. These variables may diverge in certain regions of the operating space or exhibit steep gradients and discontinuities, especially near low-flow conditions. For example, hemolysis metrics often involve division by flow rate and tend toward infinity as the flow approaches zero, violating the smoothness assumptions of the NIPCE method. For this reason, such variables were excluded from the present analysis, which spans the full theoretical operating range. Their suitability will be reassessed under clinically realistic flow conditions in Chapter 4, where the use of ROM techniques for modeling hemolysis-specific quantities is addressed in more detail.

Beyond scalar quantities, the applicability of NIPCE to spatially resolved data was also explored. The method was used to predict two-dimensional velocity distributions within a cross-sectional plane. While good agreement was observed for most flow regions, systematic deviations were noted in the diffuser region. This zone is particularly sensitive to minor boundary condition variations due to the presence of decelerated flow, and is therefore more challenging to model accurately. Given the broad operating range considered in this chapter and the limited number of training points, local flow features of high sensitivity remain difficult to capture. A narrower operating range combined with denser sampling and a higher polynomial order could potentially improve accuracy in such regions.

In principle, the NIPCE method can be applied to full three-dimensional datasets, including all variables obtained from CFD simulations. However, this approach substantially increases the number of degrees of freedom, requiring not only higher polynomial orders but also a much larger number of training simulations than the 30 used in this chapter. As a consequence, the associated memory requirements increase sharply. For example, applying NIPCE to the existing 30 three-dimensional CFD datasets would already require more than 180 GB of random access memory (RAM), exceeding the capacity of standard computing systems. Since 30 training points are insufficient for a meaningful three-dimensional prediction at higher polynomial orders, the actual memory and computational demands would be considerably

greater. Therefore, the application of NIPCE to full three-dimensional CFD data is currently not feasible without the development of specialized algorithms and access to high-performance computing infrastructure.
Despite the substantial reduction in computational effort compared to evaluating one million operating points, considered in Figure 3.3B, via CFD, NIPCE still requires a non-negligible number of training and test simulations. In the present case, generating 30 training simulations and 20 test simulations required approximately 15,600 core hours. This level of effort must be considered when planning applications of the method. Nevertheless, the number of training simulations can be optimized. While Hosder *et al.* [131] proposed doubling the minimum training points to ensure a well-posed regression problem of equation (3.5) as a general guideline, the present results indicate that satisfactory predictions of pressure head are achievable with fewer data points. Specifically, as few as 20 when using a fourth-order polynomial. To further reduce validation effort, cross-validation techniques could be applied, although they may limit the ability to assess extrapolation performance.

3.5 Conclusion

In conclusion, the ROM approach NIPCE was successfully applied for the first time across the entire operating range of an RBP. The influence of key modeling parameters was systematically investigated, and the following guidelines were derived to support comprehensive NIPCE modeling of RBP performance:

- Continuity of the training data should be verified prior to model construction.
- The appropriate number of training points and polynomial order must be determined through rigorous validation using independent test data.
- Inappropriate parameter settings can lead to substantial prediction errors.

These guidelines highlight the necessity of rigorous validation with test data, as incorrect NIPCE parameter settings can result in prediction errors that may remain undetected without such validation. When applied correctly, the NIPCE method offers a computationally efficient approach to estimate an RBP's performance and hemocompatibility across its full operating range. These high-resolution predictions enable the generation of full characteristic RBP footprints to support clinicians and manufacturers in evaluating and comparing individual RBP designs. This capability is particularly relevant in the context of increasing reliance on *in-silico* data for regulatory assessment, whether for uncertainty quantification, sensitivity analysis, or to inform *in-vitro* hemocompatibility testing aimed at identifying the most hemolytic operating point, in accordance with ASTM F1841 [79].
Chapter 4 builds on this foundation by applying the NIPCE-based ROM to compute shear-induced hemolysis across a wide, clinically relevant range of operating points. This enables, for the first time, a comprehensive and efficient comparison between numerical hemolysis predictions and *in-vivo* data across the full clinical spectrum.

4 Comparing in-silico hemolysis predictions and in-vivo measurements

The general findings of this chapter are based on the article "Blood trauma in veno-venous extracorporeal membrane oxygenation: low pump pressures and low circuit resistance matter" published in *Critical Care* [161].

4.1 Introduction

Mechanical blood trauma remains a central concern in the design and evaluation of blood-contacting medical devices. As outlined in Chapter 2.1.1, hemolysis arises from a multifactorial interplay of mechanical stress, RBC deformability, and biological variability. Under supra-physiological shear stress, RBCs become susceptible to membrane rupture, releasing free hemoglobin into the plasma (pfHb). This process not only reflects hemoylsis but results in broader blood damage contributing to downstream complications, including thrombosis, inflammation, and multi-organ dysfunction.

To anticipate and mitigate such risks, hemolysis models have been developed that link fluid-mechanical stress to RBC damage. As discussed in Chapter 2.3.1, these models typically rely on empirical, macroscale formulations calibrated on averaged *in-vitro* data. While such models provide a practical and computationally efficient approach, their applicability to patient-specific clinical scenarios remains limited due to biological variability and deviations from idealized flow conditions. Nevertheless, they offer valuable insights into average trends and operating conditions that increase the likelihood of hemolysis.

These kind of numerical modeling approaches are particularly relevant in the context of ECMO therapy, where blood is continuously circulated through an artificial circuit and exposed to elevated shear forces. A common modality of ECMO support is VV ECMO, which provides temporary respiratory support in patients with severe forms of acute respiratory distress syndrome (ARDS) [162–164]. As outlined in Chapter 2.1.2, VV ECMO exposes blood to non-physiological flow conditions and supraphysiological shear stresses, particularly within the RBP, which is considered the primary source of RBC damage in this context [165–167]. The clinical relevance of this damage is substantial: hemolysis has been independently associated with increased mortality in ECMO patients [168]. Moreover, blood damage contributes to major complications such as bleeding and clotting, which remain common and potentially fatal side effects of ECMO therapy [1, 2].

In an ECMO circuit, the RBP generates a pressure head, defined as the pressure difference

between the pump inlet and outlet. This pressure head represents the energy transferred from the RBP to the blood to overcome hydraulic resistance and drive blood flow through the circuit. Resistance is primarily caused by circuit components such as cannulae, tubing, and the oxygenator. The cumulative resistance, therefore, determines the amount of pressure head required to maintain a given flow rate. As resistance increases, due to patient-specific factors, component choices, or clinical conditions, the RBP must generate a higher pressure head to sustain the same flow rate. Consequently, the operating point of the RBP, defined by the combination of flow rate and pressure head, directly influences the mechanical load imposed on the blood and thereby the risk of hemolysis.

Previous *in-vitro* and *in-silico* studies have demonstrated that hemolysis strongly depends on the RBP's operating point, with lower flow rates generally leading to increased blood damage [169–171]. This is particularly important for clinical scenarios such as the weaning phase, patients with low body temperature, small body surface area, isolated CO_2 removal, or pediatric ECMO cases, as all of these require reduced blood flow rates and are therefore more susceptible to shear-induced hemolysis. However, these studies typically investigated only narrow ranges of operating conditions and often maintained constant pressure head or fixed rotational speeds. As a result, they do not represent the full spectrum of RBP use in clinical ECMO therapy, where flow rates between 1–7 L/min and pressure heads up to 400 mmHg are common [9, 10]. They also lack direct validation against clinical hemolysis measurements. Given these clinical risks and the widespread use of hemolysis models in device design, their validation has become essential to demonstrate reliability and enable regulatory acceptance. As outlined in Chapter 2.2.2, computational models are increasingly considered in regulatory processes, provided they are supported by rigorous VVUQ. In particular, hemolysis models must be evaluated against real-world clinical data to ensure their reliability for risk assessment and device optimization.

This chapter addresses this unmet need by focusing on the second research question of this dissertation:

2. *How accurately do in-silico hemolysis models capture clinical in-vivo hemolysis data?*

To address this question, a large, high-resolution clinical dataset from 580 VV ECMO patients is analyzed in combination with stress-based *in-silico* hemolysis predictions. These predictions are generated using the reduced-order modeling approach introduced in Chapter 3, allowing efficient evaluation of hemolysis across the full spectrum of clinically relevant RBP operating points. This enables, for the first time, a systematic validation of numerical hemolysis models against real-world *in-vivo* data, including low-flow and low-pressure head scenarios that have been underrepresented in previous studies. By closing this validation gap, the results presented in this chapter strengthen the reliability of *in-silico* hemolysis modeling and support the derivation of clinically meaningful conclusions. In doing so, they advance the current literature on the hemocompatibility of specific RBP operating conditions in ECMO support toward a more holistic assessment across the full range of clinically relevant operating points.

4.2 Materials & Methods

To assess the impact of VV ECMO support on hemolysis, clinical *in-vivo* patient data were used as part of a single-center, retrospective analysis. The dataset included RBP operation metrics and routinely recorded blood parameters obtained from high-resolution electronic health records. *In-vivo* hemolysis was assessed solely through pfHb levels. However, to provide

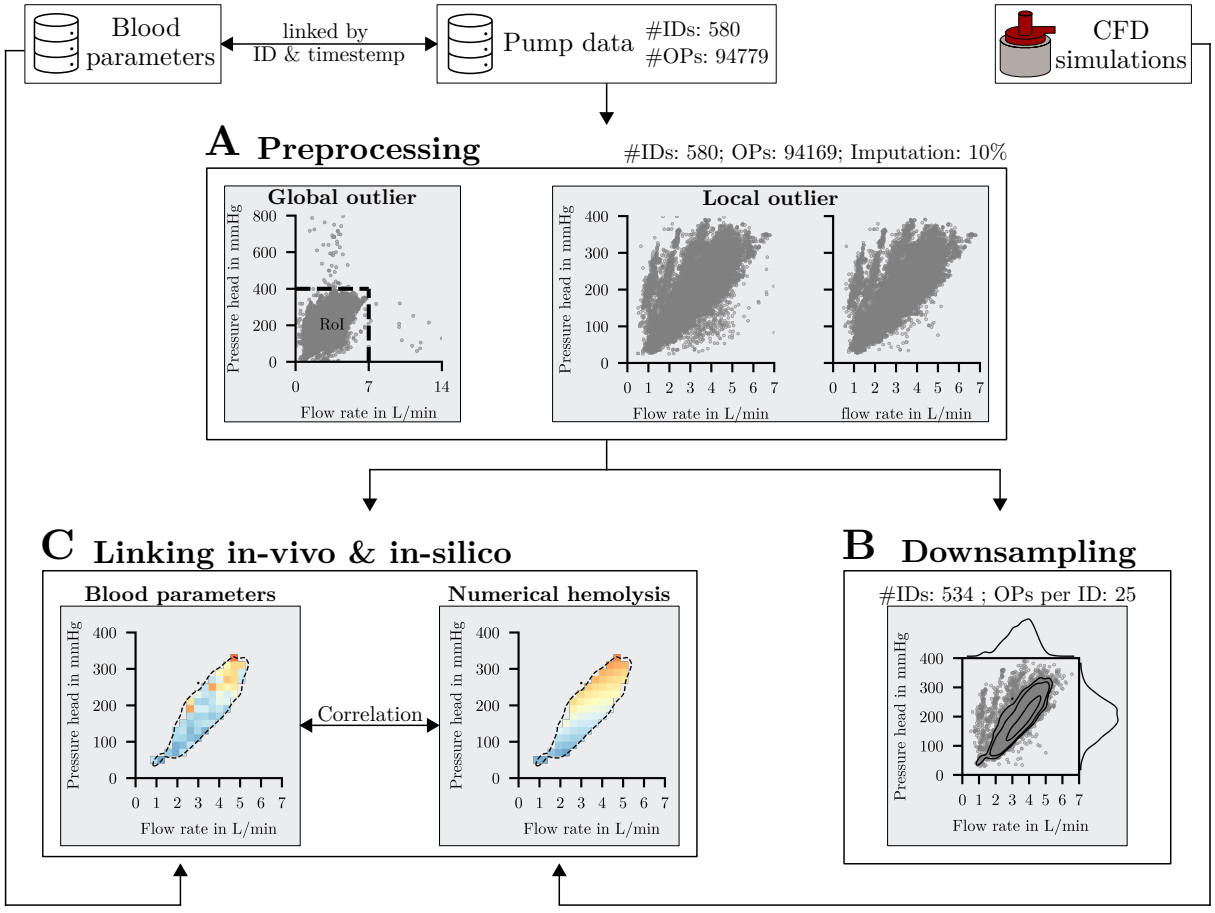

Figure 4.1: Graphical overview of the used *in-vivo* blood parameters, rotational blood pump (RBP) data, and *in-silico* computational fluid dynamics (CFD) databases and clinical data handling, detailing the steps of preprocessing, downsampling, and linking *in-vivo* and *in-silico* data, as illustrated in Panels A, B, and C, respectively. The number of unique patients is abbreviated as "#ID" and the RBP operating points as "OP".

a more comprehensive analysis that reflects the diverse scenarios encountered in the intensive care unit (ICU), additional blood parameters were included: lactate dehydrogenase (LDH), bilirubin, oxyhemoglobin (oxyHb), deoxyhemoglobin (deoxyHb), and D-dimer. While all parameters except D-dimer showed an association with hemolysis, D-dimer was included as a negative control to confirm that pump thrombosis did not significantly influence the results. *In-silico* hemolysis at clinically observed operating points was determined using the reduced-order modeling approach based on CFD calculations of RBPs, as introduced in Chapter 3. This enabled a comprehensive evaluation of clinically relevant RBP operating scenarios. The *in-vivo* pressure head, defined as the pressure difference before and after the RBP, was

approximated from the clinical dataset using the difference between drainage cannula pressure (P_{ven}) and post-pump pressure (P_{int}) (*see* Figure 4.7).
An overview of the data processing steps applied in this chapter is provided in Figure 4.1. It illustrates the integration of clinical and computational data, as well as the procedures for preprocessing, downsampling, and linking operating points. These steps are described in more detail in the following sections and in Appendix A.

4.2.1 Clinical in-vivo data

The clinical dataset was collected between 2012 and 2022 at the ECMO Center Cologne-Merheim and comprised 813 patients treated with ECMO support. A total of 233 patients were excluded based on the following criteria: age under 18 years ($n = 24$), absence of VV ECMO treatment ($n = 21$), ICU length of stay shorter than one day ($n = 29$), and insufficient RBP metrics data ($n = 159$). This resulted in a final cohort of 580 patients, the majority of whom were diagnosed with ARDS. Figure 4.2 illustrates both the inclusion and exclusion

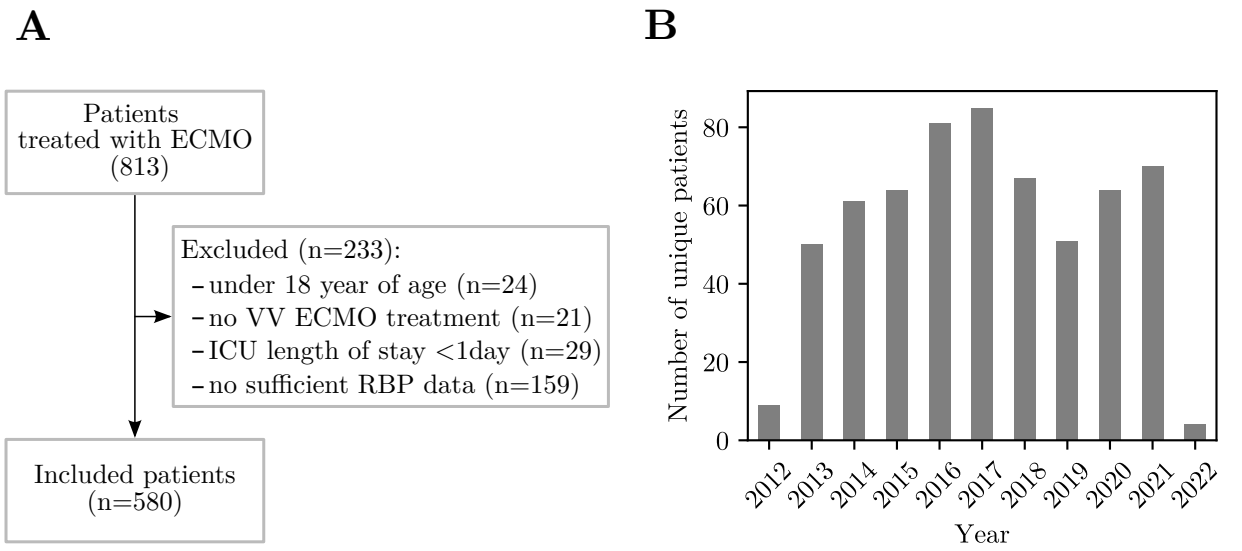

Figure 4.2: Shows in A the inclusion and exclusion criteria of the study's cohort. Of 813 patients treated with ECMO, a total of 233 patients were excluded based on four exclusion criteria, resulting in a final cohort of 580 patients. Panel B shows the distribution of patients included in the cohort across a 10-year period, with approximately 60 patients per year.

criteria (panel A) as well as the distribution of included patients over the 10-year data collection period (panel B), during which approximately 60 patients per year were included. To assess the potential impact of the Coronavirus Disease 2019 (COVID-19) pandemic on the data, a sensitivity analysis was performed by excluding all patients admitted to the ICU after March 2020. These patients represented approximately 20 % of the total cohort, though not all were COVID-19 positive. The analysis indicated that the exclusion of this subgroup did not noticeably affect the observed relationship between *in-vivo* pfHb levels and *in-silico* blood damage.
The dataset consisted of 94,779 individual RBP operating points, corresponding to a cumula-

tive ECMO support duration of 11,314 days. On average, each patient contributed a median of 125 operating points over a median ECMO support duration of 18 days. These values reflect the clinical and technical depth of the dataset, supporting detailed investigation of hemolysis under varying operating conditions. For this, two primary databases were used: one containing routine blood parameter measurements and the other recording RBP operating points. The datasets were linked using patient ID and timestamp, which enabled an integrated analysis of clinical and technical variables. The details of this process are provided in subchapter 4.2.3. An overview of key cohort characteristics, including demographics, clinical outcomes, baseline values during the first 48 hours of ECMO support, and operating point variables, is provided in Appendix A.1. The data are stratified by ICU survival, and a graphical summary of key parameter distributions is also included.

To ensure the robustness of the operating point data, a dedicated preprocessing step was applied prior to analysis (Figure 4.1, Panel A). Initial data cleaning involved the removal of global outliers by defining a Region of Interest (RoI) between 0–7 L/min for flow rate and 0–400 mmHg for pressure head. This RoI retained 99.4 % of all operating points and excluded extreme values, likely caused by unit mislabeling or decimal errors. Within the defined RoI, local outliers were identified and corrected using the Hampel filter [172], which involved sorting the dataset chronologically and applying a sliding window of size 5. Data points deviating more than three standard deviations from the median within this window were classified as outliers and replaced. This resulted in an imputation rate of approximately 10 %. Following this procedure, the total number of operating points was reduced from 94,779 to 94,169, with all 580 patients retained for further analysis.

4.2.2 In-silico methods

The *in-silico* hemolysis assessment employed CFD simulations of the Rotaflow (Getinge, Gothenburg, Sweden) and DP3 (Xenios, Heilbronn, Germany) centrifugal pumps, based on a previously validated numerical setup developed by Gross-Hardt *et al.* [152]. To improve computational efficiency, stationary simulations were used, which have been shown to yield hydrodynamic results comparable to those of transient simulations of the original setup. Hemolysis prediction was carried out using the Eulerian Garon and Farinas approach [114], incorporating hemolysis model parameters defined by Heuser *et al.* [108]. This approach results in a numerical hemolysis measure analogous to the Modified Index of Hemolysis (MIH), as defined by the ASTM F1841-97(2017) standard [79].

While CFD simulations provide detailed predictions of flow and hemolysis, they are computationally expensive. Exhaustively simulating hemolysis across the entire operating point range, covering nearly 95,000 cases, would be infeasible using CFD alone. To overcome this limitation, the ROM strategy based on NIPCE, introduced in Chapter 3, was employed. In contrast to Chapter 3, the training data in this chapter was limited to clinically relevant operating points above 0.5 L/min, as suggested by the *in-vivo* data of this chapter. This approach enables rapid evaluation of pressure heads, flow rates, and hemolysis values while maintaining high predictive accuracy. Following the guidelines established in Chapter 3, the ROM was trained on 30 CFD simulations across the RoI and validated against 9 independent test cases. The results of this validation are shown in Figure 4.3, which demonstrates the ROM's accuracy in predicting hydrodynamic variables and hemolysis across both RBP types.

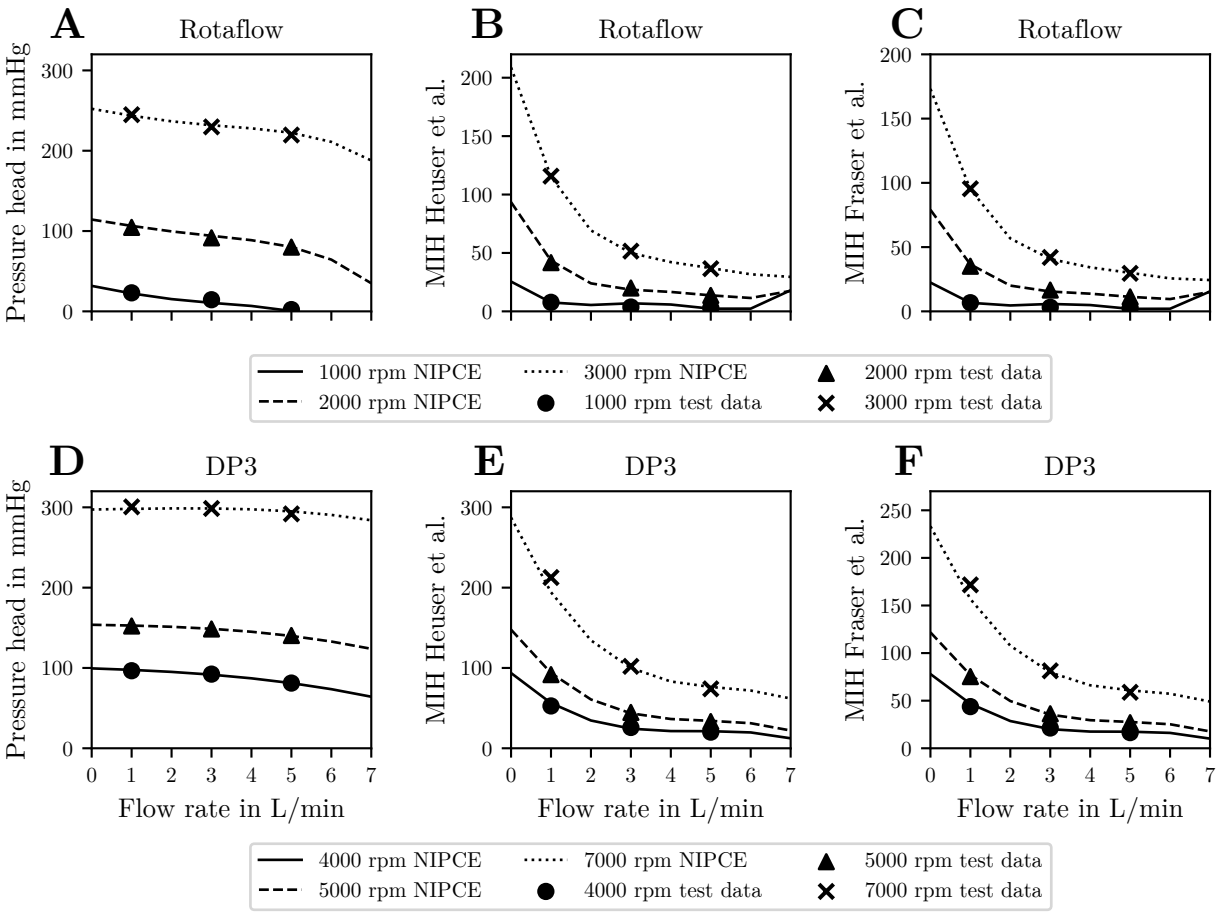

Figure 4.3: Validation overview of reduced order method non-intrusive polynomial chaos expansion (NIPCE). All lines are predictions of the NIPCE method, while cross, triangle and circle symbols are validation points, that were not included in the training data. Panel A shows the validation of pressure and flow rate predictions while B and C show the validation of hemolysis predictions using Heuser [108] and Fraser [173] model coefficients for the Rotaflow (Getinge, Gothenburg, Sweden) model. Panels D–F depict the same for the DP3 (Xenios, Heilbronn, Germany) model.

Panel A compares the predicted flow rates and pressure heads (lines) with the CFD validation points (dots,triangles and crosses) for the Rotaflow setup, showing strong agreement between model output and reference data. Panels B and C validate hemolysis predictions using the Heuser [108] and Fraser [173] hemolysis model parameters, respectively. For the DP3 setup, validation results are shown in Panels D–F, again demonstrating minimal deviation between the NIPCE model and the corresponding CFD simulations. These results confirm that the ROM accurately captures both hydraulic and hemolytic behavior across all relevant operating range.

Based on the validated ROM, hemolysis predictions in terms of MIH could be efficiently computed for all 94,779 clinical operating points without the need for additional CFD simulations. A detailed description of the reduced-order modeling technique and its mathematical formulation is provided in previous Chapter 3 and the corresponding publication [145].

4.2.3 Linking in-vivo and in-silico data

As indicated in Figure 4.1B, the Raposo downsampling algorithm [174] was applied to ensure equal patient representation in the operating point probability distribution. Patients contributing fewer than 25 operating points were excluded, and those with more were uniformly downsampled to this threshold. Although this procedure reduced the dataset from 580 to 534 patients, it enabled a more balanced and representative assessment of the overall operating point probability distribution. In this context, it is important to note that this downscaling was applied solely for calculating the operating point distribution that equally represents all patients across the cohort. All other analyses in this chapter were performed using the full set of 94,779 operating points. Figure 4.4 presents a sensitivity analysis of this downsampling

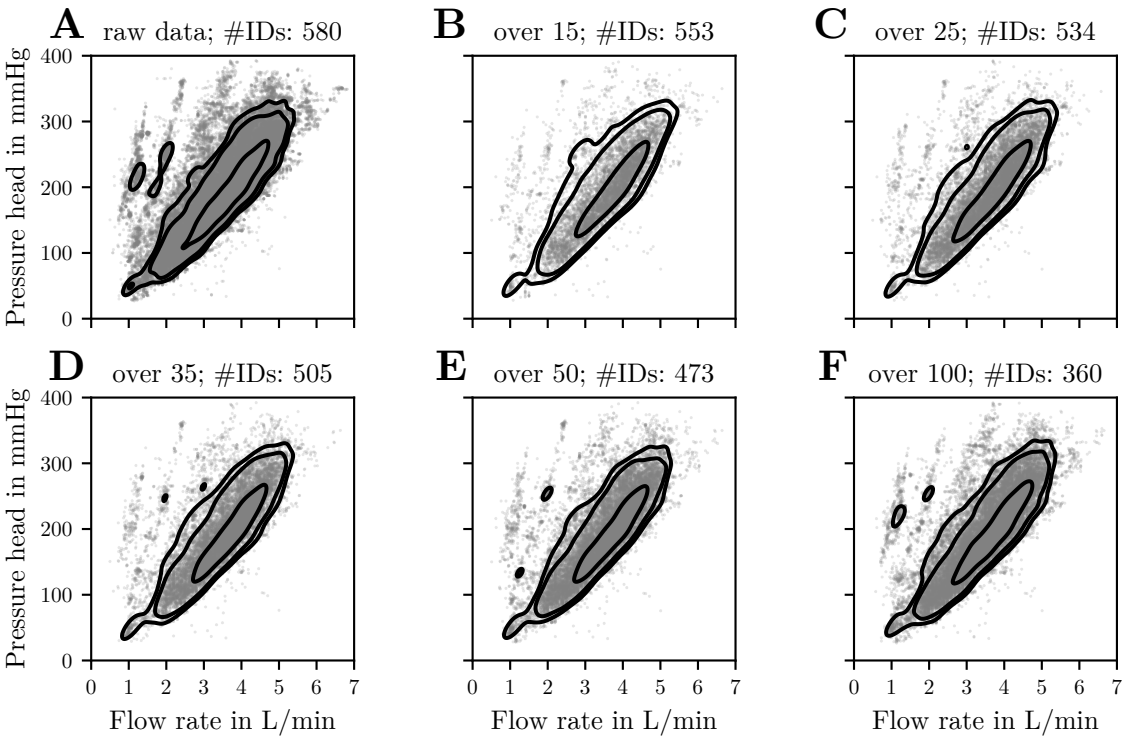

Figure 4.4: Sensitivity of Raposo [174] downsampling algorithm. Panel A depicts all 94,779 operating points of the whole cohort (580 patient IDs) with grey scatter points and the 95, 90 and 50 % confidence intervals with black contour lines. B–F shows the downsampling to 15, 25, 35, 50 and 100 operating points per patient, respectively.

approach. Panel A displays all 94,779 operating points from the full cohort, with grey scatter points and 95, 90 and 50 % confidence interval contours in black. Panels B–F show the results of applying varying downsampling thresholds (15, 25, 35, 50, and 100 operating points per patient), demonstrating that the overall shape and distribution of the contours remain largely consistent across different thresholds.

In addition to operating point data, a separate database stored routine blood test results, which were linked to the operating point data via patient identification (ID) and timestamp (*see* Figure 4.1C). To achieve temporal alignment of these two data bases, operating points were averaged within an 8-hour interval and assigned to the nearest corresponding blood

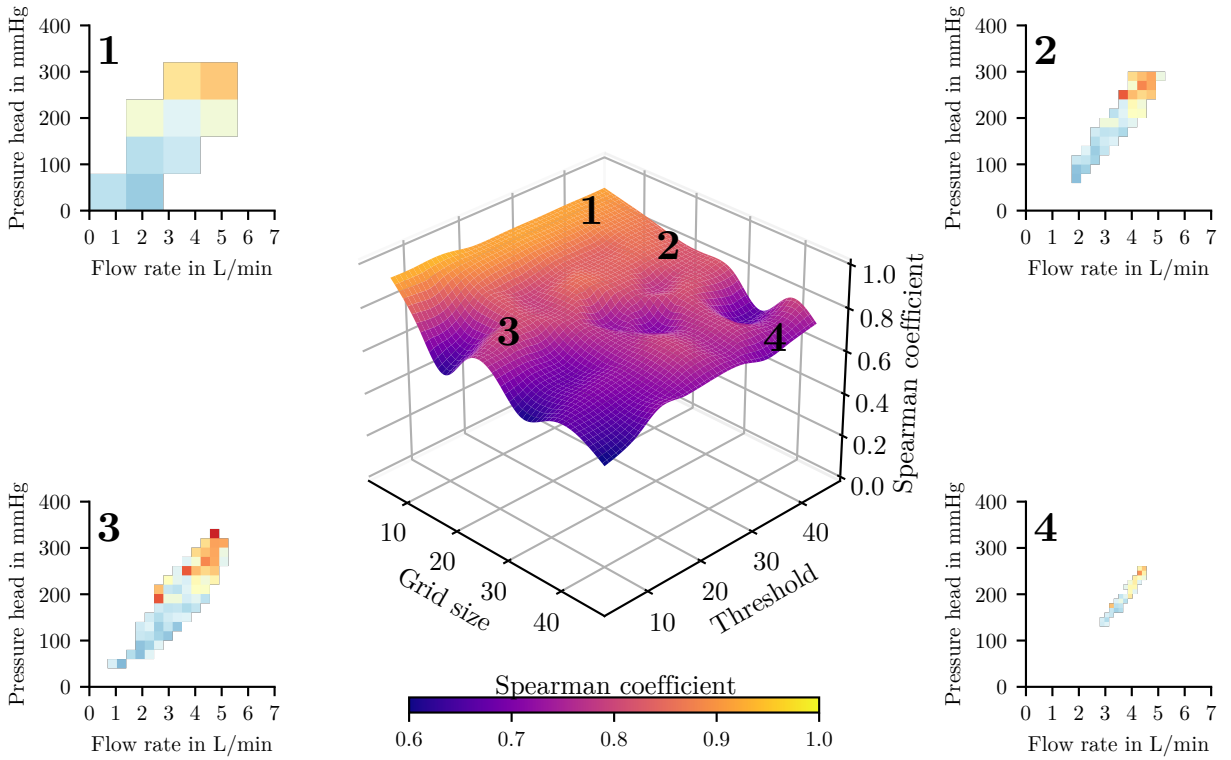

Figure 4.5: Sensitivity of the Spearman correlation coefficient to variations in grid size (5–45) and operating point threshold values (5–45). Plasma free hemoglobin (pfHb) images corresponding to specific grid and threshold combinations are shown in zoom windows: (1) 5×5, threshold 40; (2) 20×20, threshold 40; (3) 20×20, threshold 10; and (4) 40×40, threshold 40. Each configuration illustrates trade-offs between correlation strength, resolution, and data coverage.

parameter measurement of a given patient. A sensitivity analysis was performed to justify the selection of the 8-hour window. As shown in Figure A.7, 95 % of all flow rate and pressure head values varied by less than 0.5 L/min and 50 mmHg, respectively, over an 8-hour timeframe. This corresponds to approximately 7–10 % of the total RoI in the operating point space, and was deemed an acceptable trade-off between data availability and temporal resolution. To support this, measurement frequencies of the relevant blood parameters and RBP data are analyzed in Figure A.6, showing that pfHb, bilirubin, and LDH were typically sampled every 24 hours, while oxyHb and deoxyHb were measured every 2 hours. RBP operating points were recorded approximately every 2 hours, ensuring that the resolution of RBP data matched or exceeded that of blood parameter measurements.

To statistically compare *in-silico* hemolysis predictions to *in-vivo* pfHb measurements, the RoI was divided into a 20×20 grid (Figure 4.1C), with each cell covering 5 % of the total operating point space. Within each cell, the median of all assigned values was calculated to reduce the influence of short-term fluctuations and individual outliers. To ensure statistical robustness of these median values, a minimum threshold of 10 operating points per cell was imposed. Assuming predominantly unidirectional shifts in RBP settings, the observed operating point variability of approximately 7–10 % over the selected time window of 8 hours effectively translates to about 3.5–5 % in one direction. This aligns with the chosen 5 % grid resolution, ensuring that temporal variations largely remain within a single grid cell.

As a result, the setup enables a structured and interpretable comparison between *in-silico* predictions and heterogeneous *in-vivo* data while minimizing misalignment errors.
The sensitivity of this correlation method to grid size and operating point thresholds for each grid cell was further analyzed in Figure 4.5. Variations in grid resolution (5 to 45) and operating point threshold values (5 to 45) produced only minor changes in the resulting Spearman correlation coefficient between *in-silico* predictions and *in-vivo* measurements of pfHb, indicating robustness of the correlation results to these parameters. Four zoom windows highlight representative trade-offs:

- **Zoom 1** (5×5 grid, threshold 40): Displays the coarsest resolution. A high correlation coefficient (0.9) is observed, but the low spatial resolution limits interpretability of pressure–flow relationships and obscures finer data structure.
- **Zoom 2** (20×20, threshold 40): Reduces the grid size while maintaining a high threshold. This improves resolution but narrows spatial coverage, as fewer grid cells meet the inclusion criteria.
- **Zoom 3** (20×20, threshold 10): The selected configuration used in this analysis. It achieves a practical balance between resolution, data inclusion, and statistical robustness, enabling consistent correlation results across the full operating point range.
- **Zoom 4** (40×40, threshold 40): Offers the finest resolution but restricts the analysis to regions with the highest data density, potentially overlooking less frequent yet clinically relevant operating points.

4.3 Results

This chapter presents the results of linking clinical blood parameters to RBP operating points, validating numerical hemolysis predictions, and examining variability across patients to assess the consistency and relevance of *in-silico* models in the context of VV ECMO therapy.

4.3.1 Linking blood parameters and RBP operating points

The distribution of blood parameters across different combinations of pressure head and flow rate can be visualized by connecting them with the RBP operating point data (*see* Figure 4.6). Specifically, pfHb, bilirubin, and LDH (panels A–C) demonstrate markedly elevated median values at the maximum pressures across the entire range of flow rates (1–7 L/min), particularly in regions of overall high pressures (> 200 mmHg). In contrast, Figure 4.6D–E depict opposing trends for oxyHb and deoxyHB, with the latter showing elevated levels in high-pressure zones and the former exhibiting reduced levels in these same areas. The D-dimer parameter, shown in Figure 4.6F, presents a more subtle trend, with occasional spikes in values at the highest pressures of certain flow rates. The number of matched blood parameter measurements differs between markers. Approximately 5000 pfHb values were available, compared to about 10,000 for bilirubin and LDH, 20,000 for D-dimer, and around 70,000 for oxyHb and deoxyHb. This variation reflects the differences in clinical measurement frequency across parameters.

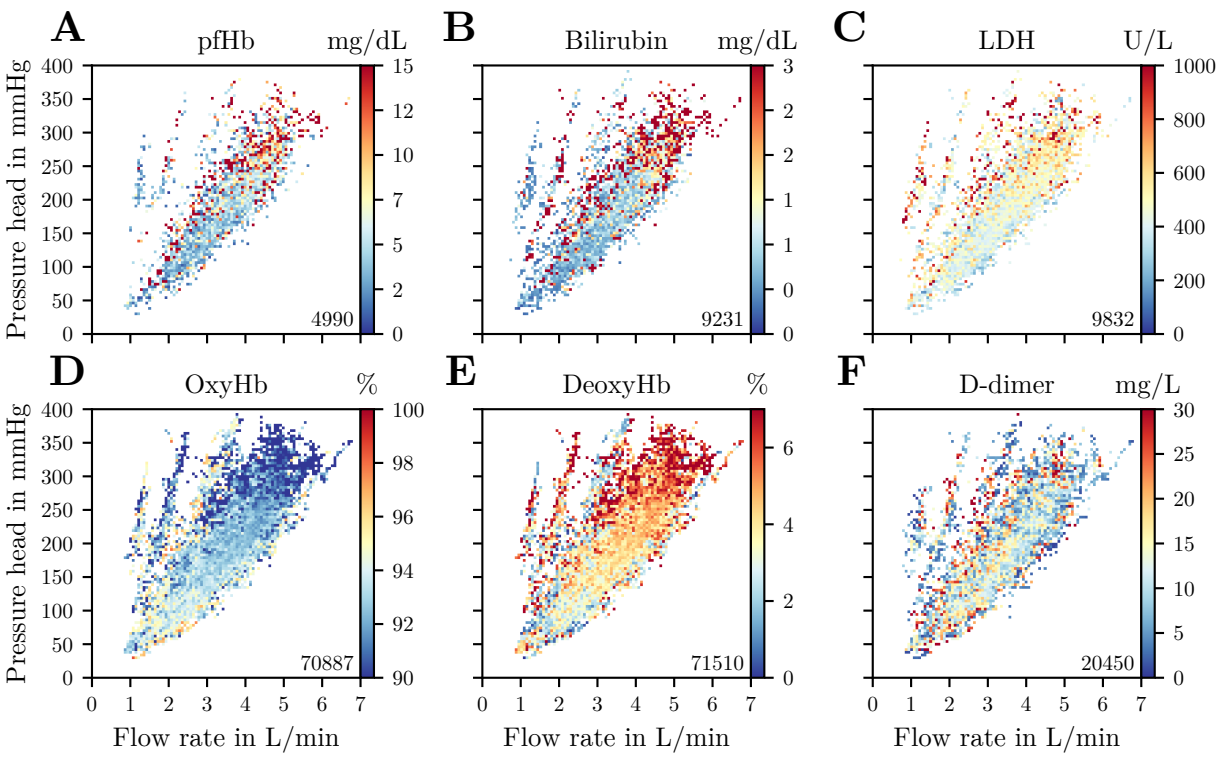

Figure 4.6: Illustrates the relationship between rotational blood pump (RBP) parameters and *in-vivo* blood parameters for plasma free hemoglobin (pfHb), bilirubin, lactate dehydrogenase (LDH), oxyhemoglobin (oxyHb), deoxyhemoglobin(deoxyHb), and D-dimer, shown on panels A–F, respectively. To minimize overlay effects from individual measurements, the median values from small grid cells (90×90) are presented. The lower right corner of each panel displays the count of individual blood parameter measurements that have been matched with RBP operating points. The streaks in the low flow and high-pressure area represent individual data points within the cohort that cannot be clearly associated with a specific ECMO support situation. As shown in Figure 4.7C, these streaks correspond to a very small number of unique patients. Therefore, could also be data artefacts that were not removed in the data preprocessing.

4.3.2 Relevant operating points and pressure measurements

To place these raw data in the context of clinically relevant operating points, Figure 4.7A illustrates the marginal distributions for flow rate and pressure head, along with the down-sampled operating points of the whole cohort. The black contour lines represent the 95, 90 and 50 % confidence intervals of the joint probability distribution of patient operating points. The operating points (grey scatter) are predominantly clustered within an elliptical zone spanning from the lower left (1 L/min; 400 mmHg) to the upper right of the plot (5.3 L/min; 340 mmHg). The red cross at a flow rate of approximately 4.1 L/min and pressure head of 225 mmHg indicates the operating point with the highest likelihood across the entire cohort. The confidence intervals of the probability distributions are transferred to Figure 4.7B, which showcases the numerically predicted hemolysis within the RoI ranging from 0–7 L/min flow rate and 0–400 mmHg pressure head. Notably, the low-flow and high-pressure region in the upper left corner shows a much higher potential for hemolysis compared to all other operating points. As previously discussed in Chapter 3, this region represents an unrealistic operating condition within the extrapolation zone of the NIPCE method, where predictions may yield

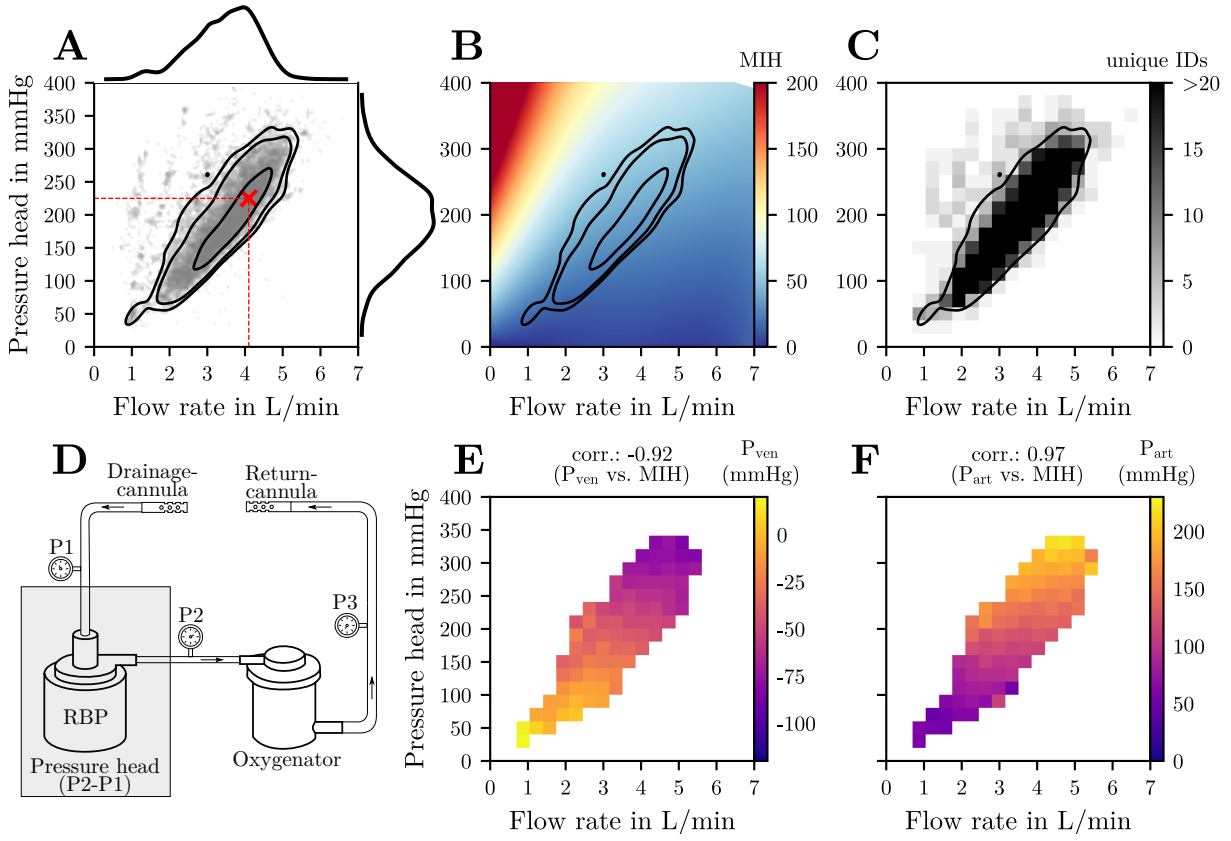

Figure 4.7: Panel A shows the *in-vivo* operating point distribution for 534 equally weighted patients (25 operating points each). Grey dots indicate individual points, with black lines denoting the 95, 90 and 50 % confidence intervals. Marginal distributions for flow rate and pressure head are shown along the top and right axes, and the most likely operating point is marked with a red cross. Panel B presents the *in-silico* hemolysis prediction across the full operating range. Panel C shows the number of unique patient IDs per cell within the 20×20 correlation grid used for pfHb analysis. Panel D schematically illustrates the VV ECMO circuit, indicating pressure measurements at the drainage cannula (P_{ven}), post-pump segment (P_{int}), and return cannula (P_{art}). Panels E and F display the *in-vivo* measurements for P_{ven} and P_{art}, respectively, with Spearman correlations to predicted hemolysis noted above.

unphysiologically high hemolysis values due to the lack of training data. To further focus the analysis on clinically relevant operating conditions, the number of unique patients per grid cell is shown in Figure 4.7C. Regions outside the 95 % confidence interval are influenced by less than five unique patients per grid cell, while grid cells inside the 95 % confidence interval have values over 20 unique patients per grid cell. To explore the components of the variable pump pressure head, Figure 4.7D presents a schematic overview of all VV ECMO circuit components influencing its resistance levels. It is shown that the *in-vivo* pressure head is composed of both the drainage cannula pressure (P_{ven}) and the post-pump pressure (P_{int}). Figure 4.7E–F display the measured distributions of drainage cannula pressure P_{ven} (E) and return cannula pressure P_{art} (F) based on the clinical cohort data. A correlation analysis with the numerical predicted hemolysis (*see* Figure 4.7B) inside the region of relevant pressure measurements reveals that both the negative drainage cannula pressure and the return cannula pressure are highly associated with numerical hemolysis, indicating that both

pressures are relevant for hemolysis. This is quantified by very strong correlation coefficients of -0.92 and 0.97, respectively (*see* Figure 4.7E–F).

4.3.3 Linking in-vivo and in-silico data

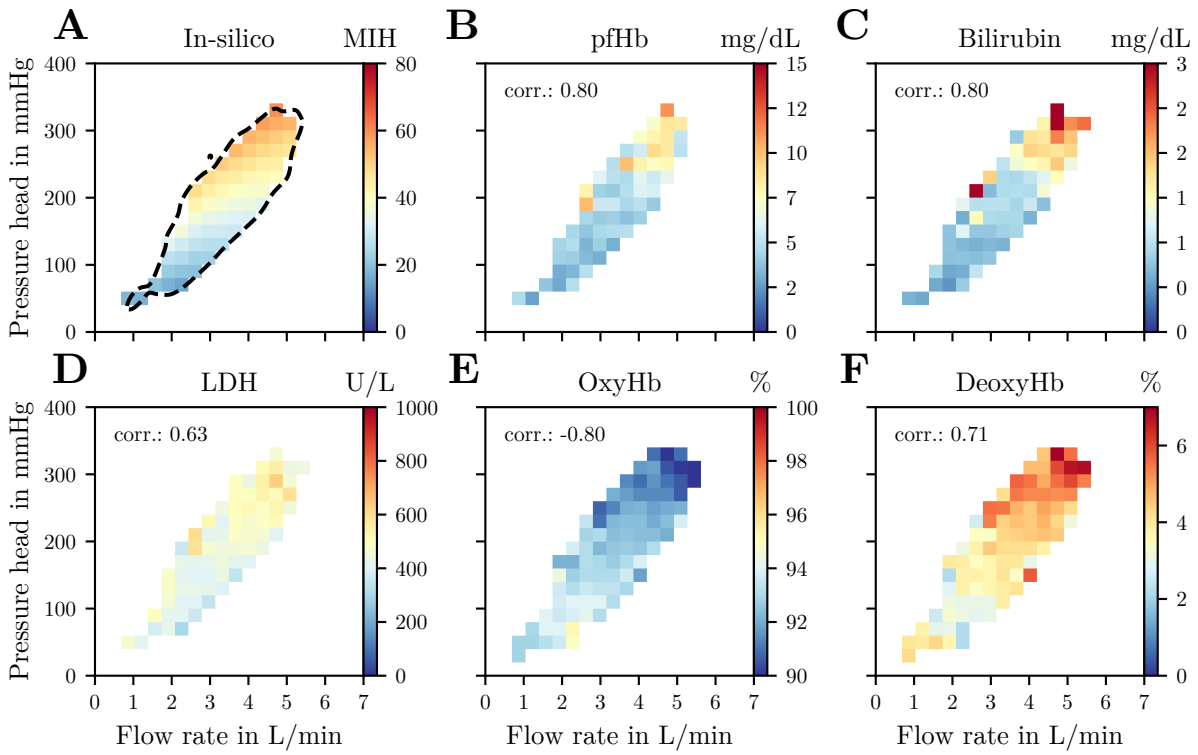

Figure 4.8: In A, the *in-silico* numerical prediction of hemolysis is displayed in a 20×20 grid. B–F show the distribution of *in-vivo* parameters plasma free hemoglobin (pfHb), bilirubin, lactate dehydrogenase (LDH), oxyhemoglobin (oxyHb), deoxyhemoglobin (deoxyHb), and D-dimer on the same grid. The abbreviation "corr." describes the Spearman correlation coefficient of the numerical prediction and the specific blood parameter. For all grid points the median value from at least 10 data points is displayed.

Narrowing down the numerical prediction of hemolysis from Figure 4.7B to clinically relevant operating conditions within the 95 % confidence interval of the cohort's operating point distribution, Figure 4.8A displays the predicted hemolysis potential in regions supported by at least 10 recorded operating points. This targeted selection ensures that the analysis is grounded in well-represented clinical data, enabling a meaningful and accurate comparison with observed blood parameter values. These are displayed in Figure 4.8B–F by illustrating the spatial distribution of pfHb, bilirubin, LDH, oxyHb and deoxyHb across the same grid, with an indicated Spearman correlation coefficient to the numerical prediction of hemolysis of 0.80, 0.80, 0.63, −0.80 and 0.71, respectively. The not shown D-dimer parameter has a correlation coefficient of -0.13.
Figure 4.9 combines the observations from the analysis of the *in-vivo* data and the *in-silico* prediction of hemolysis in the RBPs Rotaflow and DP3 in the Panels A and B, respectively. It classifies the determined blood trauma into four relative risk regions: low, moderate, high,

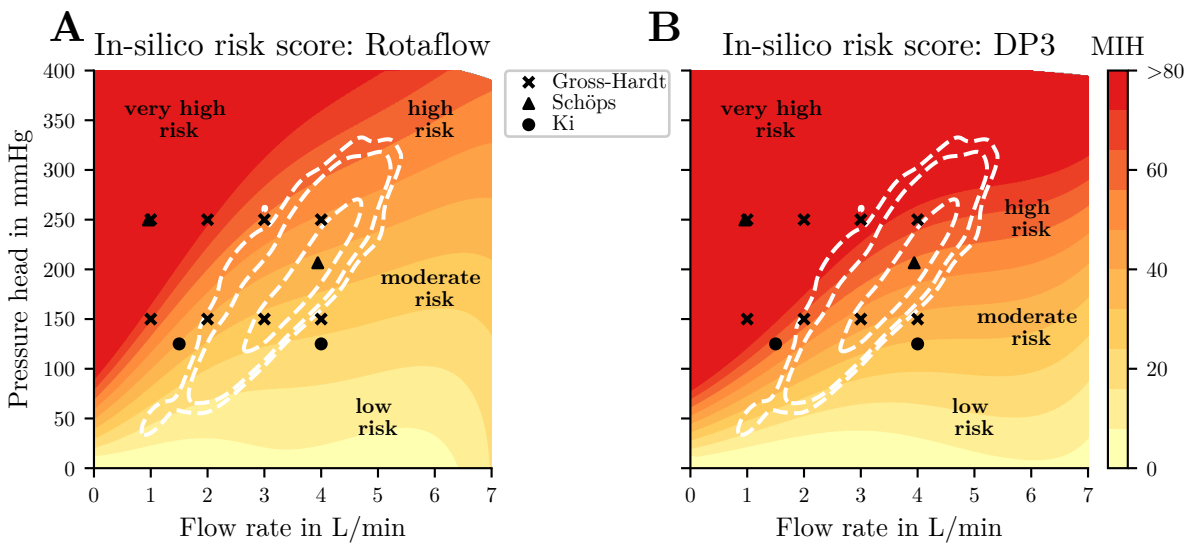

Figure 4.9: Provides in A and B an overview of the *in-silico* numerical hemolysis potential across the entire range of operating points for the Rotaflow (Getinge, Gothenburg, Sweden) and DP3 (Xenios, Heilbronn, Germany) pumps, respectively. The colormap is divided into four sections: low, moderate, high, and very high risk, to better highlight differences in the clinically relevant area. The 95, 90 and 50 % confidence intervals of the operating point probability distribution are marked in dashed white lines. Both plots share the same colorbar showing numerical hemolysis calculated with the Heuser *et al.* [108] hemolysis parameter set displaying values of MIH from 0 to > 80. Furthermore, the examined operating points from the studies by Gross-Hardt *et al.*, Schöps *et al.* and Ki *et al.* [169–171] are marked with crosses, triangles, and circles respectively.

and very high. These classifications were selected to simplify the presentation of the data with arbitrarily chosen thresholds for each class and should not be interpreted as absolute risk scores. By incorporating the probability distribution of patient operating points, Figure 4.9 highlights this information using dashed white contour lines, indicating the area in which the numerical hemolysis model previously showed a strong association with *in-vivo* hemolysis (*see* Figure 4.8). The color gradient clearly indicates that at low flow rates, a small change in pressure head results in larger variations in hemolysis than the same change in pressure head would cause at higher flow rates. Additionally, areas of high flow rate and consequently high pressure head within the 95 % confidence interval of all operating points exhibit an increased potential for hemolysis compared to areas with low flow rate and low pressure head, such as those typically encountered during the weaning process. These observations hold for both the Rotaflow and DP3 pump, despite the latter generally exhibiting a higher hemolysis signal. Areas previously studied in low-flow ECMO pump analyses are highlighted in black with crosses, triangles, and circles. It becomes evident that the tested conditions of Gross-Hardt *et al.*, Schöps *et al.* and Ki *et al.* [169–171] are sparsely distributed and cannot represent the clinically relevant region of VV ECMO therapy. Especially the investigated low flow conditions lie outside of clinically observed operating points (outside of the 95 % confidence bounds).

While Figure 4.9 features only the numerical hemolysis model of Heuser *et al.* [108] for

Table 4.1: Displays the Spearman correlation coefficient for plasma free hemoglobin (pfHb) and modified index of hemolysis (MIH) using a 20×20 grid and a threshold of 10 data points across (Getinge, Gothenburg, Sweden) and DP3 (Xenios, Heilbronn, Germany) setups, including all employed hemolysis models (Giersiepen [99], Fraser [173], Heuser [108], and Zhang [106]).

RBP	**Giersiepen** [99]	**Fraser** [173]	**Heuser** [108]	**Zhang** [106]
Rotaflow	0.80	0.79	0.80	0.79
DP3	0.80	0.78	0.79	0.80

the Rotaflow and DP3 pump setups, Table 4.1 shows Spearman correlation coefficient for pfHb and MIH across all investigated combinations of RBP setups and hemolysis model coefficients. The resultant numerical hemolysis predictions differ in absolute hemolysis values but the overall hemolysis trend in regard to pressure and flow rate changes, as indicated by the spearman correlation coefficient, stay the same across all RBP and hemolysis models. While the strong correlation between predicted hemolysis and *in-vivo* pfHb observed at the cohort level reflects the overall trend, individual-level analysis reveals a more heterogeneous picture. Figure 4.10 presents three representative examples of patient-level correlations. Panel A shows a patient with high agreement between numerical predictions and measured

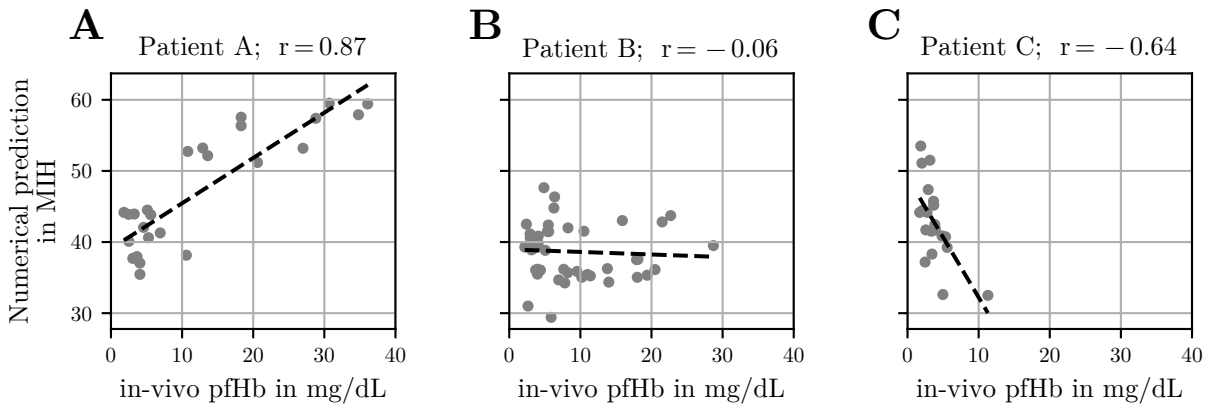

Figure 4.10: Examples of patient-level correlations between *in-silico* hemolysis predictions and *in-vivo* plasma free hemoglobin (pfHb) measurements. Panel A shows a representative patient with strong positive correlation ($r = 0.87$), while Panel B displays no correlation ($r = -0.06$), and Panel C an inverse correlation ($r = -0.64$). These examples highlight the variability in individual patient responses despite the overall strong cohort-level trend.

pfHb values, consistent with the majority of the cohort. Panel B illustrates a patient with no clear correlation, and Panel C shows an inverse trend, demonstrating that, despite a robust overall association, patient-level variability (inter-variability) exists within the dataset, with some individuals deviating from the general pattern. The spread of data points within each panel further indicates within-subject variability (intra-variability), suggesting fluctuations in pfHb levels even under similar shear conditions.

4.4 Discussion

With the aim of addressing the second research question of the dissertation, this chapter examines how well stress based *in-silico* hemolysis predictions align with clinical observations across a wide range of VV ECMO operating conditions. The combined use of high-resolution *in-vivo* clinical data from 580 patients and *in-silico* reduced-order hemolysis modeling enabled a comprehensive comparison across a broad range of operating conditions, including scenarios that had previously been underrepresented, such as low flow rate and low pressure head. Through this validation-focused approach, several clinically relevant findings were derived, highlighting not only the reliability of the model but also its potential to inform evidence-based decision-making in VV ECMO. Three key observations were identified:

- ECMO circuit resistance emerged as a key contributor to hemolysis by forcing the RBP to generate higher pressure heads.
- Low flow scenarios were found to be particularly critical, as even minor increases in circuit resistance led to increased hemolysis.
- A strong association between *in-silico* hemolysis and averaged *in-vivo* pfHb measurements was observed, supporting the model's predictive value.

The analysis in this chapter represents the first validation of numerical hemolysis predictions against clinical hemolysis data from a large patient cohort by providing an association between *in-silico* hemolysis predictions and *in-vivo* hemolysis data. This approach extends beyond current *in-vitro* correlations [8, 109, 170], integrating clinical *in-vivo* data of pfHb for a more comprehensive understanding and a better transferability into clinical practice. This strengthens the validity of these models, which will be instrumental in developing novel ECMO components that could potentially induce less hemolysis.

Through the interdisciplinary study design of combining *in-silico* and *in-vivo* findings, increased ECMO circuit resistance, and consequently higher pump pressure head, was identified as the primary contributor to mechanically induced hemolysis during VV ECMO support. This finding has immediate clinical relevance, as it underscores the importance of selecting ECMO circuitry with the lowest possible resistance.

This includes cannulae, with further analysis distinguishing the impact of drainage and return cannula pressures (*see* Figure 4.7) and revealing that each is substantially associated with hemolysis. This finding reinforces the conclusion that, in clinical practice, it is crucial to consider the RBP's pressure head, which represents the cumulative effect of these pressures, rather than evaluating them separately.

Consequently, minimizing ECMO circuit resistance by selecting oxygenators with the lowest possible resistance, using the shortest feasible circuit tubing, and choosing cannulae with the largest possible diameter is expected to reduce the RBP's pressure head and the risk of pump-induced mechanical hemolysis. Furthermore, apart from low cannula resistances, unfavorable flow conditions such as flow separation and recirculation should be avoided. All these aspects influence the choice of the cannula and should be carefully considered.

Previous *in-vitro* and *in-silico* investigations in the field of RBP induced hemolysis in ECMO support, notably those conducted by Schöps *et al.*, Gross-Hardt *et al.*, and Ki *et al.* [169–171],

highlighted low flow conditions as more hemolytic as high flow conditions. These studies compared few operating points at varying flow rates under constant pressures heads or rotational speeds. The present analysis broadens this view by extending the scope of constant rotational speeds or pressure boundary conditions, covering a more holistic operating range of VV ECMO (*see* Figure 4.9). In clinically relevant scenarios (within the 95 % confidence interval of operating points), the combination of high pressure head and high flow rate conditions pose a greater risk of hemolysis compared to low flow rate and low pressure head situations. It is especially evident that the low flow rate and low pressure head operating points, not always lead to increased hemolysis. These findings demonstrate that substantial hemolysis at low flows is predominantly a consequence of simultaneously occurring high pressure heads. Additionally, it can be seen that at low flow rates, a small change in pressure head results in larger variations in hemolysis than the same change in pressure head would cause at higher flow rates. This is true irrespective of the specific RBP type or hemolysis model used, suggesting a universal applicability across all centrifugal pumps used in ECMO support (*see* Table 4.1). The novelty of this finding, compared to existing literature, lies in the shift in perspective from "low flow" being inherently hemolytic to identifying "high resistance" as the key factor contributing to hemolysis in these scenarios. The low flow itself is not inherently hemolytic, as demonstrated by the low hemolysis values observed in areas with low pressure head. This nuanced understanding provides more actionable insights for clinical practice by highlighting the narrow therapeutic window in the low flow area, where optimizing ECMO support can help minimize blood trauma. This is of particular importance for the weaning phase of ECMO support, patients with low body temperature, small body surface area, carbon dioxide removal, or pediatric ECMO, as all these scenarios encounter decreased blood flow rates. Supporting this, a recent meta-analysis demonstrated that ECMO pumps at low flow rates, whether during pediatric ECMO or carbon dioxide removal in adults, are associated with a significantly higher rate of hemolysis compared to higher flow rates[175], which could potentially be explained by the current findings of increased sensitivity to higher resistances in these low flow scenarios.

Further analysis of the clinical data presented in the Appendix A.1 revealed that non-survivors exhibit higher levels of hemolysis compared to survivors (*see* Figure A.2 and Figure A.3). This observation aligns with current reports in the literature [168, 176–178]. Despite this, the global distributions of pressure head and flow rate between the two sub-cohorts are largely similar, as illustrated in Figure A.5, indicating that pump parameters alone are unlikely to account for observed differences in hemolysis between survivors and non-survivors.

Additionally, the *in-vivo* dataset shows a notable variability in individual measurements, indicating that patient-specific factors play a crucial role in hemolysis outcomes. This is further supported by the observation that not all of the 580 patients in the cohort follow the numerically predicted hemolysis trend (*see* Figure 4.10). This suggests that the investigated numerical models provide reasonable predictions for averaged operating point scenarios but are not capable of capturing individual patient-specific deviations.

These findings highlight the complexity of the ICU environment and the importance of patient-specific circumstances, which prompted the inclusion of additional parameters beyond pfHb. LDH, bilirubin, oxyHb, and deoxyHb were included in the analysis, as these were associated with *in-vivo* hemolysis. D-dimers were also included as a negative control to confirm that pump thrombosis did not significantly influence the results. In the ICU setting,

LDH can only be considered an indirect surrogate marker for hemolysis. Although increased LDH levels are known to be associated with hemolysis [55, 56], they can also result from other factors, such as cellular necrosis or increased tissue turnover [56]. Nevertheless, the trends observed in LDH levels may provide valuable insights in daily clinical practice. To the best of currently available knowledge, there is no established evidence linking oxyHb or deoxyHb levels directly to increased hemolysis. However, their strong association with hemolysis in the analysis is noteworthy. One possible explanation is that patients in the ICU who require higher oxygenation levels often receive higher blood flow rates as a priority, regardless of the pressure head. This could increase the RBP's mechanical stress on the blood, leading to greater hemolysis. However, it remains unclear whether this increased hemolysis is due to the higher energy input required by the RBP or the oxygenation state of the hemoglobin itself. Further research is needed to clarify the specific impact of hemoglobin oxygenation on hemolysis, as any causal relationship between these parameters remains speculative at this point.

The presented analysis has several limitations, both from a clinical and a modeling perspective. From a clinical perspective, the findings reflect average outcomes across the patient cohort, but individual patient-specific conditions may lead to significant deviations. Moreover, the analysis was based on a retrospective single-center dataset. Expanding the dataset with comparable *in-vivo* data from other institutions could help validate and generalize the results. Further clinical investigations may also explore differences between specific patient groups, such as VA versus VV ECMO configurations, with regard to hemolysis risk.

From a modeling perspective, the current framework is based on a stress-based hemolysis model, which, while supported by strong correlations with *in-vitro* data across different operating points [109], primarily provides relative rather than absolute predictions. Although the model was validated against *in-vivo* data, its ability to predict absolute hemolysis values remains limited. The field of blood damage modeling continues to evolve, and it would be valuable to investigate alternative hemolysis model formulations. Comparing their predictive performance against this extensive clinical *in-vivo* data may reveal whether some models offer improved accuracy or better generalizability under clinical conditions.

4.5 Conclusion

From a modeling standpoint, this work presents the first validation of numerical *in-silico* hemolysis predictions against clinical *in-vivo* data across a large ECMO patient cohort. The results demonstrate that stress-based numerical hemoylsis models are capable of predicting average hemolysis trends across varying ECMO operating conditions.

From a clinical standpoint, ECMO circuit resistance was identified as a primary driver of increased hemolysis, as it forces the RBP to generate higher pressure heads. This finding highlights actionable strategies for minimizing mechanical hemolysis during VV ECMO support. Low flow scenarios were found to be particularly critical, as even minor increases in circuit resistance under such conditions resulted in disproportionate increases in hemolysis. To mitigate these risks, low suction in the drainage cannula and low return pressures should be maintained at all times by selecting oxygenators with the lowest possible resistance, the shortest feasible tubing, and cannulae with the largest possible diameter. This is particularly

relevant in clinical settings involving low blood flow, such as during weaning, in patients with low body temperature or small body surface area, carbon dioxide removal, or pediatric ECMO, where current RBPs are not yet optimized for such conditions.

It is important to emphasize, however, that these conclusions apply to average trends across the cohort. Substantial intra- and inter-individual variability was observed in the clinical dataset, and the numerical model was not capable of capturing these patient-specific deviations. This limitation underscores the need for further investigation into the factors contributing to variability in hemolysis outcomes. In response to this need, the following chapter builds directly on this limitation by addressing the challenge of quantifying variability in the underlying *in-vitro* hemolysis measurements used to construct stress-based hemolysis models.

5 Quantifying variability in in-vitro hemolysis measurements

The general findings of this chapter are based on the article "Quantifying Experimental Variability in Shear-Induced Hemolysis to Support Uncertainty-Aware Hemolysis Models" published in *Annals of Biomedical Engineering* [179].

5.1 Introduction

As demonstrated in the previous Chapter 4, stress-based hemolysis models are capable of predicting the overall trend of *in-vivo* hemolysis across a wide range of operating conditions in VV ECMO patients. However, these predictions remain limited to cohort-averaged outcomes. The models neither reproduce absolute hemolysis values nor account for patient-specific variability. This limitation presents a significant barrier to the broader applicability of hemolysis models, particularly in clinical and regulatory contexts where model precision and individual risk assessment are essential.

As outlined in Chapter 2.2.2, *in-silico* modeling has become a powerful tool for medical device development, allowing efficient exploration of design alternatives while reducing the need for extensive *in-vitro* and *in-vivo* testing [81]. Regulatory authorities, such as the FDA, have recognized this potential and emphasize the need for rigorous VVUQ to support the use of *in-silico* models in the evaluation of blood-contacting medical devices. To this end, the FDA endorses the ASME V&V 40 standard [84], which outlines a structured approach for assessing model credibility. Within this framework, uncertainty quantification plays a central role, particularly when translating model predictions into clinically relevant decisions.

Central to any meaningful uncertainty quantification is a thorough understanding of the sources of uncertainty in the system being modeled. This includes not only uncertainties in model parameters and boundary conditions, but also in the choice of model formulation and the empirical data used for model calibration and validation. In the context of hemolysis modeling, such uncertainties are particularly relevant, as current macro-scale models are fundamentally based on empirical relationships between shear stress, exposure time and hemolysis. A major contributor to overall model uncertainty is therefore the experimental data on which these relationships are derived. While these data form the foundation for calibrating and validating hemolysis models, their inherent variability and measurement uncertainty remain largely uncharacterized. One potential source of this uncertainty is biological variability inherent to hemolysis measurements, which may arise both within the same donor across

repeated measurements (intra-donor variability) or between different donors (inter-donor variability). Despite its relevance, the relative contribution of these two sources is rarely quantified in existing hemolysis studies, even though they may substantially affect model calibration and generalizability.
Without a systematic quantification of this underlying experimental uncertainty, it is not possible to perform a credible uncertainty analysis of hemolysis predictions. This critical gap motivates the third research question addressed in the following chapter:

3. *How can uncertainties in experimental hemolysis data be quantified?*

To inform numerical hemolysis models, various experimental strategies have been pursued over the past decades. At the microscopic scale, several studies have investigated the structural membrane properties of individual RBCs and their response to applied mechanical stress, aiming to correlate membrane rupture with specific loading conditions [180–182]. While these approaches offer valuable insights into cell-level biomechanics, both their findings and the associated uncertainty are difficult to translate to the macroscopic flow conditions of cardiovascular devices, where RBCs experience vastly different and less controlled mechanical environments.
In contrast to cell-level studies, macroscopic experiments have aimed to characterize hemolysis under varying flow regimes, such as laminar versus turbulent conditions [112], and to define threshold values for hemolysis onset [183–185]. However, these measurements are highly sensitive to the specifics of the experimental setup and prevailing flow characteristics, which limits their reproducibility and complicates comparisons across different studies.
A more commonly adopted approach in macroscopic experiments is to expose whole blood samples to controlled shear stress levels for defined durations using specialized shearing devices. Since the pioneering work of Giersiepen *et al.* [99], this methodology has been widely applied across different species [107] and experimental setups [104, 106, 108]. The resulting data are typically used to derive empirical power law relationships (*see* equation (2.3)) between shear stress, exposure time, and hemolysis. These models are straightforward to integrate into numerical simulations, yet no universally valid set of power law coefficients has been identified. Instead, the fitting parameters vary substantially between studies, likely due to differences in experimental protocols, sample handling, and measurement techniques. Moreover, due to constraints on time and resources, most of these studies are designed to explore a wide range of shear conditions with only a limited number of replicates per condition (often just three). For instance, Zhang *et al.* [106] examined shear stresses ranging from 30–330 Pa and exposure times of 0.03–1.5 s across 45 conditions, each tested in triplicate. With mean coefficient of variation (CV) of the triplicate measurements for all conditions being approximately 30 %, indicating not only fitting parameter variability across different studies but also substantial within-study variability, making it challenging to identify universally valid fitting parameters.
Although these foundational studies have supported the use of numerical hemolysis modeling as a valuable tool for comparative device optimization, large discrepancies in predicted outcomes remain a substantial limitation when absolute hemolysis values are desired. Taskin *et al.* [109] demonstrated that predicted hemolysis can differ by several orders of magnitude depending on the modeling approach (Eulerian vs. Lagrangian) and the chosen fitting coefficients. Similarly, Yu *et al.* [120] identified the selection of fitting parameters as the primary source of deviation between model predictions and experimental measurements.

Beyond uncertainties in parameter selection, recent studies have raised fundamental concerns about the structure and assumptions of current hemolysis models. These include the treatment of exposure time in accelerating flows and the use of oversimplified shear stress representations [116, 121]. In addition, Craven *et al.* [110] showed that the basic power law formulation may not yield unique solutions for a specific operating point, further complicating the interpretation of model outputs.
These challenges make it clear that achieving reliable predictions of absolute hemolysis values requires addressing not only parameter uncertainty but also underlying conceptual limitations in model design. However, even the most sophisticated modeling approaches remain constrained by the quality of the experimental data used for calibration. From an experimental perspective, a root cause of this limitation is the substantial and often unquantified uncertainty embedded in existing datasets. The conventional practice of fitting a single mean curve to variable experimental values of a given study neglects inherent variability, potentially obscuring the full range of plausible outcomes and leading to models that systematically underestimate uncertainty. As a result, identifying truly representative model parameters may be unfeasible under current methodologies.
Addressing this gap requires a focused investigation of experimental variability under controlled conditions, in line with the third research question of this dissertation concerning the quantification of uncertainty in hemolysis data. To achieve this, the analysis concentrates on a single, well-defined scenario instead of distributing resources across a broad range of shear stress and exposure time combinations. By allocating a substantial number of replicates ($n = 100$) to this condition, it becomes possible to examine key components of experimental uncertainty. In particular, the relative influence of intra-donor and inter-donor variability is assessed, along with the impact of sample size on the precision and reliability of hemolysis measurements. These aspects are critical for characterizing the variability embedded in hemolysis data and for enabling the development of more robust, uncertainty-aware *in-silico* models, ultimately contributing to more credible predictions for both device development and regulatory evaluation.

5.2 Materials & Methods

To investigate intra- and inter-experimental variability in hemolysis measurements, a high-precision shearing device was developed based on a rotational rheometer. Using this setup, a total of $n = 5$ experiments were conducted with human blood obtained from healthy donors, following approval by the Institutional Review Board of the University Hospital Aachen (approval number 23-199). For preliminary testing and device calibration, porcine abattoir blood was used in accordance with previously established acquisition protocols [186].

5.2.1 Test setup

Rotational rheometers are well-established instruments known for their ability to generate controlled and well-defined flow conditions. They provide exceptional measurement precision, stable temperature control, and superior reproducibility. In the present setup a rotational rheometer (502e, Anton Paar, Austria) was equipped with a custom-made rotational bob,

capable of a maximum rotational speed of 3000 rpm. The objective was to create reproducible, supraphysiological shear stress conditions inside the rheometer. This required an optimal balance between gap size, sample volume, geometry diameter, and manufacturing tolerances. The resulting configuration was a Mooney-Ewart shearing geometry, which is essentially a combination of cylinder-cylinder and cone-plate designs. The geometry was designed to maximize the proportion of the blood sample exposed to a uniform shear rate, resulting in 98.5 % of a 180 µL blood sample experiencing a consistent shear rate of approximately 45,000 s^{-1}. Assuming a dynamic viscosity of 3.5 mPa s, the corresponding shear stress was approximately 160 Pa, mirroring mechanical circulatory support conditions [169, 187]. The sample volume of 180 µL was sufficient to perform duplicate hemolysis measurements. Figure 5.1A provides a

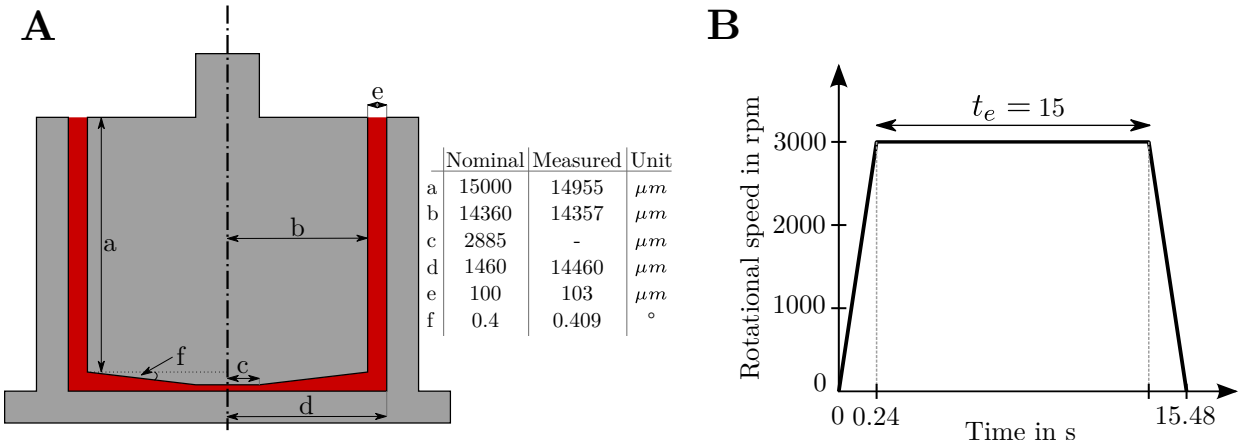

Figure 5.1: Panel A illustrates the Mooney-Ewart shearing geometry, consisting of a rotating inner bob and a stationary outer cylinder (gray). The blood sample, shown in red, fills the narrow space between the two surfaces. Nominal and measured dimensions of the geometry are listed in the accompanying table. Panel B shows the rotational speed profile: the system accelerates to 3000 rpm within 0.24 seconds, maintains this speed for an exposure time of 15 s, and then decelerates over another 0.24 seconds.

schematic of the Mooney-Ewart geometry, with a nominal gap size of 100 µm and a nominal cone angle of 0.4°. The inner and outer cylinders were machined in a single process, ensuring a concentricity of ±3 µm under operating conditions, including the maximum runout of 1 µm of the rheometer's air-bearing coupling, as certified by the manufacturer. The measured dimensions yielded an actual gap size of 103 µm, corresponding to a maximum shear rate of approximately 43,600 s^{-1} at 3000 rpm. Given the system's moment of inertia, the rheometer can accelerate to 3000 rpm in 0.24 s and decelerate to zero in the same time interval. To assess flow stability at maximum speed, the Reynolds number (170) and the Taylor number (0.0095) were calculated. Both values are well below their respective critical thresholds for the onset of flow instabilities. In particular, the Taylor number is several orders of magnitude below the critical value of 1708, confirming that the flow remains laminar. Under these conditions, shear is applied uniformly across the narrow gap without the formation of secondary flow structures such as Taylor vortices.

The exposure time (t_e) was defined as the duration during which the rheometer maintains a constant rotational speed, as illustrated in Figure 5.1B. This choice was informed by preliminary experiments conducted with porcine blood, in which exposure times ranging

from 0.5–25 seconds were tested. Based on previous literature [106, 107], it is known that increasing shear stress and exposure time tends to raise both the HI and its variability. To allocate experimental resources toward a single, well-defined operating point for assessing intra- and inter-donor variability, an exposure time of 15 seconds was selected to reliably produce measurable hemolysis while maintaining reproducibility. A target HI of approximately 0.5–1.0 %, corresponding to the higher HI range reported in the literature, was chosen based on experiments with porcine blood at a hematocrit of 30 % (*see* Figure B.1). This decision accounted for the reduced hematocrit of 25 % in the proposed experimental setup, along with the corresponding decrease in viscosity and shear stress, as well as the greater shear resistance of porcine blood compared to human blood [188]. Additional preliminary tests with porcine blood, involving only acceleration and deceleration phases without a steady shear period, showed no significant increase in hemolysis ($n = 6$), confirming the assumption that hemolysis occurs only during the plateau phase of the shear rate profile.

5.2.2 Blood collection and handling

Human blood was collected from healthy volunteers via venipuncture into 9 mL ethylenediaminetetraacetic acid (EDTA) tubes (S-Monovette K2E Gel, Sarstedt, Germany) to ensure cell stability and allow for accurate blood counts, as recommended by the International Council for Standardization in Hematology [189]. After 10 minutes of gentle mixing, the blood was centrifuged at $5000\,g_0$ for 15 minutes to produce platelet-poor plasma, which was in a next step recombined with whole blood at the appropriate ratio to achieve a hematocrit of 25 ± 1 %, reflecting values relevant to ECMO conditions adhering to best clinical practice with a transfusion trigger of 7 g/dL hemoglobin [58]. Following hematocrit adjustment, the blood was transferred using an automated pipette (Xplorer Plus, Eppendorf, Germany) at the lowest speed setting with 1000 µL tips having a minimal opening diameter of 0.77 mm. This approach minimized wall shear stress to physiological values (<10 Pa, assuming a viscosity of 3.5 mPa s) and allowed for maximal consistency in blood handling. The blood was then transferred into 1.5 mL microtubes (Safe-Lock Tube, Eppendorf, Germany) and mounted in a custom 3D-printed cylindrical holder. This assembly was gently rotated at 10 rpm (BTR5-12V tube roller mixer, Ratek, Australia) throughout the experiment day to ensure uniform mixing. Before and during the experiments, blood samples were routinely analyzed for standard RBC metrics (ProCyte Dx, Idexx, USA) as well as pH, lactate, and glucose levels (ABL 825 Flex blood gas analyzer, Radiometer, Germany). All sampling and handling steps were performed under identical conditions to ensure maximal consistency and comparability across experiments.

5.2.3 Experimental protocol and cleaning procedure

All experiments were conducted at 37 °C to replicate *in-vivo* viscosity conditions. Viscosity measurements, taken after dilution to a hematocrit of 25 ± 1 %, were performed at the start and end of each experimental day, using the same rotational rheometer. Across all five donors, the viscosity at a shear rate of $1000\,\mathrm{s}^{-1}$ had a mean value of 2.9 mPa s with a standard deviation of 0.09 mPa s. Although direct measurements at the experimental shear rate ($45{,}000\,\mathrm{s}^{-1}$) were not feasible, blood is known to exhibit only minor changes in viscosity at

shear rates above $1000\,s^{-1}$ [190, 191]. Therefore, the value measured at $1000\,s^{-1}$ is considered representative of the conditions applied in the experimental setup. For each experimental run, blood from a single 1.5 mL microtube was divided into two samples: a zero sample (180 µL) transferred into a 200 µL microtube and a rheometer sample (180 µL) subjected to 15 s of shear at 3000 rpm in the rheometer. Before each new measurement, the rheometer geometry was cleaned with deionized water to ensure osmotic hemolysis eliminated any residual RBCs. After flushing with deionized water, the system was rinsed with NaCl solution (B. Braun, Germany) and dried with purified, pressurized air. The geometry was then reheated to 37 °C and allowed to equilibrate for 10 minutes. The relative vertical distance between the rotating and static parts of the Mooney–Ewart shearing geometry was redefined before each measurement by determining the zero-gap position at a force threshold of 1 N. Subsequently, 180 µL of blood was placed at the center of the outer cylinder's bottom surface, using reverse pipetting to ensure bubble-free application. The inner cylinder then approached the measurement gap at a controlled rate, starting at 8000 µm/s and slowing progressively to 1 µm/s near the final position. This procedure minimized pre-shearing and ensured even blood distribution within the gap. Once the gap was set, the temperature was confirmed to remain within ±0.1 °C for at least 20 seconds before initiating shearing. After the shearing step, the inner cylinder was retracted following a similar speed profile, starting slowly (1 µm/s) to minimize additional shear before increasing speed once the gap geometry no longer influenced the blood sample. The sheared blood was then transferred to a 200 µL microtube and, together with the zero sample, centrifuged at $2000\,g_0$ for 15 minutes. The separated plasma was centrifuged again ($2000\,g_0$, 15 minutes) to ensure removal of any residual cellular components, then stored at −30 °C until subsequent hemolysis analysis.

5.2.4 Analysis of hemolysis

Hemolysis was quantified in accordance with DIN 58931:2021-09 using the cyanmethemoglobin (HiCN) method. Plasma samples were thawed in a 37 °C water bath for 8 minutes and diluted 1:5 (v/v) with HiCN conversion solution (fHb (HiCN), Bioanalytic GmbH, Germany). The dilutions were prepared in duplicate on a 96-well microplate (96 Well Half Area Microplate, Greiner Bio-One, Austria). Following a 30-minute incubation, pfHb was measured photometrically at 540 nm, with a reference wavelength of 680 nm, using a microplate reader (Spark Multimode Microplate Reader, Tecan, Switzerland). All readings were blank-corrected and converted to pfHb in mg/dL following Tapernon *et al.* [192]. Quality controls of pfHb at concentrations of 12.5, 39, 50 and 250 mg/dL were included on each plate to confirm measurement validity and calibration range. Each sample was measured in duplicates, and the mean value of the duplicates was used for subsequent analyses. To facilitate direct comparisons with literature values, the previously described HI, as defined in equation (2.2), was calculated. First, a ΔpfHb value was determined by subtracting the pfHb level of the zero sample from the pfHb level of the sheared sample ($\Delta\text{pfHb} = \text{pfHb}_{\text{sheared}} - \text{pfHb}_{\text{zero}}$). Importantly, both the sheared and zero samples were derived from the same 1.5 mL microtube, ensuring they shared the same blood handling and time history, thereby minimizing variability between paired measurements. This ΔpfHb was then adjusted for the individual hematocrit (HCT) and hemoglobin (Hb) values, which varied slightly between donors (*see* Table B.1 and Figure B.2). Mean HCT and Hb values from the entire experimental day were used for the

adjustment. The resulting HI thus provides a normalized measure of hemolysis that accounts for variations in blood composition, enabling a more meaningful comparison to established literature standards.

5.2.5 Statistical analysis

All statistical analyses were performed using Python (version 3.11.7) and standard scientific computing libraries, including NumPy for numerical computations [157], pandas for data handling [158], and Matplotlib for data visualization [160].
To assess whether variability in the measurements was primarily driven by inter-donor or intra-donor effects, a variance components analysis was conducted. This method allows the total variability in the data to be decomposed into distinct sources. Each measurement Y_{ij} was modeled according to equation 5.1:

$$Y_{ij} = \mu + d_i + \varepsilon_{ij} \tag{5.1}$$

where Y_{ij} denotes the j-th measurement from donor i, μ is the overall mean, d_i is the random effect associated with donor i (with variance θ^2), and ε_{ij} captures the intra-donor measurement noise (with variance δ^2). To quantify the proportion of variance attributable to donor differences, the intraclass correlation coefficient (ICC) was calculated as defined in equation 5.2:

$$ICC = \frac{\theta^2}{\theta^2 + \delta^2} \tag{5.2}$$

To characterize uncertainty and to assess how the number of replicate measurements per donor affects the estimated mean HI, a bootstrap analysis was conducted. Bootstrapping is a non-parametric resampling technique that does not assume an underlying distribution of the data. In this study, the bootstrap samples were generated with replacement, ensuring that each resampled dataset could be selected multiple times. Each bootstrap run consisted of 10,000 resampling iterations, and the 95 % confidence intervals were determined by taking the 2.5^{th} and 97.5^{th} percentiles of the resulting bootstrap distribution.

5.3 Results

This section presents the variability observed in HI measurements across donors and experimental runs, further explores intra- and inter-donor variation, and examines how sampling frequency influences the accuracy of HI mean estimates. To complement this analysis, baseline characteristics and the temporal progression of experimental parameters are provided in Appendix B.

5.3.1 HI variability among donors and experiments

Figure 5.2A–E presents the HI results for each of the five donors across 20 consecutive experimental runs. Within each donor, individual hemolysis measurements ranged approximately from 0.2–1.1 % HI, with a similar degree of variation observed across all donors. For all

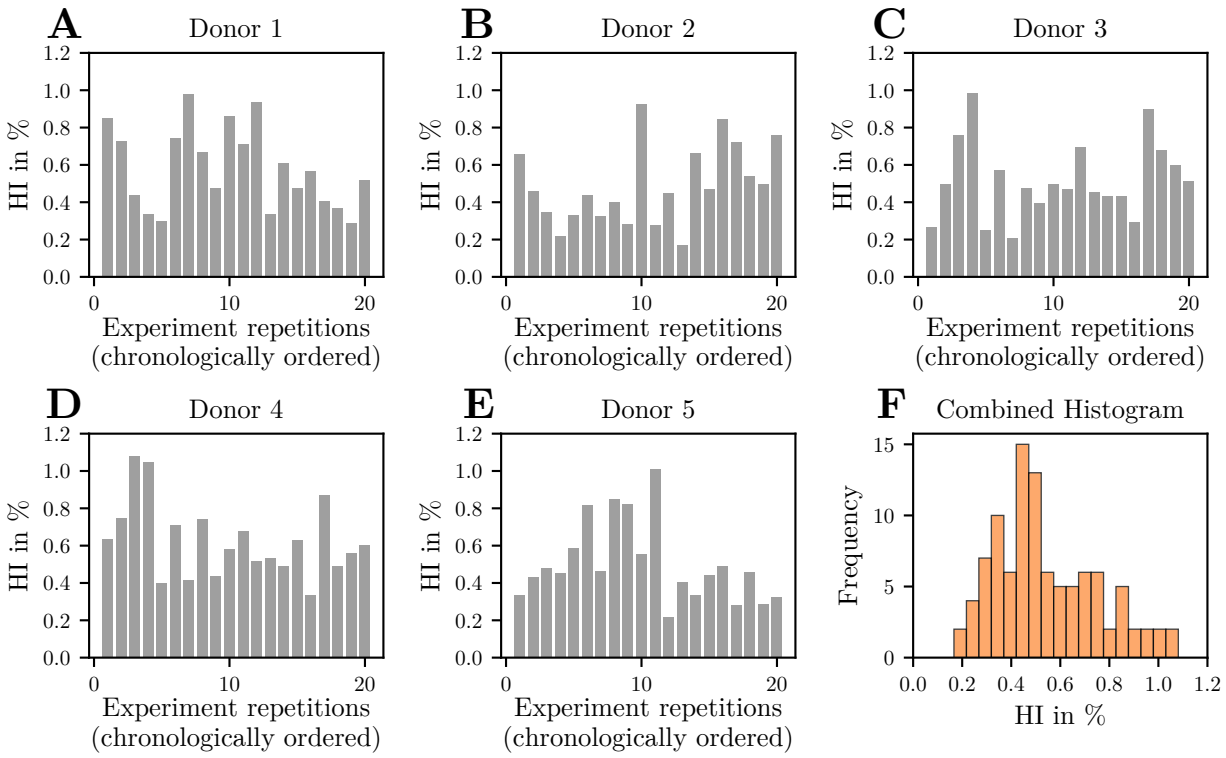

Figure 5.2: Chronologically ordered hemolysis index (HI) measurements for each donor are shown as bar charts in panels A–E, with 20 individual measurements numbered 1–20 for each donor. Panel F displays the combined histogram of all HI measurements across all donors, represented in orange.

donors no clear temporal trend was observed over the course of the 20 runs. The combined distribution of all 100 measurements (20 runs × 5 donors) is shown in Figure 5.2F, indicating a positively skewed distribution with a mean HI near 0.5 %.

5.3.2 Intra- and inter-donor variability

Figure 5.3 shows a boxplot comparing the distribution of HI values for each donor and the variation of the HI measurements between the donors. The means and standard deviations for each donor are provided at the bottom of Figure 5.3. As has already been observed in Figure 5.2, there is a similar degree of variation for each donor with similar means and standard deviations. However, the variation within each donor's measurements (intra-donor variability) is higher than the variation between donor means (inter-donor variability). Table 5.1 presents a quantification of the variation across potential sources of uncertainty, revealing an intra-donor CV of 39.3 % and an inter-donor CV of 9.5 %. Additional sources of variability that could be quantified included the measurement variability across duplicate measurements (2.2 %) and the and the photometric readout variability (0.08 %), determined by repeated absorbance measurements of the same well plate using the microplate reader. When disregarding minor variability sources, the ratio of intra- to inter-donor CV was approximately 4.1. This is further highlighted by the variance components analysis applied to the measurement data which produced $\sigma^2 = 0.044424$ and $\tau^2 = 0.001116$, resulting in an ICC of 0.0245. This indicates

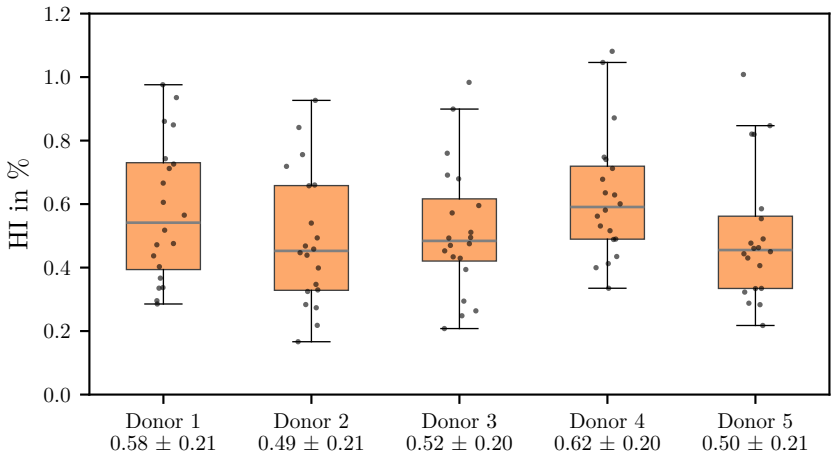

Figure 5.3: Hemolysis index (HI) measurements in percentage for each donor, displayed as box plots. The lower and upper edges of each box represent the 25^{th} and 75^{th} percentiles, respectively, while the whiskers indicate the 95^{th} percentile. The mean value is shown as a horizontal line within the box and is also provided numerically along with the standard deviation in the description of the abscissa for each donor. Individual measurements are represented by scatter points overlaid on the box plots

Table 5.1: Coefficient of variation (CV) in percentage (%) for various potential sources of variation, including intra-variability, inter-variability, measurement variability, and photometer measurement variability.

	CV in %
Intra-variability (average over 5 donors)	39.3
Inter-variability (across donor means)	9.5
Measurement variability (average over 100 duplicates)	2.2
Photometer accuracy (triplicate measurement)	0.08

that only 2.45 % of the total variance can be attributed to inter-donor differences, while the remaining 97.55 % is due to intra-donor variability.

5.3.3 Impact of sampling frequency on HI mean estimates

To assess how the number of replicate measurements per donor affects the estimated mean HI, a bootstrap analysis was conducted. Figure 5.4A–E shows the distributions of bootstrapped mean HI values when randomly selecting 3 out of the 20 measurements per donor. Across donors, mean values derived from subsets of measurements ranged from approximately 0.2–1.0 % HI, with the distribution centers aligning closely with the overall mean of all 20 measurements (dashed black lines). In other words, if experiments were conducted with only three measurements per donor, the inherent variability could result in selecting three outliers by chance. This would lead to a mean HI value at the edges of the distribution, such as 0.2 or 1.0 %, highlighting the risk of underestimating or overestimating the true mean due to limited sampling. Figure 5.4F displays the possible variation in mean HI as a function of sample size (2–20). With only two measurements, the 95 % confidence interval (CI) around the mean was ±0.3 % HI. Increasing the sample size to five measurements narrowed the CI to ±0.2 % and increasing it further to twenty measurements reduced the CI to ±0.1 %. This pattern

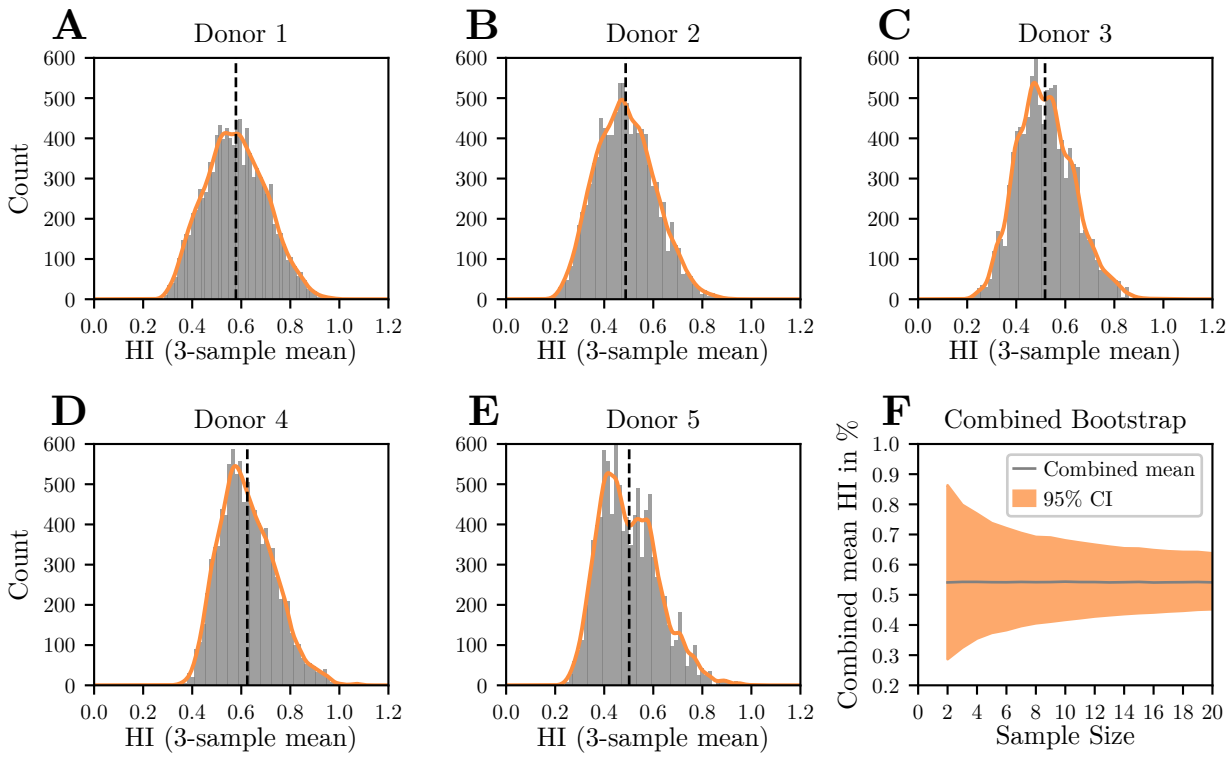

Figure 5.4: Bootstrapped distributions of possible mean values for a theoretical sample size of 3 for each donor are shown in panels A–E. The distributions are represented by gray histograms and approximated using a Gaussian kernel density estimate (orange). The black dashed line indicates the mean value for each donor based on the original 20 samples measured in this study. Panel F displays the bootstrapped confidence interval progression across sample sizes ranging from 2–20 for all combined samples (orange), with the total mean of all donors represented as a gray line.

demonstrates an asymptotic decrease in variability as the sample size grows. For illustration, increasing the sample size from twenty to a theoretical one hundred measurements (requiring eighty additional samples) would reduce the CI to $\pm 0.05\,\%$ HI.

5.4 Discussion

The aim of the analysis was to dissect the variability inherent in hemolysis measurements under a single, well-defined shear stress condition and exposure time, with a focus on distinguishing between intra-donor and inter-donor sources of variation. A total of 100 measurements were conducted under identical conditions, comprising 20 replicates per donor across 5 donors. This approach enabled not only quantification of the mean hemolysis level, but also an assessment of the underlying variability. The resulting insights have important implications for the design of future experiments and for the incorporation of uncertainty into numerical hemolysis models. Based on the findings of Chapter 4 it was hypothesized that using a low sample number in experimental hemolysis measurements may obscure results due to substantial intra-donor and inter-donor variation.
Within the present experimental analysis intra-donor variability was found to be approxi-

mately four times greater than inter-donor variability. This challenges the commonly held assumption that donor-specific factors [193] or inter-laboratory handling protocols [140] constitute the dominant sources of variability. Despite all donors being healthy, having similar physiological baseline parameters and undergoing identical sample handling procedures, the range of hemolysis index values within a single donor exceeded the variation observed between donor averages. Consequently, the expected donor-specific differences may be overshadowed by these intra-donor factors. One potential explanation for this inherent stochasticity observed across all donors could be RBC fragility [117, 120]. Such stochasticity may be further amplified by the small sample volume of 180 µL, where the composition of RBCs or the influence of outliers may have a disproportionately large effect. Each of the 180 µL samples contained approximately 500 million RBCs and a HI of 1 % within this sample volume corresponds to the release of hemoglobin equivalent to the content of approximately 2 million RBCs. This estimate assumes complete rupture of the affected cells, without accounting for potential pre-lytic hemoglobin release through transient pores or membrane thinning. Based on this assumption and the sensitivity of the measurement, a preliminary analysis (*see* Figure B.3) indicates that even subtle variations in the measured cell volume distribution across samples may contribute to intra-donor differences in HI measurements of approximately ±0.1 %. While this exploratory analysis does not fully account for the total observed intra-donor variability, it provides a plausible explanation and identifies a potential direction for future studies aimed at investigating the origins of variability in hemolysis measurements. It should also be noted that despite substantial effort and resources invested in developing a highly precise shearing geometry, absolute accuracy and reproducibility cannot realistically be attained in practice. Given that only a small fraction of RBCs needed to rupture to influence the results, small deviations in operating tolerances such as micrometer-scale variations in gap size or mechanical runout in the cone-plate region may have also contributed to the observed variability.

In the context of macroscopic blood flow through blood-contacting medical devices, other factors such as device variability, geometric uncertainties, blood handling, operator technique, and circuit variability contribute to experimental variability, making it challenging to determine whether these findings are directly translatable to such experiments. Nonetheless, the experimental setup was designed to minimize all other sources of variation, allowing intra-donor variability to be isolated and quantified under conditions of exceptional experimental control.

An analysis of the influence of sample count on the precision of mean hemolysis estimates revealed that increasing the number of replicates substantially narrows the confidence intervals. Two measurements yielded broad variability (±0.3 % HI around the total mean), but just five measurements reduced this uncertainty to ±0.2 % HI. As shown in Figure 5.3, the combined mean HI of all donors was 0.55 % HI, thus only performing two measurements might over- or underestimate the true HI by up to 60 %. From an experimental standpoint, this challenges the conventional practice of using only a few replicates (often triplicates), which may underrepresent the true variability. From a modeling perspective, this underscores the importance of incorporating uncertainty quantification measures to represent the full distribution of possible outcomes, rather than relying solely on mean values. These findings align with growing regulatory expectations from regulatory bodies such as the FDA, which emphasize VVUQ as essential components of robust computational models for medical devices.

Since most numerical hemolysis models currently rely on empirical relationships (e.g., power law models), the quality and representativeness of the underlying experimental data are critical. Existing literature often attributes differences in reported hemolysis values to factors such as species disparities [107] or device-specific designs and consequently different blood flow conditions [117]. The findings of this chapter demonstrate that substantial intra-donor variability can occur even under fixed conditions involving a single species, controlled shear stress, and exposure time. This observation implies that some discrepancies across studies may be attributable not only to inter-species differences or variations in device design, but also to limited replication at specific operating points. Supporting this, preliminary experiments with porcine blood across a range of exposure times (*see* Figure B.1) suggest that intra-donor variability is not restricted to a single shear condition, but remains present across different flow regimes and species. Relying on as few as three samples to represent a single condition, as is common in many published studies, may lead to considerable uncertainty and potentially misleading model calibration. This finding is not only relevant for numerical hemolysis calibration but also potentially applicable to a wide range of experimental *in-vitro* investigations involving blood, as previously demonstrated in the context of material hemocompatibility with an intra-variability of 40 % CV [194].

By highlighting the inherent variability in hemolysis measurements, the findings of this chapter encourage a re-evaluation of current experimental protocols. Rather than distributing limited resources across numerous operating points with minimal replication, greater benefit may be achieved by increasing replication at selected key conditions. This approach can reveal variability patterns that would otherwise remain undetected in sparsely replicated datasets. In particular, reliance on a small number of replicates, such as the common use of triplicates, may underestimate the true uncertainty due to substantial intra-donor variability. For instance, mechanistic studies or efforts to reduce technical noise may benefit from increased replication per donor, whereas studies focused on capturing population-level trends or stratifying donor responses may prioritize a broader donor base. These results aim to support experimentalists in making informed trade-offs between the number of donors and the number of replicates per donor, depending on the specific aims of their study.

Furthermore, incorporating variability directly into hemolysis models, rather than relying on a single mean curve, would represent a substantial step toward developing uncertainty-aware *in-silico* predictions. These more comprehensive models could assist device developers and regulators in better understanding the range of plausible hemolysis outcomes, ultimately informing safer and more reliable device designs.

Despite these advances, several limitations must be acknowledged. First, the analysis was restricted to a single shear stress and exposure time condition, using one device configuration and a consistent blood handling protocol. While this enabled strict experimental control and reproducibility, the generalizability of the findings to other shear rates, exposure times, sample volumes, or device geometries remains to be demonstrated. Second, the focus was placed on hemolysis outcomes, without an in-depth examination of the underlying cellular or biochemical mechanisms. Future research may benefit from combining high-replication protocols with detailed cellular analysis or advanced imaging techniques to identify the specific sources of intra-donor variability at the cellular or molecular level.

5.5 Conclusion

In summary, the chapter demonstrates that intra-donor variability represents a substantial and previously underrecognized source of uncertainty in hemolysis measurements. Although the experiments in this chapter were limited to a single shear stress and exposure time condition and are therefore not intended for fitting new hemolysis models directly, they provide critical insights into the inherent variability of experimental data. This variability has important implications for the interpretation and application of hemolysis models, particularly in the context of parameter fitting, where assuming low measurement noise can lead to biased estimates. The findings emphasize the need for both increased experimental replication and robust uncertainty quantification to improve the credibility of numerical predictions. Building on these results, the following chapter introduces an uncertainty-aware modeling framework that incorporates this inherent variability into the computational prediction of hemolysis.

6 Uncertainty-aware hemolysis modeling

The general findings of this chapter are based on the article "Towards Uncertainty-Aware Hemolysis Modeling: A Universal Approach to Address Experimental Variance" published in the *International Journal for Numerical Methods in Biomedical Engineering* [195].

6.1 Introduction

The previous Chapter 5 revealed considerable variability in controlled hemolysis experiments, demonstrating that even under strictly defined shear stress and exposure time conditions, both intra-donor and inter-donor variability can arise. This challenges the practice of using averaged experimental data to calibrate macroscopic hemolysis models, as it implicitly assumes that such data offer a consistent and low-noise basis for model fitting. As a result, these models may yield misleading averaged predictions and lack the transparency required for applications where predictive uncertainty must be quantified and trusted. Recognizing this limitation gives rise to the fourth central research question of the dissertation:

4. *How can these experimental uncertainties be incorporated into uncertainty-aware in-silico hemolysis models?*

To address this question, the present chapter investigates the sensitivity of fitted model parameters and introduces an approach to explicitly incorporate experimental uncertainty into numerical hemolysis predictions.
As outlined in Chapter 2.3.1, fully resolving cellular damage mechanisms at the device scale is computationally unfeasible. Therefore, current models approximate the cumulative effects of shear-induced damage using empirical relationships fitted to *in-vitro* data. A widely used approach is the power law model introduced by Giersiepen *et al.* [99] (*see* equation (2.3)), which expresses hemolysis as a function of shear stress and exposure time. Experimental data are used to fit the parameters in $HI = C \cdot \tau^{\beta} \cdot t^{\alpha}$, which is then applied in numerical simulations to predict hemolysis throughout the flow field. Despite its utility, the power law method primarily allows for relative comparisons rather than predicting absolute hemolysis values. Numerous studies have attempted to refine this approach by investigating blood from different species [99, 107, 108], employing various apparatuses [104, 105, 108, 173], applying different numerical techniques [101, 124, 126, 196, 197], or using advanced fitting algorithms for parameter identification to enable device-specific formulations [110, 111]. However, no

universal parameter set for the power law method has been established yet. A detailed review of hemolysis prediction with CFD models is presented by Yu *et al.* [120] and a more general overview of shear-induced hemolysis has been performed by Faghih *et al.* [117].
As outlined before, a critical limitation of power law models is their inability to account for inherent variance in experimental data, as the fitting process results in a loss of this variance, leading to a representation that only captures an average scenario of hemolysis. This raises concerns about the robustness of the fitting approach used to determine the optimal model parameters ($C,\ \alpha,\ \beta$), as it should accurately mirror the original hemolysis signal of the underlying experiment. Notable attempts to address this robustness include Mohammadi *et al.'s* [150] adjustment of model parameters to calculate probability distributions of the model parameters and Craven *et al.'s* [110] device-specific hemolysis model, that relied on a combination of computational models and experimental device hemolysis data. Both studies highlighted a notable sensitivity of resulting hemolysis to changes in model parameters, underlining the importance of finding the correct parameters. However, both approaches lacked consideration of experimental variance in their models.
To address the limitation outlined above and respond to the fourth central research question of this dissertation, the chapter begins with a sensitivity analysis of the power law parameters, demonstrating how small fluctuations in the underlying data can lead to substantial variation in predicted hemolysis. Building on these findings, a probabilistic modeling framework is introduced that incorporates experimental variability directly into the parameter estimation process. Instead of relying on single best-fit values, this approach infers parameter distributions that reflect the observed uncertainty in hemolysis measurements. The resulting model enables uncertainty-aware predictions that better capture the plausible range of outcomes, offering a more robust and transparent foundation for simulating shear-induced blood damage in blood-contacting medical devices.

6.2 Materials & Methods

Figure 6.1 provides a visual summary of the analyses conducted and methods employed in this chapter. Published experimental data describe the relationship between the Hemolysis Index (HI), exposure time (t), and shear stress (τ), with three repeated measurements ($n = 3$) at each operating point. The measurements were obtained using a custom-built Couette-type shearing device, as described in [106], which enabled precise control of flow conditions. Exposure times ranged from 0.03–1.5 s, and shear stresses from 30–330 Pa, using ovine blood as the test species. Building on this data this chapter introduces a novel probabilistic approach to hemolysis prediction and compares it to the conventional deterministic method currently used for predicting hemolysis in blood-contacting devices. At first the deterministic approach is analyzed in detail, revealing the sensitivity between experimental observations and model predictions. Then the probabilistic approach including the Bayesian parameter estimation method MCMC and the ROM technique NIPCE of Chapter 3 is introduced. In the end, the deterministic and probabilistic hemolysis models are compared using data from the FDA benchmark RBP [151]. For easy reproducibility of the presented probabilistic approach, a python script to create the MCMC model as well as synthetic data to test the model can be found in the corresponding publication [195].

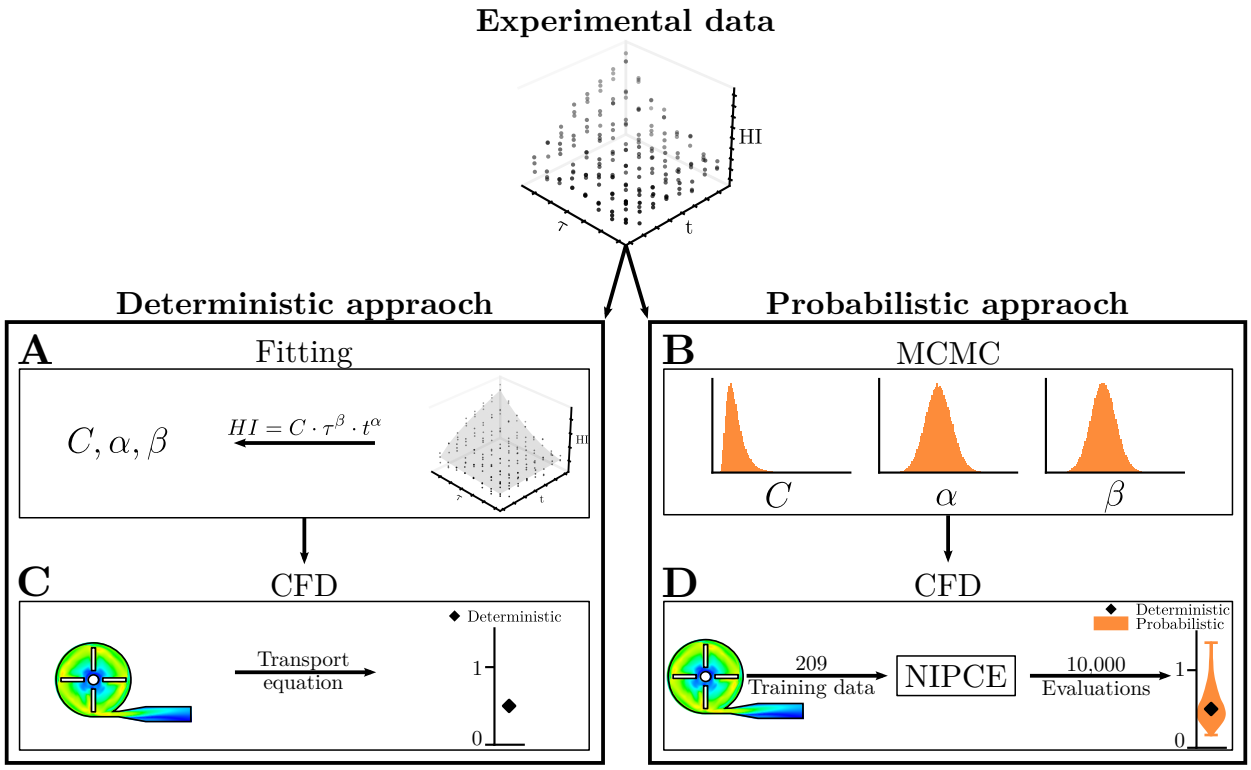

Figure 6.1: Graphical overview of deterministic and novel probabilistic approach. Starting from experimental data, A and C show the traditional deterministic approach with experimental data fitting and CFD solution steps respectively. In B and D the probabilistic approach is detailed with MCMC and CFD solution including NIPCE calculation, respectively.

Deterministic approach

Using the power-law equation (2.3), the deterministic approach begins with fitting a power law model to the experimental data describing the relationship between shear stress, exposure time and hemolysis under well-controlled experimental conditions. For this, various fitting algorithms can be employed to achieve an optimal solution based on all experimental data, described by the optimal parameters C, α, and β of equation (2.3). This fitted equation can than be utilized to predict deterministic hemolysis values (e.g., MIH) in more complex flow scenarios typically encountered in realistic blood-contacting devices like RBPs. To achieve this, equation (2.3) is reformulated into a transport equation and solved using CFD under specified boundary conditions, as described in more detail in Chapter 2.3.1.
For all results presented using the deterministic approach, the original model parameters determined by the multivariate regression method from Zhang *et al.* [106] were used ($C = 1.228 \times 10^{-5}$; $\alpha = 0.6606$; $\beta = 1.9918$).
But regardless of the fitting method, the sum of squared errors (SSE) between the experimental observations and model predictions can be calculated:

$$\mathrm{SSE} = \sum_{i=1}^{n} \left(y_i - \hat{y}_i \right)^2 \tag{6.1}$$

Here, y_i represents the experimental observation of HI, and $\hat{y}_i$ denotes the model prediction of equation (2.3) with specific parameters C, α, and β. The variable n describes the number of observed experimental values. To explore the possible parameter space and identify potential multiple optimal solutions of the deterministic approach of equation (2.3), an initial global grid search with broader parameter ranges was conducted to determine appropriate ranges for a more detailed analysis. Subsequently, a refined grid search was performed over the narrowed values for $C \in [1 \times 10^{-5}, 1 \times 10^{-3}]$, $\alpha \in [0.1, 1]$, $\beta \in [1.2, 2.5]$, using 1000 evenly-spaced values for α and β and 12 values for C, resulting in 12,000,000 possible combinations that were analyzed.

Probabilistic approach

The probabilistic approach uses the same experimental data of Zhang *et al.* [106] as the deterministic approach, employing the MCMC method to determine the posterior distributions of the model parameters in equation (2.3) with the highest likelihood of observing the provided experimental data. This is in contrast to the deterministic approach which minimizes one objective function and delivers only one fitting parameter set. The samples drawn from the posterior distributions of C, α, and β are transformed into hemolysis values using CFD, similar to the traditional approach. To achieve a meaningful posterior distribution of hemolysis values, this sampling process must be conducted very frequently. Therefore, the NIPCE-based ROM approach introduced in Chapter 3 was used to efficiently predict hemolysis based on varying model parameters.

6.2.1 MCMC implementation

MCMC methods in Bayesian statistics estimate posterior distributions of model parameters by updating prior assumptions based on observed data. The prior distribution represents the initial beliefs about the possible values of the parameters before incorporating experimental data. The posterior distribution, in contrast, is the updated probability distribution of the parameters after integrating the experimental evidence using Bayes' theorem. Through repeated sampling, MCMC approximates this posterior distribution, allowing an unbiased exploration of the parameter space while accounting for uncertainties in both the prior knowledge and the observed data. Further introductory details of this method can be found in Chapter 2.3.3 and [198, 199]. By leveraging MCMC, this chapter accounts for uncertainties in experimental hemolysis data, refining the parameter fitting process and improving the reliability of numerical hemolysis predictions. To this end, the Python package PyMC [200] was used to implement MCMC, with experimental data from Zhang *et al.* [106] and equation (2.3) forming the basis for sampling the posterior distributions. An additional parameter σ was introduced to capture the variability in the experimental data that cannot be explained by the model parameters C, α, and β. To avoid introducing excessive domain knowledge into the prior assumptions, uninformed priors were used. Specifically, the prior distributions were set as uniform (0–1) for C and normal (mean = 0, std = 1) for the other parameters. While defining normal distributions around deterministic values could potentially improve MCMC convergence, an uninformed approach was chosen to ensure that the posterior distributions were determined solely by the experimental data. This decision prevents prior assumptions

from influencing the results and ensures an unbiased exploration of parameter uncertainties. The likelihood of observing the HI was modeled by a Student's t-distribution:

$$HI \sim t\Big(\mu = \hat{y}_i,\ \sigma = \sigma,\ \nu = \tfrac{1}{30}\Big) \tag{6.2}$$

where μ represents the central value around which the data is distributed. For each combination of C, α, and β this central value is given by equation (2.3). The scale parameter σ describes the spread of the data around the central value capturing the measurement noise variability during the sampling process. The heaviness of the tails (degrees of freedom) is described by ν following an exponential distribution with a rate parameter of 1/30. The Student's t-distribution was chosen for its robustness to outliers, effectiveness with small sample sizes, and ability to handle uncertainty in variance estimates. During MCMC sampling, σ is sampled alongside C, α, and β. The sampler explores the joint posterior distribution of all parameters, accounting for the interdependencies between them. This simultaneous sampling ensures that the uncertainty expressed in σ is properly propagated to the estimates of C, α, and β. This means that σ directly affects the probability of observing the data given the model parameters and that the posterior distributions of C, α, and β incorporate the uncertainty in σ. Supporting information on this statement can be found in the Appendix C (*see* Figure C.2). Using the NUTS algorithm [144], four chains with 50,000 samples each were sampled after a burn-in period of 1000 samples with a target acceptance rate of 0.95. The resulting 200,000 samples per MCMC run provided the basis for the posterior distributions. Chain convergence was assessed using trace plots and Gelman-Rubin Diagnostic, with an example trace plot provided in the Appendix C (*see* Figure C.1). Models with a constant C parameter were also computed. For different values of C, the distributions of α, β, and σ took on different forms. Bayesian optimization identified the C value that yielded the smallest median value of the σ distribution. For this C was varied within the bounds of 1×10^{-6}–1×10^{-3}, and the optimal value of $C = 3.515 \times 10^{-3}$ was found within 100 iterations using the *gp_minimize* function of the Python package *scikit-optimize* [201] (*see* Figure C.5).

6.2.2 CFD model

To compare the probabilistic and deterministic methods, this chapter again utilized the FDA benchmark RBP geometry [151]. This RBP has been extensively investigated for hemolysis in a multicenter study [8, 92], making it an ideal basis for comparing hemolysis predictions from the models with experimental and simulation data. As already described in Chapter 3, the CFD model was previously validated by Gross-Hardt *et al.* [152], where detailed information on the mesh and simulation parameters is provided. To summarize, the geometry was meshed using the ANSYS 2021 R1 Meshing tool (ANSYS Inc., Canonsburg, USA), employing unstructured tetrahedral elements with prism layers. This meshing approach resulted in a total element count of 9.3 million cells, corresponding to the "medium" mesh configuration described in the earlier study [152]. The CFD simulations were performed using the ANSYS CFX solver, with the k-ω shear stress transport model selected for turbulence modeling. Unlike Gross-Hardt *et al.* [152], steady-state simulations were conducted to reduce computational effort for generating training data. This approach was validated against experimental data [92], demonstrating that the simplification to steady-state simulations is

justified due to the minimal differences observed when compared to unsteady simulations and the validation dataset. For hemolysis calculations the shear stress tensor was reduced to a scalar shear stress (SSS) using equation (2.6). With this SSS the hemolysis value for each simulation was calculated in a post-processing step using the volume integral formulation (equation (6.3)) of Garon and Farinas [114] resulting in a numerical equivalent of the MIH.

$$MIH = \left(\frac{1}{Q} \int_V \left(C\, SSS^{\beta} \right)^{1/\alpha} \mathrm{d}V \right) 10^6 \tag{6.3}$$

With Q being the volume flow rate of the inlet and V the volume of the whole computational domain. With this formulation, Gauss's theorem ensures that, in a steady-state flow and assuming zero HI at the inlet, the volume integral of the transported HI species over the entire computational domain is equivalent to its surface integral over the outlet. As a result, once the flow field has been computed, hemolysis can be evaluated for different sets of fitting parameters through a simple post-processing step by calculating the relevant volume integral. This approach streamlines the analysis, as it avoids the need to rerun the full CFD simulation for each parameter combination.

6.2.3 Reduced-order model (ROM)

Reduced-order modeling is a computational technique used to approximate computationally intensive simulations with substantially lower computational cost while maintaining accuracy. In the context of this chapter, NIPCE was applied, as a ROM approach that efficiently calculates hemolysis values for 10,000 variations of fitting parameters. More detailed information on this technique can be found in Chapter 2.3.2 and Chapter 3. Instead of solving equation (6.3) for each new fitting parameter set, NIPCE expresses the output quantity (MIH) as a weighted sum of polynomial basis functions of the input parameters. The coefficients of this polynomial expansion are determined through regression, using a set of training data generated from evaluations of equation (6.3). Once trained, the NIPCE model provides rapid evaluations of new parameter sets by simply evaluating the polynomial approximation rather than running the more computationally expensive task of solving equation (6.3). The MCMC posterior distributions were fitted using log-normal distributions truncated at the lower 1 % and upper of the 99 % intervals to avoid extreme values in the subsequent sampling process (*see* Figure C.3). These distributions were then combined into a multivariate distribution. Next, 200 samples were drawn from this multivariate distribution using Latin Hypercube Sampling. These samples, along with combinations of the maximum and minimum values of the individual distributions, were used as training data for the NIPCE model. Hemolysis was determined for these parameter combinations using the CFD approach (equation (6.3)) described above. The resulting dataset of fitting parameters and corresponding hemolysis values was used to train the NIPCE model with a polynomial order of 4, transforming a large number of MCMC model parameter samples into a distribution of hemolysis values within seconds. As described in Chapter 3, verifying the NIPCE model is essential to ensure reliable predictions. The corresponding verification process and sampling process are illustrated in Appendix C (*see* Figures C.3 and C.4).
Generating the training data points using CFD post-processing on a standard laptop took

approximately 60 seconds, and NIPCE training with these samples took 10 seconds. Hemolysis with uncertainty quantification was determined through 10,000 forward simulations of NIPCE for each simulation condition taking less than a second for all six conditions simultaneously. This approach reduced the post-processing time for 60,000 volume integrals from approximately 5 hours to just over a minute, demonstrating the efficiency of the ROM.

6.2.4 Model comparison

To compare the deterministic and probabilistic models, experimental hemolysis data of Conditions 1–6 from Ponnaluri *et al.* [8] was used. RIH values were extracted from Figure 11 c) of [8] using a graph digitization tool. According to Equation 4 in Ponnaluri *et al.* [8], the RIH values correspond to the MIH values, normalized to Condition 5 of the study. For the probabilistic approach, this normalization involved dividing by the median of the posterior MIH distribution of Condition 5.
In addition to the comparison with the experimental data of Ponnaluri *et al.* [8], the probabilistic model was also evaluated using synthetically generated training datasets to assess the influence of varying numbers of repetitions per operating point on model performance. The synthetic data were created by discretizing the parameter space spanned by shear stresses from 50–320 Pa and exposure times from 0.04–1.5 s into 80 equidistant operating points, arranged on a uniform grid. At each of these points, a reference HI value was computed using the original model parameters reported by Zhang *et al.* [106]. To simulate measurement variability, a normal distribution was defined around each reference value. The variance of these distributions was scaled linearly with the HI magnitude, starting from 1×10^{-4} for the lowest HI values and increasing up to 1×10^{-1} for the highest. This approach reflects the experimentally observed trend of greater variability at higher levels of hemolysis. Based on these distributions, three synthetic training datasets were constructed by randomly sampling 3, 5, or 10 values from each of the 80 operating points. Sampling was performed from truncated normal distributions restricted to the central 99 % confidence interval, thereby avoiding extreme outliers while retaining realistic variability. Reproducibility was ensured by assigning each operating point a unique deterministic seed, derived from its coordinates and a fixed global base seed. This method guarantees consistent sampling across repeated runs while preserving local randomness in the generated datasets.

6.3 Results

This section analyzes the sensitivity of the power law model to variations in fitting parameters, examines dependencies between them, and evaluates probabilistic model predictions in comparison to deterministic predictions and experimental data across multiple operating conditions of the FDA benchmark RBP geometry [92].

6.3.1 Fitting parameter sensitivity

Figure 6.2A shows discrete contour slices of the SSE in a three-dimensional plot for the parameter space α (0.1–1), and β (1.2–2.5), with 12 fixed, equidistant C values in the

interval (1×10^{-5}–1×10^{-3}). The size of the low SSE regions decreases from high to low C values. To maintain a similar SSE at a given C level, β can only vary slightly (± 0.02), while α can vary more (± 0.3). This indicates that the optimal fit for HI is more sensitive to β

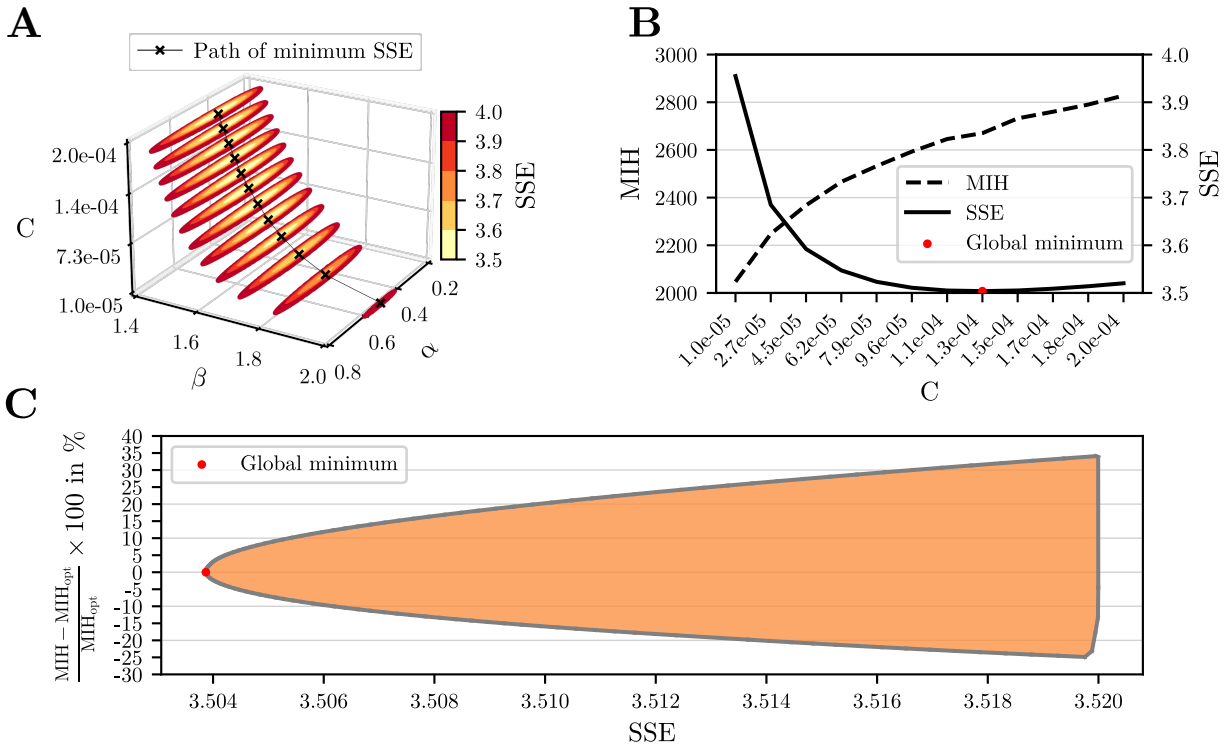

Figure 6.2: Three-dimensional contour plots depicting the Sum of Squared Errors (SSE) of experimental data against the power law model across the parameters C, α, and β. The black line with cross symbols indicates the minimum SSE value for each level of C (panel A). In Panel B the minimum SSE value for each C level is plotted in solid black and the corresponding Modified Index of hemolysis (MIH) for Condition 5 of the FDA RBP benchmark is shown in dashed black. In Panel C the range of possible hemolysis predictions for Condition 5 of the FDA RBP benchmark setup is displayed for increasing SSE values. The range of possible hemolysis predictions (min, max) is shown in orange and calculated as the difference to the hemolysis value of the global SSE minimum parameter set, which is depicted by a red dot.

than to α. Notably, despite different C levels, several regions exhibit similarly low SSE values between 3.5 and 3.6. This is further highlighted in Figure 6.2B, displaying the optimal SSE value for each C level in solid black, along with the modified index of hemolysis (MIH) of the FDA pump setup (Condition 5) for each optimal SSE value in dashed black. The global SSE minimum is marked by a red dot. Notably, the gradient of this minimum with respect to C is small, resulting in a relatively flat region around the optimum. This suggests that a broader range of C values can yield nearly identical SSEs, indicating a degree of flexibility in fitting parameter selection. Figure 6.2C extends this analysis by considering the full parameter space, incorporating values around the global minimum of not only C but also α and β. The plot visualizes the range of possible hemolysis predictions for Condition 5 of the FDA pump setup, demonstrating how small variations in SSE affect the prediction results. The plot shows that even a slight increase of around 0.5 % in SSE from the global minimum can lead to hemolysis predictions deviating by up to 60 % from the optimal value. This highlights the substantial

impact of minor SSE variations on hemolysis estimations and underscores the sensitivity of the model to fitting parameter selection.

6.3.2 Fitting parameter correlation

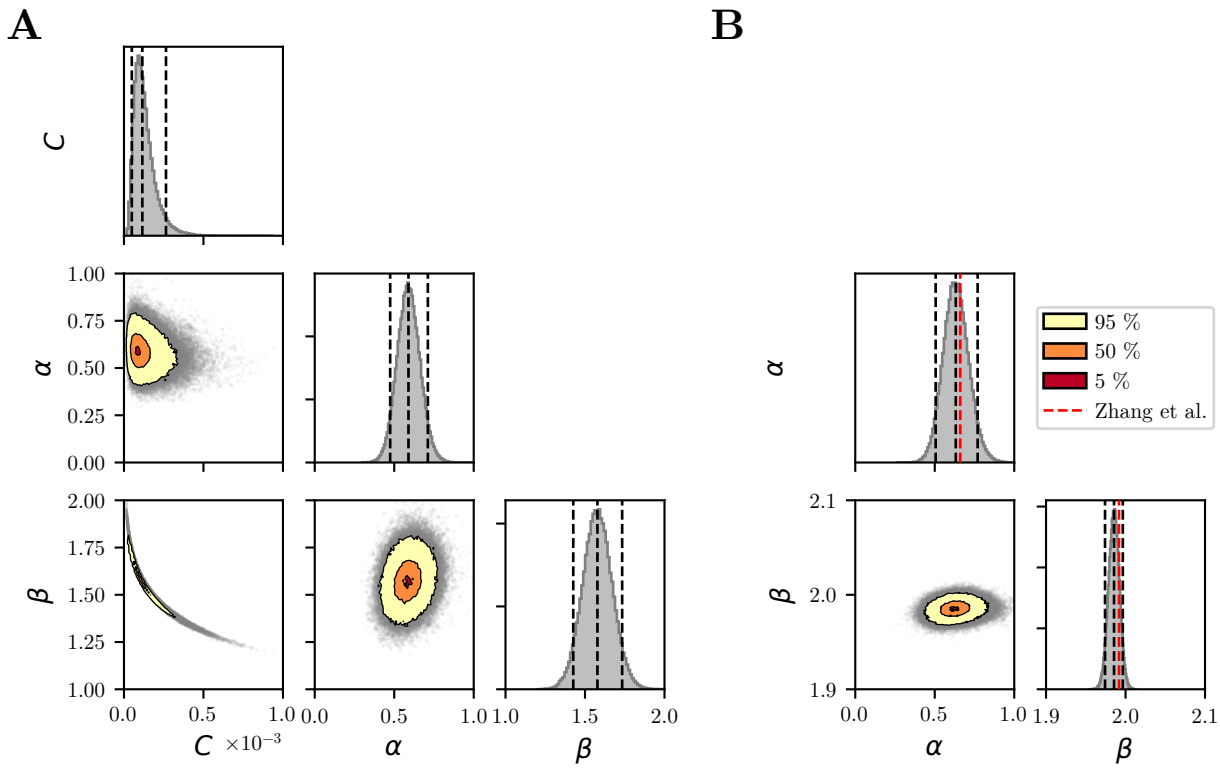

Figure 6.3: Displays the corner plot of the MCMC method for the power law model with parameters C, α, and β in Panel A, and for the model with fixed $C = 1.228 \times 10^{-5}$, with variable α and β in Panel B. The histograms show the median, along with the 5 % and 95 % intervals, indicated by dashed lines. The scatter plots are overlaid with contour plots representing the 95, 50 and 5 % confidence intervals

Figure 6.3A presents the results of the MCMC method in the form of a corner plot. It displays the individual distributions of the parameters C, α, and β, as well as their relationships in a combination of scatter and contour plots. The contour levels of the 95, 50 and 5 % confidence intervals are depicted in yellow, orange and red inside the scatter plots, respectively. It is clearly visible that the parameters C and β are correlated. This implies that no uniquely identifiable set of parameter distributions can be found, indicating redundancy in the model. An increase in one parameter can be compensated by the decrease in the other one. This observation indicates that defining C with a constant value and fitting the other two remaining parameters is a better strategy than optimizing all three parameters. In Figure 6.3B, the corner plot with a constant $C = 1.228 \times 10^{-5}$, as used in the original Zhang *et al.* [106] model, is shown. It can be observed that through MCMC sampling, the median of the α posterior distribution has settled at 0.632, and the median of the β posterior distribution at 1.985. Additionally, no correlation between the parameters is evident. However, with a constant C model, it is not guaranteed that the underlying experimental data can be

optimally reproduced. To examine the influence of the constant C on model performance, the distribution of σ, the parameter in the MCMC model that quantifies the residual error between predicted and measured hemolysis data, was evaluated across a range of fixed C values. Figure 6.4 shows in panels A–C the various distributions of α, β, and σ with

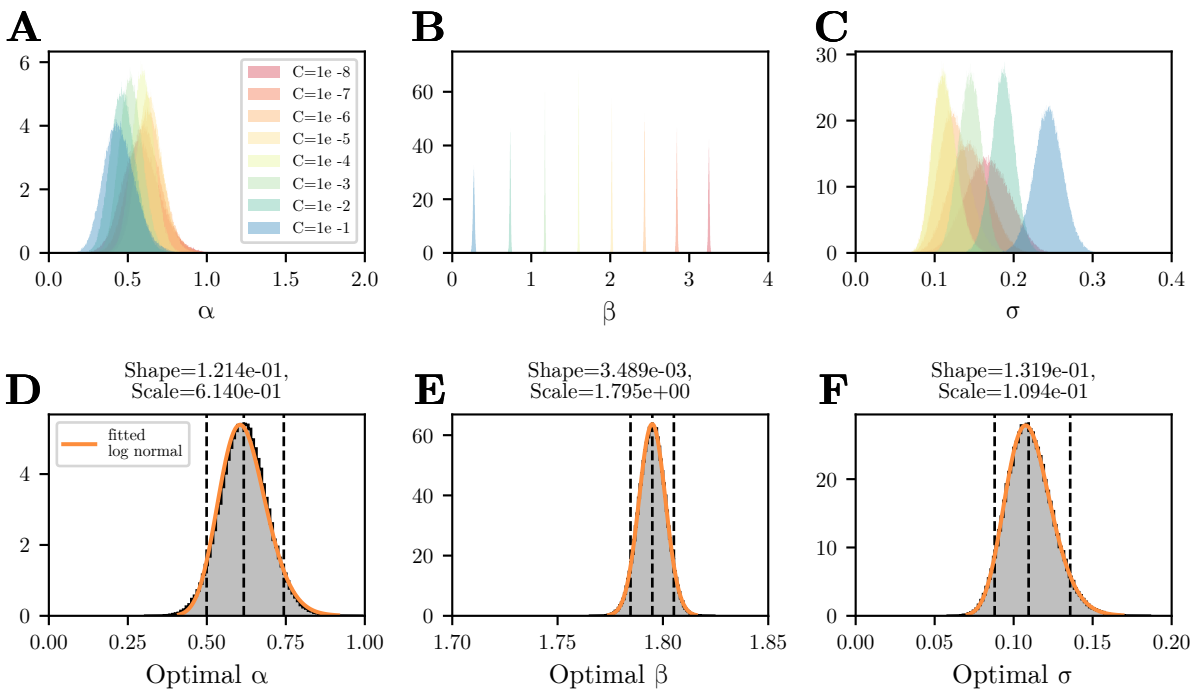

Figure 6.4: Panels A–C illustrate the influence of different constant C levels, ranging from 1×10^{-1} to 1×10^{-8}, on the parameter distributions of α, β, and σ, respectively. Panels D-F depict the parameter distributions for the optimal $C = 3.515 \times 10^{-3}$ for α, β, and σ. The dashed lines indicate the 5 %, 95 %, and median values of the individual distributions.

different constant C values ranging from 1×10^{-1} to 1×10^{-8}. It is evident that as the C value increases, the median value of the β distribution decreases, while the width remains relatively constant. For σ it is observed that there must be a minimal value between 1×10^{-1} to 1×10^{-8}. To identify the optimal constant C that minimizes σ and yields the best fit to the experimental data, a Bayesian optimization was performed on σ (*see* Figure C.5). This minimization of the median of the σ distribution through Bayesian optimization led to an optimal C value of 3.515×10^{-3}. The parameter distributions of this optimal model, providing the best fit to the data and minimizing the residual error between the observed and predicted hemolysis values, are shown in Figure 6.4D–F. Alongside the histogram created from the MCMC samples, the 5 % and 95 % interval boundaries as well as the median values are indicated with dashed lines. With the given shape and scale parameters of the log-normal distribution, as well as the constant location parameter set to 0, the depicted distributions of the optimal model can be accurately reproduced.

6.3.3 Comparing deterministic and probabilistic with experimental data

This optimal model is compared to the traditional deterministic approach and experimental data from the FDA RBP benchmark study [8] in Figure 6.5. The RIH with respect to

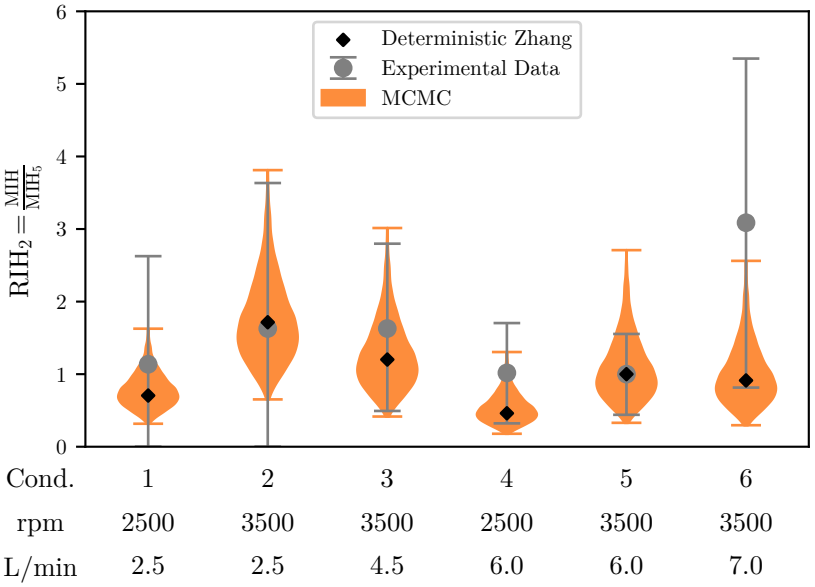

Figure 6.5: This figure compares the deterministic and probabilistic approaches with experimental data from the FDA RBP benchmark study under six different operating conditions [8]. The deterministic model [106] is represented by black diamonds, the probabilistic model by orange violin plots, and the experimental data's mean ± SD is shown with grey error bars.

Condition 5 is plotted for boundary conditions 1–6. The comparison between the mean values of the experimental data from Ponnaluri *et al.* [8] (grey error bars) and the deterministic data of Zhang *et al.* [106] (black diamonds) shows that all deterministically calculated RIH values lie within the experimental measurements (mean ± SD). The novel probabilistic method (orange violin plots) shows a similar pattern regarding the median values of the distribution. In addition to the median values, the distribution of possible hemolysis values can be observed through the propagated uncertainty of the underlying experimental data. It is also evident that Conditions 1 and 4, which correspond to a rotational speed of 2500 rpm, have a smaller distribution width compared to the other conditions with a rotational speed of 3500 rpm. In addition to the MCMC analysis based on published experimental data from Zhang *et al.* [106], which primarily consisted of triplicate measurements per operating point, a supplementary analysis was performed using artificially generated datasets with varying numbers of repetitions per operating point. The results are presented in Figure 6.6.
Panels A–C depict three synthetic datasets comprising 80 operating points spanning shear rates from 50–320 Pa and exposure times from 0.04–1.5 s, with 3, 5 and 10 repetitions per operating point, respectively. Panels D–F show the corresponding posterior distributions of the parameters α, β, and σ, assuming a fixed value of $C = 1.228 \times 10^{-5}$, for each of the training datasets shown in Panels A–C. An increase in the number of repetitions per operating

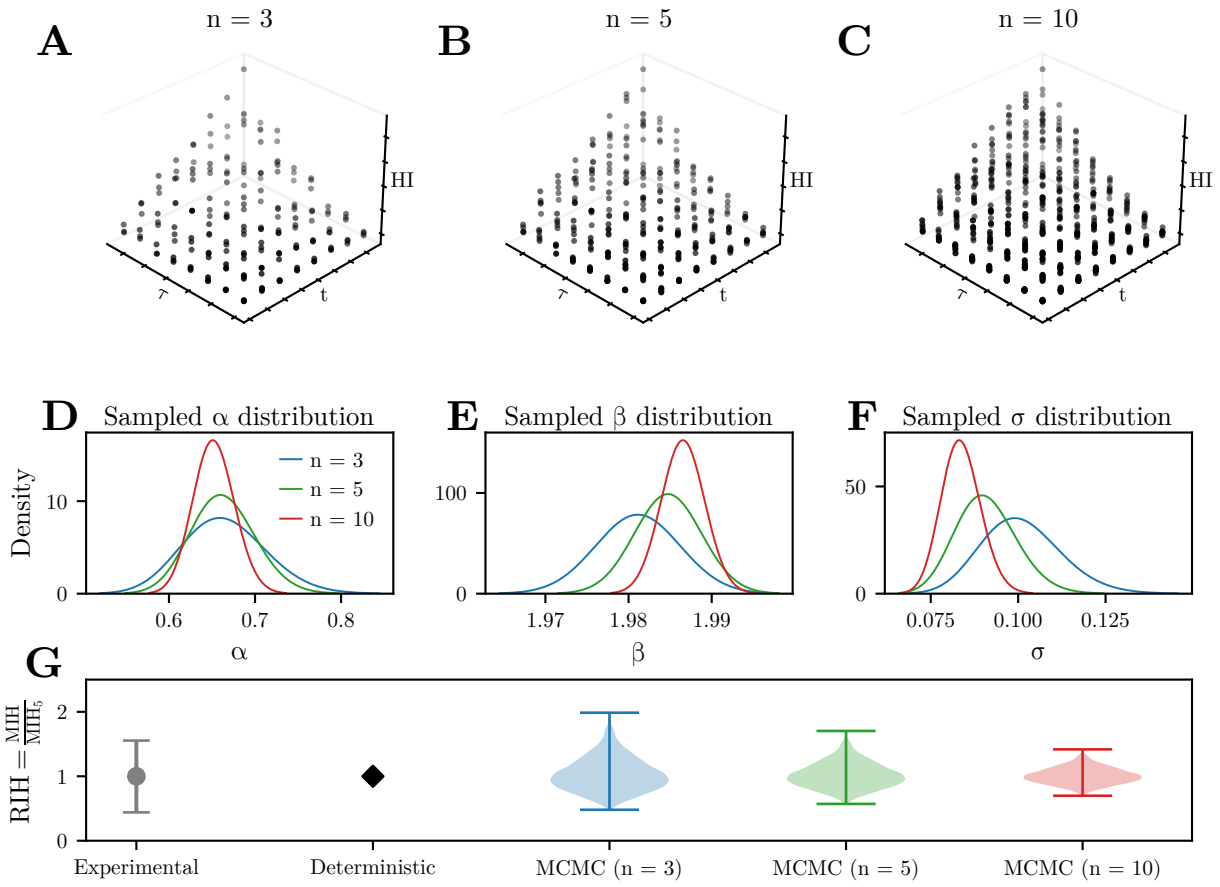

Figure 6.6: Panels A–C show synthetic data generated from the power law model by Zhang *et al.* [106] across 80 operating points (shear stresses: 50–320 Pa, exposure times: 0.04–1.5 s), with 3, 5 and 10 samples per condition. Samples were drawn from truncated normal distributions (1^{st}–99^{th} percentile), centered on the power law model HI values, with variances linearly scaled from 1×10^{-4} to 1×10^{-1} based on the lowest and highest HI values predicted by the model. Panels D–F show posterior distributions of α, β, and σ for each dataset ($n = 3, 5, 10$). Panel G shows the corresponding RIH predictions for Condition 5, compared with experimental data from Ponnaluri *et al.* [8] (gray error bars) and deterministic results from Zhang *et al.* [106] (black diamonds).

point leads to a noticeable reduction in the variance of the posterior distributions, indicating increased confidence in the parameter estimates. This effect is further illustrated in the lower row of the figure (Panel G), which shows the resulting probabilistic predictions of RIH for condition 5 based on each training dataset. The distributions are color-coded according to the number of repetitions in the training data, with blue representing 3, green 5, and red 10 repetitions. The increased confidence in the parameter estimates is propagated through the simulation, resulting in a corresponding decrease in the uncertainty of the predicted RIH as the number of repetitions increases. To contextualize these results, the predicted distributions are shown alongside the experimental data of Ponnaluri *et al.* [8] (gray error bar) and the deterministic prediction from Zhang *et al.* [106] (black diamond) for the same condition.
This analysis closely relates to the experimental investigation presented in Chapter 5, where Figure 5.4 demonstrated, that the variance in experimental data decreases with an increasing

number of measurements at each operating point. Interestingly, the synthetic MCMC results presented in Figure 6.6E exhibit a similar trend: the 95 % confidence intervals for predicted hemolysis becomes increasingly narrow with 3, 5 and 10 repetitions. To enable a direct comparison between the synthetic and experimental results, the 95 % confidence interval ranges are normalized by their respective maximum values, yielding relative CI ranges. For the repetition numbers 3, 5 and 10 both datasets show a consistent decrease across repetitions, with values of [1.00, 0.758 and 0.487] for the synthetic data and [1.00, 0.770 and 0.549] for the experimental data. This indicates a very similar trend in the relative reduction of uncertainty. A Pearson correlation coefficient of $r = 0.999$ ($R^2 = 0.998$) confirms an almost perfect linear relationship between the two normalized trends, further supporting the consistency between synthetic and real-world observations.

6.4 Discussion

The aim of this chapter was to develop a modeling approach that incorporates experimental uncertainty into hemolysis predictions, addressing a key limitation of conventional power law-based models. By first examining the sensitivity of the power law fitting process and then introducing a probabilistic framework based on Bayesian inference and reduced-order modeling, this chapter demonstrated how uncertainty in the underlying data can be systematically propagated into model outcomes. Three central findings emerged from this analysis:

- The fitting process of uncertain experimental hemolysis data exhibits a flat global optimum, resulting in substantially different hemolysis predictions when allowing a small margin of fitting error.
- The strong correlation between C and β prevents the MCMC approach from converging on a uniquely identifiable set of parameter distributions.
- Considering the variance from the underlying experiments enhances the robustness of hemolysis model predictions and simplifies the comparison between experimental and simulated data.

To explore the first finding, a grid search was conducted by systematically varying the three power law fitting parameters (C, α, β). For each parameter combination, the SSE between predicted and measured hemolysis values was calculated to assess the quality of fit. The resulting SSE landscape (*see* Figure 6.2) revealed a flat global optimum for the fitted parameters. In other words, small deviations in SSE can lead to substantially different fitting parameter combinations. One reason for this flat minimum was identified to be the compensatory relationship between C and β, preserving the model output while one parameter is offset by opposite changes in the other. Importantly, this is a fundamental property of multi-parameter power-law models, where fitting parameters can compensate for one another based on the data they are fitted on. However, just based on the mathematical formulation of equation (2.3) it is not clear that this compensatory effect in the fitting process will take place and which of the fitting parameters will contribute to this compensation. Beside the compensatory effect, the flatness of the optimum can be further amplified by three factors:

the relative contribution of a fitting parameter to the model output, the range over which the associated independent variable is varied, and the inherent variance in the experimental data. Figure 6.2A demonstrates that the SSE is highly responsive to changes in β compared to α. Even a slight increase in β leads to a substantial change in the SSE, whereas similar modifications in α produce only minimal variations. The reason for that is that the parameter β exerts a strong contribution on the model output ($\beta \approx 2$) and is fitted over a wide range (0–400 Pa) and is therefore more tightly constrained, leading to a steeper, more well-defined minimum in the SSE objective function landscape. Conversely, α's contribution is modest ($\alpha \approx 0.6$) and is fitted over a narrow range of conditions (0–1.5 s) and therefore the impact of small parameter variations on the model output becomes less noticeable. Furthermore, the variance in the underlying experimental data for each operating condition amplifies this effect because higher measurement noise means that subtle differences in the objective function are masked, allowing a broader set of parameter combinations to yield nearly equivalent SSE values. Together, these factors explain why, under the influence of contribution, range, and data variance, the objective function remains flat over a wide region of the parameter space. A key implication of this flat SSE minimum is that even a small deviation from the global optimal SSE can lead to substantially different hemolysis predictions for the same device and operating condition. As illustrated in Figure 6.2B–C, hemolysis outcomes may vary by up to 60 % when the SSE is allowed to increase from 3.50 to merely 3.52 (a narrow 0.5 % margin). This large variation highlights the importance of quantifying parameter uncertainty because relying solely on a best-fit parameter set can mask the possible range of plausible outcomes. Instead, acknowledging and reporting the plausible range of solutions is essential for accurately assessing hemolysis risk and comparing experimental with simulated data.
Craven *et al.* [110] have previously found that there are infinitely many different parameter combinations for a single operating point that lead to the same hemolysis value. This chapter extends that analysis, demonstrating that the proposed power law function might be an unsuitable choice for finding a robust representation of Couette shearing device experimental data with variance over a wider range of operating points. While there are other options that could potentially resolve this flatness issue, such as logistic or additive forms or imposing mechanistic constraints on parameters to reduce the likelihood of degenerate fits, this chapter introduced a Bayesian MCMC method to incorporate uncertainties from the underlying experimental data of hemolysis models into the prediction of hemolysis outcomes. By integrating the experimental variance directly into the parameter fitting process, this method aimed not only to address the flatness of the objective function but also to demonstrate how uncertainty propagates through the model. Notably, this integration of the variance from underlying hemolysis experiments into model parameter estimation has not previously been reported. However, MCMC sampling revealed a persistent and strong correlation between parameters C and β, consistent with the pattern observed in the SSE analysis (*see* Figures 6.2 and 6.3). This correlation leads to a nearly flat objective function landscape where a wide range of parameter combinations yield similar fits, thereby preventing convergence of the MCMC method to a uniquely identifiable set of parameter distributions. To address this challenge, the dimensionality of the multi-parameter power-law model was reduced by fixing C and sampling only α, β, and σ. Examining several fixed values of C, showed that the residual error, quantified by σ as the difference between the model predictions and the measured hemolysis data, reached an optimum at a particular value of C. Then the employed Bayesian

optimization was used to identify the C value that minimizes the median of σ. Minimizing σ in this context effectively minimizes the residuals and therefore enhances the precision of the parameter estimates for α and β. Consequently, the resulting probability distributions for these parameters more accurately reflect the true variance of the experimental data, leading to a more reliable model. This approach differs substantially from previous literature [150], which randomly samples probability distributions for α, β, and C from parameter intervals derived from literature values of multiple models.

The ability of the proposed probabilistic approach to faithfully reflect only the variance present in the experimental data was further validated using synthetic training data that contained no variability. Specifically, the model was trained on a dataset consisting of hemolysis values that lie exactly on the surface defined by the original model parameters of Zhang *et al.* [106], without any added noise or experimental variation. As shown in Appendix C, Figure C.2, this resulted in point estimates for the parameters α and β, with no spread in their posterior distributions. This confirms that the probabilistic representation of hemolysis as a distribution in this chapter arises solely from the variability in the underlying experimental data, and not from the inference process itself.

In a second analysis using synthetic training data, the influence of varying numbers of experimental repetitions per operating point on the resulting parameter estimates was investigated. As discussed in the theoretical foundations of Bayesian inference in Chapter 2.3.3, increasing the number of observations while maintaining the same overall variance structure leads to narrower posterior distributions and, consequently, greater confidence in the estimated parameters. This effect was clearly demonstrated in the hemolysis predictions of the probabilistic model shown in Figure 6.6E, where increasing the number of repetitions in the training data reduced the uncertainty in model output.

What is particularly noteworthy in this context is the striking agreement between this synthetically constructed trend and the experimental trend shown in Figure 5.4E of Chapter 5. The normalized 95 % confidence interval ranges for both datasets exhibit nearly identical patterns of reduction as the number of repetitions increases from 3 to 5 to 10, resulting in an almost perfect linear correlation. This agreement supports the notion that the insights gained in Chapter 5 from controlled analyses at individual operating points are transferable to the broader context of uncertainty propagation across multiple operating conditions, as encountered in power-law-based hemolysis models applied to full-device predictions. Moreover, it demonstrates that the model is sensitive to both the quality and quantity of input data in a way that reflects experimental conditions, highlighting its utility for uncertainty quantification in hemolysis prediction.

These findings highlight the current underrepresentation of experimental replicates per operating point, especially given the inherent uncertainty in hemolysis measurements. This should be kept in mind when interpreting the results of the present chapter, as the MCMC approach captures the uncertainty associated with parameter estimation based the published dataset of Zhang *et al.* [106], which relies mostly on triplicate measurements. The confidence in these estimates could be substantially improved if the experimental database were expanded through studies with higher repetition numbers.

Nevertheless, to demonstrate the practical applicability of the method, the Bayesian hemolysis model was used as a post-processing step in a CFD simulation of the FDA RBP benchmark geometry [92]. The resulting numerical hemolysis predictions were compared to experimen-

tal data [8] across six operating conditions. This comparison further highlights the benefit of incorporating experimental variance into the model fitting process, as the probabilistic representation of hemolysis predictions aligns well with the variability observed in the experimental measurements. The MCMC approach strengthens the robustness of relative hemolysis predictions by considering potential variance factors such as inter- or intra-variability of blood samples from the underlying experiments, which can also be sources of variance in the experimental investigation of MCS devices. Furthermore, the well-known strong sensitivity of predicted hemolysis values to variations in model parameters [110, 150, 202] is directly observable in the distribution representation of the probabilistic model, facilitating a more comprehensive interpretation of the model outputs. In the comparison to the FDA pump data [8] it is important to note that the modal (most probable) value of the probabilistic results closely matches the deterministic solution. This outcome reflects the fact that the deterministic solution, corresponding to the global minimum of the SSE, is identified by the probabilistic approach as the most probable solution as well. However, the probabilistic framework offers a substantial advantage by quantifying and propagating uncertainties from the underlying data. As a result, while the modal value aligns with the deterministic prediction, the probabilistic approach also captures the range of plausible hemolysis outcomes arising from uncertainties of the underlying experimental data. But not all operating conditions could be predicted equally well. In particular, Condition 6 proved challenging for the numerical model as both the deterministic solution and the probabilistic solution could not predict the mean of the experimental results. That this Condition 6 is particularly difficult to predict is also evident in the comparison of the numerical results of the FDA round robin study as none off the numerical models could reliably predict this Condition 6 [8]. This suggests that certain flow phenomena (e.g., transitional flow regimes, turbulence, flow separation) may not be fully captured by current simulations or that the experimental data themselves might be influenced by unaccounted factors. In this context it is important to note that the experimental FDA pump data may be influenced by numerous other factors (e.g. donor differences, blood handling, uncertainty in operating conditions, ...) that contribute to the measured variance in hemolysis values, many of which cannot be fully accounted for in a numerical model. However, variability is not limited to the experimental side. There are also sources of uncertainty within the numerical framework that were not within the scope of this chapter. Specifically, the probabilistic approach focuses on uncertainty in the fitted parameters of the calibrated power-law model and does not account for uncertainties related to operating conditions, geometric fidelity (including potential manufacturing deviations or wear), or numerical modeling assumptions (e.g., mesh resolution, turbulence modeling choices, and rheological models for blood). Each of these factors can influence the predictive accuracy of hemolysis models in practical RBP simulations and contribute to overall uncertainty. Accordingly, while this chapter enhances the understanding of fitting parameter-related variability, it should be considered as one component within a broader uncertainty quantification framework. Additionally, it is important to note that the raw data from Zhang *et al.* [106] examined bovine blood, while the FDA pump experiments [8] investigated porcine blood. Consequently, the comparison to the FDA pump data serves only as an illustrative example of how the probabilistic hemolysis model simplifies the comparison to experimental data. Nonetheless, despite the different donors, a relative comparison between simulation and experiment is still possible.

However, the problem of accurately predicting absolute hemolysis values is not solved by the probabilistic model. Absolute hemolysis values of the probabilistic approach are within the same order of magnitude as those predicted by classical power law models. This issue does not appear to be due to the variance in the underlying experimental data but, as discussed in the literature, seems to result from more fundamental modeling issues. In particular, the application of calibration data from simplified experimental setups to more complex flow environments introduces additional uncertainties, as controlled shear exposure experiments may not fully capture the range and type of stresses and range of residence times encountered in realistic RBP flows. Moreover, the inability to resolve all turbulence scales and individual RBC dynamics in numerical simulations further limits predictive accuracy, as key flow structures that influence hemolysis may not be fully captured [117].

6.5 Conclusion

This chapter highlights the fundamental mathematical limitations in fitting classical multi-parameter power-law models to uncertain experimental data and demonstrates how the inherent variability in hemolysis experiments can be effectively incorporated into numerical models. By capturing parameter uncertainty through MCMC, the robustness of hemolysis predictions is substantially improved. The current gold standard of relative comparisons is strengthened by accounting for the variance in underlying experiments, providing a more rigorous basis for aligning simulated hemolysis outcomes with *in-vitro* measurements. The proposed probabilistic framework supports seamless integration of new experimental data, across varying shear conditions, donor species, and experimental designs, without requiring retraining from scratch. As shown, increasing the number of observations narrows the posterior distributions, thereby improving confidence in parameter estimates. This makes the approach both flexible and scalable, with uncertainty naturally diminishing as the experimental database expands. With this, the approach has the potential to result in a comprehensive model that incorporates all the experimental knowledge of the scientific community, setting a new standard of predictive accuracy in hemolysis modeling.

7 Summary & Outlook

The overarching aim of this dissertation was to develop a computationally efficient, uncertainty-aware *in-silico* framework for predicting shear-induced hemolysis in RBPs. This goal was driven by the clinical and regulatory need for robust numerical modeling approaches that enhance device safety, inform clinical decisions, and support the assessment of RBPs across all clinically relevant operating points, with ECMO serving as the representative application throughout this dissertation. To achieve this, the dissertation was structured around four central research questions, each addressed in a dedicated chapter.

Chapter 3 – Efficiently computing a wide range of pump operating points

- *Research question:* How can fluid variables across a wide range of RBP operating points be computed efficiently?
- *Key finding:* Given the correct settings and validation procedures, NIPCE enables efficient and accurate predictions of flow variables across a wide range of operating points.

To simulate the full clinically relevant range of flow rates (0–7 L/min) and pressure heads (0–400 mmHg) in ECMO support, Chapter 3 introduced a ROM framework based on NIPCE. This method was train on and validated against high-fidelity CFD simulations and shown to accurately reproduce key fluid dynamic variables such as pressure head, velocity, and axial force. By selecting suitable polynomial orders and a sufficient number of training points, the ROM achieved substantial computational savings without compromising prediction accuracy. Together with an analysis of input variable continuity, these insights were formulated into practical guidelines that support the integration of the ROM into hemocompatibility assessments and device development pipelines. With this, the approach enables efficient parameter studies and serves as a scalable foundation for uncertainty quantification, forming the computational basis for all subsequent numerical analyses in this dissertation.

Chapter 4 – Comparing *in-silico* hemolysis predictions and *in-vivo* measurements

- *Research question:* How accurately do *in-silico* hemolysis models capture clinical *in-vivo* hemolysis data?
- *Key finding:* Strong correlation between predicted hemolysis and cohort-averaged pfHb confirmed model validity, while the underlying data revealed substantial variability at the individual patient level.

Building on the ROM framework established in Chapter 3, Chapter 4 evaluated the clinical relevance of stress-based hemolysis models. While prior validations relied primarily on controlled *in-vitro* experiments, this chapter utilized a high-resolution *in-vivo* dataset from 580 patients undergoing VV ECMO treatment. The analysis involved extensive processing of clinical data, including normalization, filtering, and mapping of RBP operating conditions to measured *in-vivo* hemolysis values, to enable a meaningful comparison. The comparison revealed on average a strong association between numerically predicted hemolysis and measured *in-vivo* pfHb, with ECMO circuit resistance or RBP pressure head emerging as a key contributor to hemolysis. By identifying this strong correlation, the findings not only strengthen the credibility of numerical stress-based hemolysis modeling in real-world clinical scenarios but also enable actionable insights for circuit design and patient management. Importantly, the results also revealed considerable inter- and intra-patient variability, underlining the need to account for uncertainty in both model inputs and clinical outcomes.

Chapter 5 – Quantifying variability in *in-vitro* hemolysis measurements

- *Research question:* How can uncertainties in experimental hemolysis data be quantified?
- *Key finding:* Inherent intra-donor variability in hemolysis measurements under fixed shear conditions is substantial and becomes more clearly characterized through increased measurement repetition.

Building on the clinical variability observed in Chapter 4, Chapter 5 examined whether similar uncertainty exists in the underlying experimental data of macroscopic stress-based hemolysis models. These models are typically derived from sparse datasets with limited replicates per condition, raising concerns about their representativeness. To investigate this, hemolysis was measured 100 times under a fixed shear condition using a rheometer-based shearing device designed for consistency and reproducibility. Data were collected across five donors to assess both intra- and inter-donor variability. Under these rigorously controlled conditions the results revealed that intra-donor variability was markedly higher than inter-donor variability, challenging the common assumption that donor differences and blood handling errors are the primary sources of measurement variation. Crucially, hemolysis variability appeared high in standard triplicate measurements due to sampling noise, but larger sample sizes revealed a more stable and representative distribution. These findings call into question the robustness of existing hemolysis models that dont incorporate this inherent variability and point to the need for more comprehensive datasets and model formulations that explicitly account for experimental uncertainty.

Chapter 6 – Uncertainty-aware hemolysis modeling

- *Research question:* How can experimental uncertainties be incorporated into uncertainty-aware *in-silico* hemolysis models?
- *Key finding:* MCMC enables uncertainty-aware predictions by propagating experimental variability to model outputs, highlighting the need for more experimental repetitions.

Building on the quantified experimental variability presented in Chapter 5, Chapter 6 addressed two key limitations of current hemolysis modeling approaches. First, conventional

models typically rely on point estimates derived from sparse and variable experimental data. This prevents the propagation of input uncertainty into the model output, limiting the ability to assess prediction confidence. To overcome this, the Bayesian parameter estimation framework MCMC was introduced. By incorporating experimental variability directly into the calibration process of fitting parameters, the method provides probability distributions for fitting parameters and enables uncertainty-aware predictions. Applied to the FDA benchmark RBP, the probabilistic model preserved the central trend of deterministic predictions but additionally quantified the range of plausible outcomes, supporting better interpretation in clinical and regulatory contexts. Secondly, the chapter highlighted a structural problem in multi-parameter power law models. Fitting such models to uncertain data can result in compensatory effects among parameters, creating flat regions in the objective function landscape. This means that small changes in fitting error may correspond to large differences in parameter values and predictions. The Bayesian approach addresses this problem by reflecting these uncertainties in the resulting parameter distributions. This leads to a more transparent and reliable model calibration that better reflects the uncertainty in the underlying data and model structure. Additionally it was shown, that the current model uncertainty can be decreased by incorporating more replicate measurements in the underlying experimental data.

To understand the full impact and future potential of these individual contributions, it is necessary to place the chapter results within the broader clinical, regulatory, and modeling context in which they will be applied:

Clinical

Chapter 4 provided the first *in-vivo* validation of stress-based hemolysis models across a wide range of clinically relevant operating points. This validation confirmed that such models can reliably capture average blood damage trends and can be used to estimate cohort-level outcomes under different therapy settings. Based on this validation, Chapter 4 already provided reliable recommendations for clinical practice, establishing the practical relevance of stress-based hemolysis models beyond theoretical studies.
At the same time, the clinical data revealed substantial inter- and intra-patient variability. Hemolysis responses differed substantially not only between patients but also within repeated measurements for the same patient under identical conditions. These findings point to a fundamental limitation of current stress-based models, which rely solely on mechanical input parameters such as shear stress and exposure time. By assuming that all RBCs respond uniformly to a given mechanical environment, these models fail to capture the biological heterogeneity observed in practice. While the probabilistic framework developed in Chapter 6 was designed to incorporate variability into the modeling process, it does not support the prediction of patient-specific responses, as these are likely influenced by additional biological factors affecting hemolysis susceptibility, such as RBC age distribution, membrane properties, inflammation, and oxidative stress.
To enhance the clinical applicability of numerical hemolysis modeling and support the shift toward patient-specific care, future work should focus on incorporating biomarkers that reflect inherent biological variability. This would improve predictive accuracy and further establish probabilistic approaches as the preferred modeling framework.

Regulatory

Chapter 6 introduced a probabilistic framework that, for the first time, quantifies uncertainty in hemolysis modeling based on the variability of the underlying experimental calibration data. This marks a substantial step forward, as it enables a more transparent and robust prediction of hemolysis outcomes and enables direct comparison to experimental measurements. In doing so, the framework aligns with growing regulatory expectations for model credibility, which emphasize the importance of VVUQ in computational models used for medical device evaluation.

While this development improves both transparency and trust in numerical predictions, it addresses only a subset of the broader uncertainty landscape. The current approach captures variability originating from experimental data but does not account for other critical sources of uncertainty. These include variations in boundary conditions, operating parameters, device geometry, numerical assumptions, or other modeling assumptions like the complex rheological behavior of blood. Each of these factors can affect model outcomes and may introduce additional layers of uncertainty that remain unquantified.

To meet regulatory expectations for comprehensive VVUQ, the current approach must therefore be extended to capture additional sources of uncertainty beyond experimental calibration. Nonetheless, this dissertation establishes an important foundation by introducing probabilistic modeling for uncertainty quantification and demonstrating the value of numerical hemolysis models through clinical validation. Building on this foundation, future efforts should focus on developing integrated VVUQ strategies that account for the full range of uncertainties relevant to regulatory evaluation.

Modeling

The modeling framework developed in this dissertation offers a structured approach to propagate uncertainty from experimental data through hemolysis simulations, overcoming a key limitation of previous deterministic models, which lacked the means to represent experimental uncertainty. This advance is made possible by a pipeline that combines Bayesian parameter estimation via MCMC sampling (Chapter 6) with the ROM introduced in Chapter 3, enabling efficient and scalable uncertainty quantification. Although the present work focuses on stress-based models due to their computational simplicity and broad adoption, the methodology itself is more general. Since it is based on data-driven parameter estimation, the framework can also be applied to other macroscopic hemolysis models, including strain-based or dissipation-based formulations, making it a versatile tool for future modeling efforts.

However, the reliability of uncertainty predictions depends heavily on the quality of the underlying experimental data. Chapter 5 showed that standard protocols with only a few replicates per condition fail to capture the true variability in hemolysis measurements. Intra-donor fluctuations exceeded inter-donor differences, revealing substantial inherent variability even under controlled conditions. This was further supported in Chapter 6, where synthetic data analysis demonstrated that increasing the number of replicates substantially reduced posterior uncertainty and improved parameter confidence. These findings point to a key limitation of the current probabilistic modeling framework: it reflects only the variability contained in the data on which it is trained. A second, equally important limitation lies in

the hemolysis models themselves. Stress-based formulations, while computationally efficient and useful for capturing relative trends, are fundamentally limited in their ability to predict absolute levels of hemolysis. Their simplified mechanics do not account for the complex biological and mechanical interactions that influence RBC damage. As a result, even a well-calibrated probabilistic model cannot overcome the structural limitations of the underlying hemolysis formulation.

Addressing these limitations will require improvements in both modeling and data generation. From a modeling perspective, future work should aim to incorporate more realistic mechanical descriptions of RBC damage. Current stress-based models are limited by their reliance on simplified scalar shear stress as an input parameter, which does not reflect how forces are actually transmitted to or experienced by the cells. As a result, they cannot capture important effects such as deformation, orientation, or the accumulation of mechanical load over time. Alternative approaches that focus on how cells respond to deformation, such as strain-based models, offer a promising way to address these limitations. However, more research is needed to further strengthen the mechanical foundations of hemolysis modeling in order to eventually predict absolute hemolysis values. In this context, the findings of this dissertation also highlight the importance of including additional biological factors to account for the substantial inter- and intra-patient variability observed in both experimental and clinical data. From a data perspective, experimental protocols could benefit from shifting focus toward fewer but carefully selected operating points with a greater number of replicates. Current approaches often prioritize broad coverage of operating conditions, typically with only a small number of replicates per point. However, this strategy falls short in capturing the inherent variability of hemolysis measurements. Concentrating limited experimental resources on a smaller set of representative conditions and increasing the number of replicates for each would enhance statistical reliability and support more robust model calibration.

A longer-term goal should be the creation of a shared, high-quality experimental database that aggregates hemolysis measurements across laboratories, devices, and donor populations. In addition to standard flow and hemolysis outcome data, this database should include detailed information on blood composition, donor health status, and relevant biochemical markers. By linking the probabilistic framework developed in this dissertation to such a continuously growing dataset, model parameters could be iteratively updated as new data become available. This would allow for ongoing uncertainty refinement without the need to retrain the model from scratch, improving both efficiency and predictive accuracy. Beyond improving calibration, this kind of well-structured database would also facilitate the identification of biological factors that contribute to hemolytic variability. Systematic analysis of this information could reveal correlations between hemolytic susceptibility and specific biochemical or cellular markers, which could provide a crucial foundation for future hemolysis models that aim to move beyond purely mechanical descriptions and incorporate patient-specific biological factors into their predictions.

Taken together, these developments in modeling would not only improve the reliability of current hemolysis predictions but also open the door to more advanced modeling applications. One particularly promising direction is the use of *in-silico* clinical trials. These trials aim to simulate large virtual patient cohorts in order to evaluate the safety and performance of blood-contacting devices across a wide range of physiological conditions. As regulatory expectations continue to move toward evidence that accounts for patient variability, such simulations can

provide an efficient and informative complement to traditional bench, animal, or human studies. They enable systematic testing under diverse and controlled scenarios, support the identification of patient-specific risk factors, and may ultimately help guide clinical decision-making before physical trials are conducted. The combination of reduced-order modeling (Chapter 3), clinical validation (Chapter 4), and probabilistic prediction (Chapter 6) developed in this dissertation forms a strong foundation for such future applications. Nevertheless, realizing this potential will require further progress in modeling, including the ability to make patient-specific predictions, a more comprehensive treatment of all relevant sources of uncertainty, and the validation of model outputs against absolute hemolysis values. Each of these steps represents a substantial challenge for future research.

In summary, this dissertation represents a first step toward a new generation of hemolysis models. It moves beyond simplified, deterministic approaches calibrated on sparse data and introduces a robust, uncertainty-aware framework that better reflects the variability of biological systems and the expectations of modern clinical and regulatory practice. The methods, findings, and limitations discussed here provide a solid basis for future work aimed at improving the safety, effectiveness, and credibility of RBPs and other blood-contacting medical devices through the use of computational models.

In line with open science principles and to support future research, code and data related to the main methodological developments have been made publicly available. For Chapter 3, the raw data and NIPCE surrogate modeling Python implementation can be found at https://zenodo.org/record/8070576. For Chapter 6, a Python implementation of the MCMC model and synthetic test data are available at https://zenodo.org/records/12820264.

A Appendix to Chapter 4

In this Appendix A, additional information about the clinical dataset is provided, along with further details on the linking process between *in-vivo* measurements and *in-silico* simulations, to complement and support the findings presented in Chapter 4, which is based on [161].

A.1 Additional analysis of clinical in-vivo data

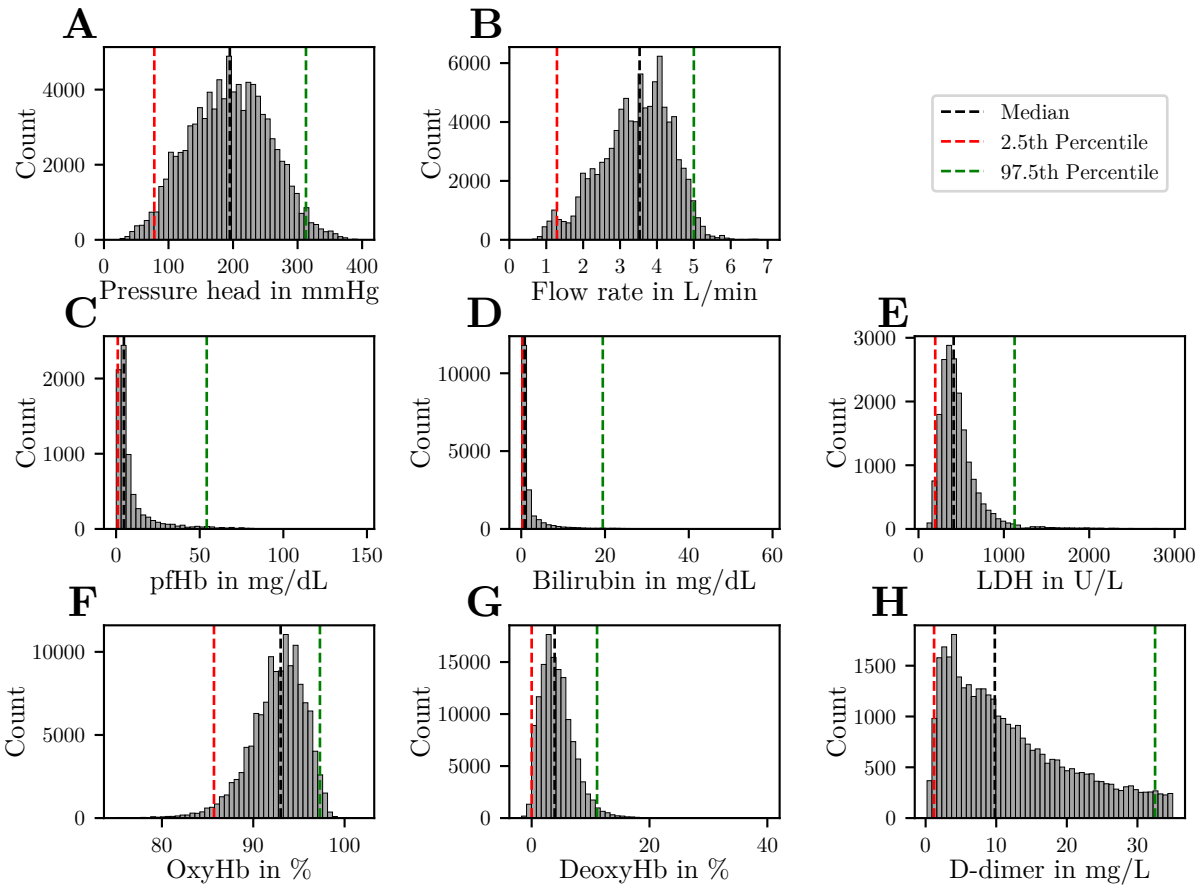

Figure A.1: Histograms of the RBP parameters pressure head, flow rate (A,B) as well as the blood parameters plasma free hemoglobin (pfhb), bilirubin, lactate dehydrogenase (LDH), oxyhemoglobin (oxyHb) and deoxyhemoglobin (deoxyHb) and D-dimer (C-H). Median (black), 2.5$^{\text{th}}$ percentile (red) and 97.5$^{\text{th}}$ percentile (green) are visualized by vertical lines.

Figure A.1 illustrates the distributions of the key RBP and blood parameters analyzed in Chapter 4 across the entire patient cohort. These histograms provide insight into the variability and range of the clinical dataset used in the analysis. Panels A–B show the distributions of the RBP parameters, pressure head and flow rate. The pressure head (Panel A) follows a roughly symmetric distribution, while the flow rate (Panel B) reveals a pronounced accumulation at lower values, likely reflecting support settings during weaning phases. Panels C–H display histograms of the main blood markers used in Chapter 4. Plasma free hemoglobin (pfHb, Panel C), bilirubin (Panel D), lactate dehydrogenase (LDH, Panel E), and deoxyhemoglobin (deoxyHb, Panel G) are all positively skewed, indicating that elevated concentrations occur infrequently but can reach high values when present. Oxyhemoglobin (oxyHb, Panel F) is negatively skewed, reflecting consistently high saturation with occasional dips. The D-dimer distribution (Panel H) is wide, with a substantial number of both low

Table A.1: Description of the cohort in the categories: demographics, outcome and multiorgan failure with additional baseline characteristics describing the first 48 hours of therapy for the whole cohort as well as the sub-cohorts survivor and non-survivor as median (25^{th} confidence interval ,75^{th} confidence interval). The asterisk (*) indicates that for body mass index (BMI) only 45.5 % of the cohort could be evaluated due to missing data. Similarly, the double asterisk (**) indicates a missing rate of 80 % for the sequential organ failure assessment (SOFA) score. Additional abbreviations used in the table can be found in the Abbreviations list at the beginning of the dissertation.

	Parameter	**All**	**Survivor**	**Non survivor**
Demographics	Age (years)	56.7 (46.5,64.18)	54.2 (46.0,61.9)	58.9 (47.6,66.1)
	Sex (% female)	34.8	36.8	32.7
	BMI (kg/m^2)*	28.0 (24.7,34.8)	29.4 (25.3,35.5)	27.2 (24.2,32.7)
Outcome	Exitus (% survivor)	52.1	100	0
	ICU Length of Stay (days)	23 (14,41)	27 (18,43)	18 (8,36)
	ECMO duration (days)	18 (14,21)	18 (14,21)	18 (14,21)
Multiorgan failure	SAPS	39.0 (32.5,46.5)	37.0 (30.25,43.0)	42.0 (35.5,50.0)
	SOFA**	11 (10,13.5)	11 (10,13)	12 (11,16)
	TISS	19 (15,22)	19 (15,21.5)	19 (15,23)
Operating point data (baseline 48 h)	Pressure head (mmHg)	212 (172,246)	208 (170,246)	215 (175,246)
	Flow rate (L/min)	3.2 (3,3.93)	3.3 (3,3.9)	3.2 (3,3.9)
	P_{art} (mmHg)	142 (119,163)	143 (121,163)	142 (117,163)
	P_{ven} (mmHg)	-45 (-56,-31)	-46 (-57,-31)	-42 (-56,-29)
	P_{int} (mmHg)	165 (140,188)	166 (141,189)	164 (137,187)
	O_2-Flow (L/min)	4.8 (3.8,6.0)	4.8 (3.8,6.0)	4.8 (3.9,6.0)
Laboratory findings (baseline 48 h)	pfHb (mg/dL)	4.9 (2.9-8.61)	4.6 (2.9-7.7)	5.1 (2.9-10.5)
	LDH (L)	475 (339,659)	459 (319,641)	486 (357,699)
	Bilirubin (mg/dL)	0.8 (0.5,1.7)	0.7 (0.5,1.4)	1.0 (0.5,1.9)
	Haptoglobin (g/L)	2.3 (1.4,3.1)	2.4 (1.6, 3.3)	2.1 (1.2, 2.9)
	pH	7.40 (7.35,7.43)	7.41 (7.36,7.43)	7.38 (7.34,7.43)
	PO_2 (mmHg)	78.1 (72.2,85.4)	78.1 (72.8,84.5)	78.1 (72.0,86.0)
	PCO_2 (mmHg)	44.5 (40.1,52.0)	44.0 (40.0,51.4)	45.0 (40.2,52.9)
	CRP (mg/L)	19.0 (11.4,26.9)	19.0 (9.1,26.6)	18.9 (12.8,27.0)
	PCT (ng/mL)	2.2 (0.8,9.8)	2.0 (0.7,8.7)	2.8 (1.0,11.7)
	Fibrinogen (mg/dL)	448 (340,599)	461 (362,606)	443 (318,595)
	D-dimer (mg/L)	5.0 (2.9,10.0)	5.0 (2.0,9.0)	6.5 (3.0,11.0)
	Lymphocytes (1/nL)	0.8 (0.5,1.2)	0.8 (0.6,1.2)	0.7 (0.4,1.0)
	Neutrophils (1/nL)	10.8 (7.0,14.9)	10.7 (6.8,15.4)	11.0 (7.6,14.6)

and elevated measurements, suggesting heterogeneous activation of coagulation pathways. In each panel, the vertical lines mark the median (black), as well as the 2.5^{th} (red) and 97.5^{th}

(green) percentiles, providing visual reference points for understanding the spread and central tendency of each parameter.
To provide a comprehensive overview of the study population and contextualize the clinical data used in Chapter 4, Table A.1 presents a detailed summary of the cohort's baseline characteristics. It includes demographic information, outcome parameters, multiorgan failure

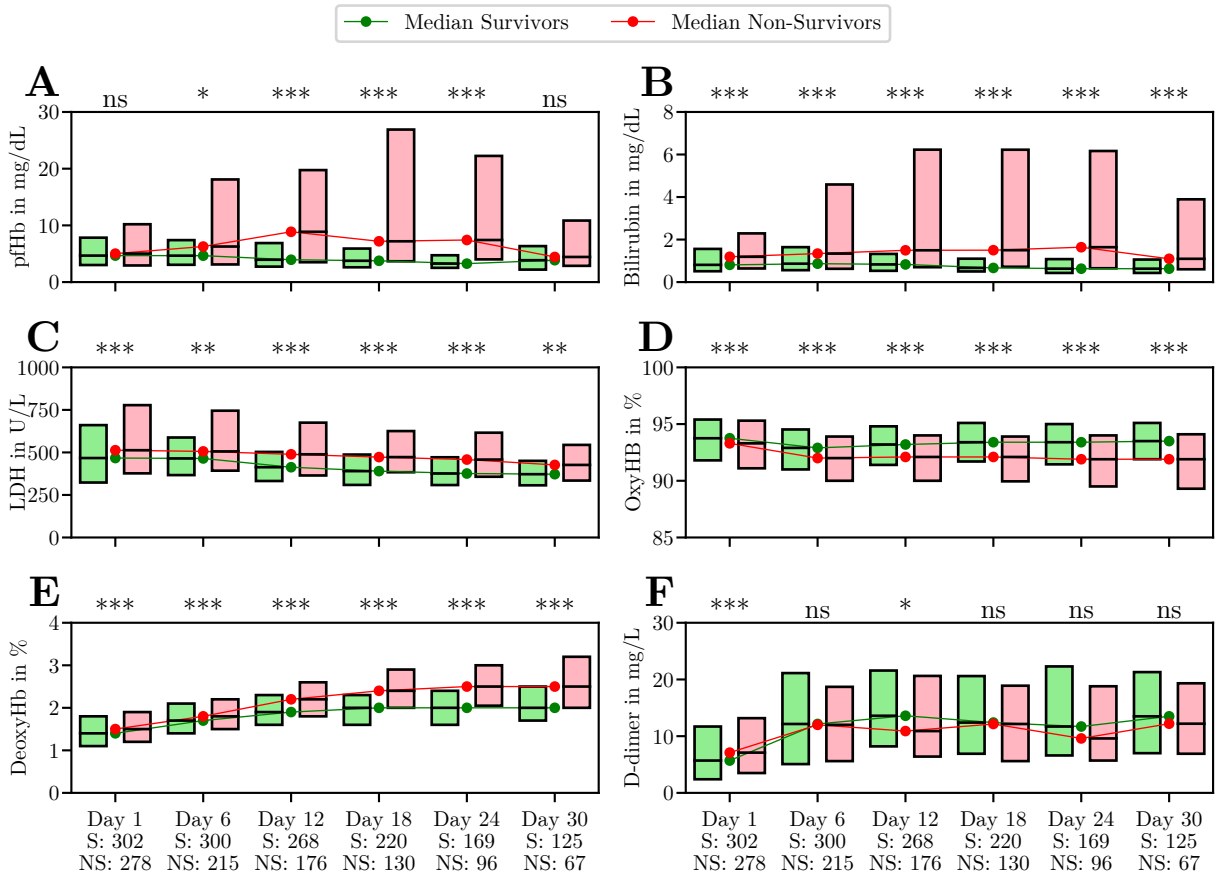

Figure A.2: Displays the temporal progression of the sub-cohorts, survivors (green) and non-survivors (red), over ECMO support days 1–30 for the *in-vivo* parameters plasma free hemoglobin (pfHb), bilirubin, lactate dehydrogenase (LDH), oxyhemoglobin (oxyHb), deoxyhemoglobin (deoxyHb), and D-dimer in the form of boxplots that indicate the 25^{th} and 75^{th} percentiles. Survivors and non-survivors are differentiated by the ICU survival. On the abscissa, ECMO support days and the number of unique patients per sub-cohort (S: survivor, N: non-survivor) are indicated. The significance level between the sub-cohorts of each ECMO support day is determined by the Mann-Whitney U test and indicated as follows: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\text{ns}}p \geq 0.05$.

scores, as well as baseline RBP operating conditions and laboratory findings during the first 48 hours of ECMO support. For this purpose, the cohort was stratified into survivors and non-survivors based on ICU survival status. Whereas Figure A.1 illustrates the overall distribution and variability of key RBP and blood parameters across the entire cohort, Table A.1 complements this by presenting the same parameters alongside additional clinical markers that reflect the complexity and heterogeneity of critically ill patients in the ICU setting. These include indicators of disease severity, inflammatory response, and organ dysfunction, as

well as demographic and outcome-related variables. This structured presentation helps reflect the complexity of the patient population, which includes varying comorbidities, degrees of organ dysfunction, and clinical courses. Together, the statistical overview (Figure A.1) and the grouped summary (Table A.1) provide complementary perspectives that support the interpretation of hemolysis markers and their potential relevance in the context of ECMO therapy. Table A.1 also notes missing data for selected parameters (e.g., BMI*, SOFA score**), which should be taken into account when interpreting results and performing statistical analyses. All abbreviations used are defined in the Abbreviations section at the beginning of the dissertation.

Not only are baseline characteristics and overall parameter distributions important for understanding the cohort, but the temporal progression of key clinical markers throughout the course of ECMO therapy is equally relevant. Observing how these values develop over time offers additional insight into patient trajectories and their association with outcome. Figure A.2 illustrates the progression of selected blood parameters over the first 30 days

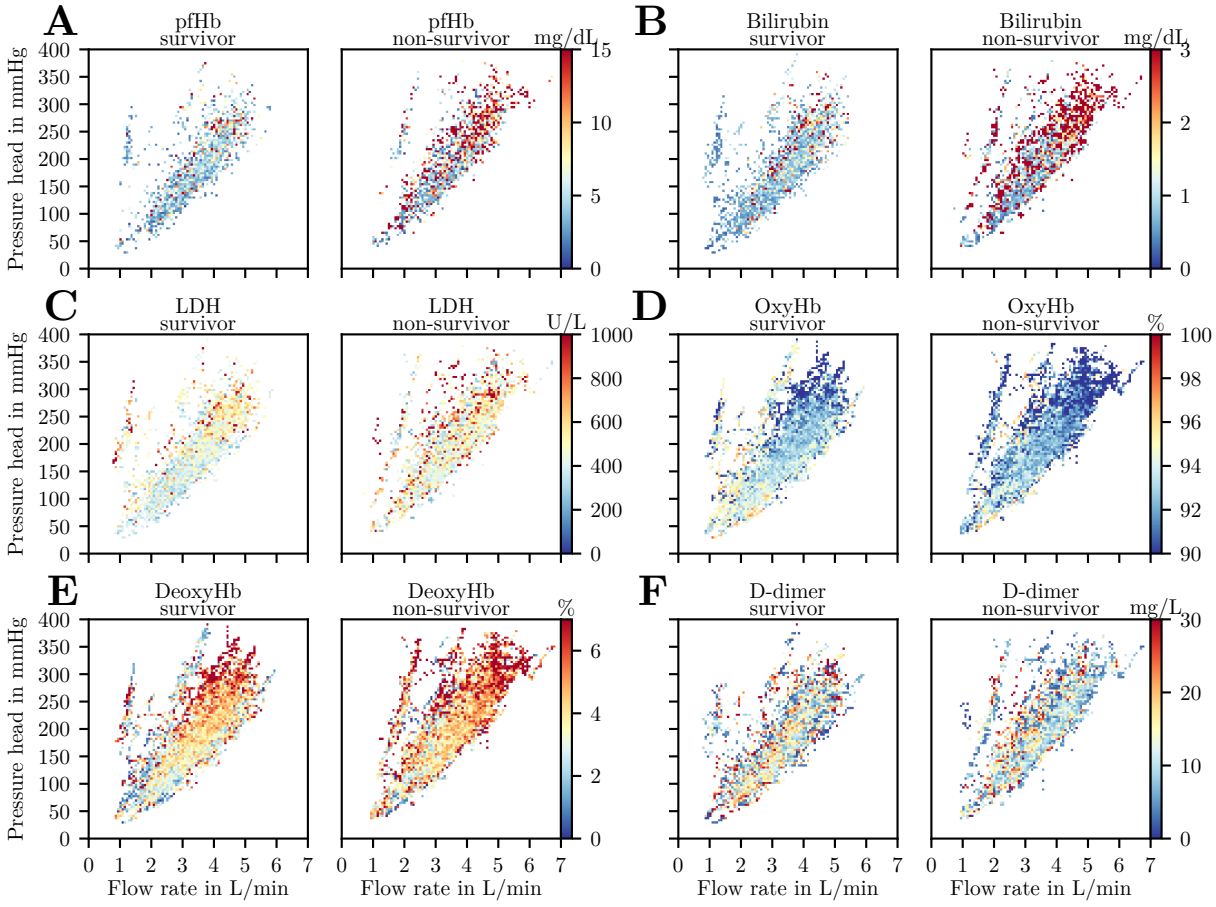

Figure A.3: Displays differences between blood parameter data of plasma free hemoglobin (pfhb), bilirubin, lactate dehydrogenase (LDH), oxyhemoglobin (oxyHb) and deoxyhemoglobin (deoxyHb) and D-dimer points stratified in the subgroups of survivor and non-survivor on panels A–F, respectively. For this figure median values are shown on a 90×90 grid.

of ECMO support, including pfHb, bilirubin, LDH, oxyHb, deoxyHb, and D-dimer. The

sub-cohorts survivors (green) and non-survivors (red) are compared based on ICU survival status. Boxplots represent the 25^{th} and 75^{th} percentiles, with the number of patients per group displayed along the x-axis. Statistical differences between the sub-cohorts at each time point were assessed using the Mann–Whitney U test and are marked accordingly: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\text{ns}}p \geq 0.05$. Throughout most of the support period, pfHb, bilirubin, LDH, and deoxyHb show significantly higher values in the non-survivor group. In contrast, oxyHb is consistently lower in non-survivors. D-dimer levels increase initially and then remain relatively stable, with a slight but non-significant upward trend in survivors. Particularly in Panels A and B, the presence of pronounced upper outliers in the non-survivor group suggests positively skewed distributions, where elevated values occur infrequently but can reach high levels when present. This is further analyzed in Figure A.3, where the main blood parameters are shown in the pressure head-flow rate plane, revealing differences between the subcohorts of survivors and non-survivors. It can be observed that for all except of the LDH and the D-dimer parameters the non-survivor subgroup consistently shows higher median values in the upper presssure head regions compared to the survivor subgroup. This is particularly

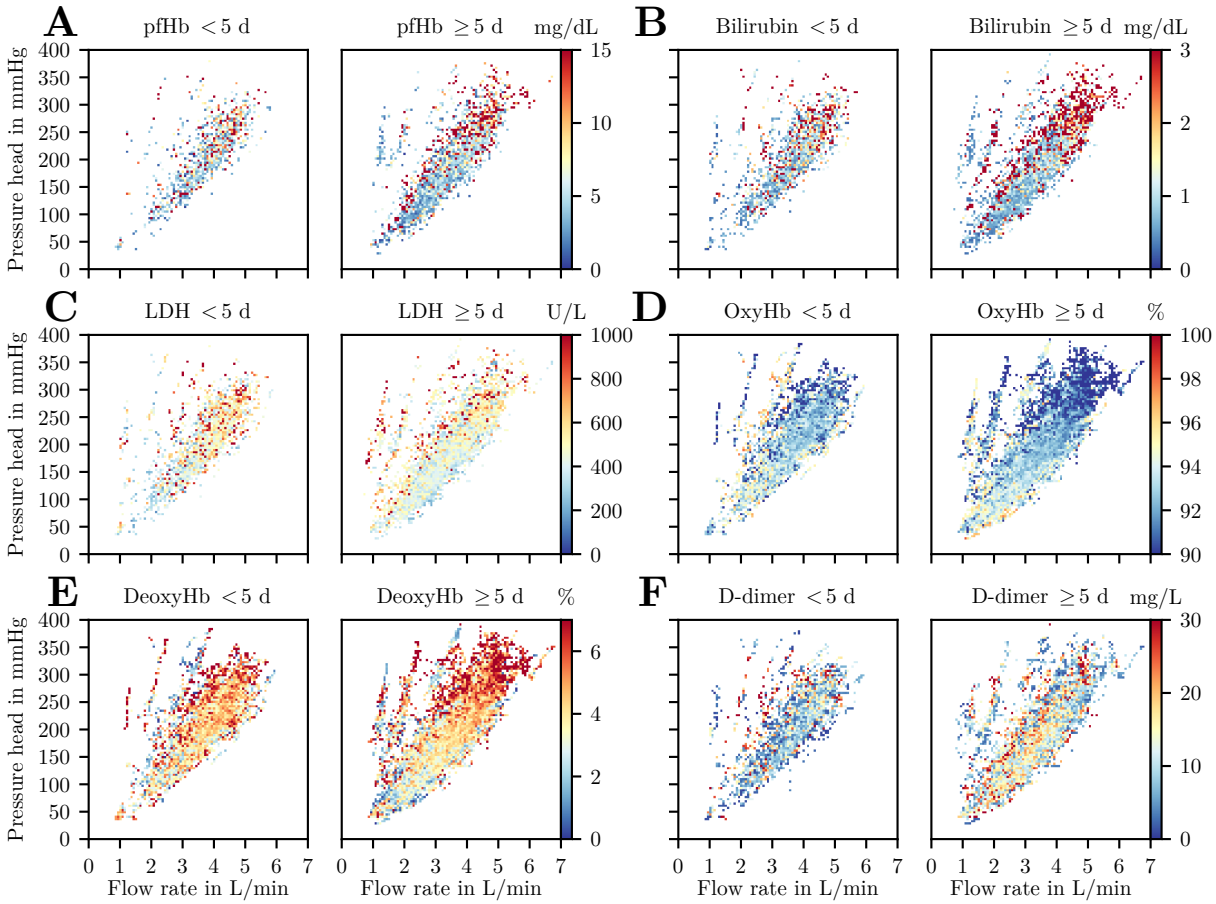

Figure A.4: The figure displays differences between all blood parameter data points up to the 5th day of ECMO treatment and after the 5th day of ECMO treatment for the blood parameters plasma free hemoglobin (pfhb), bilirubin, lactate dehydrogenase (LDH), oxyhemoglobin (oxyHb) and deoxyhemoglobin (deoxyHb) and D-dimer on panels A–F, respectively. For this figure median values are shown on a 90×90 grid.

evident in the clinical hemolysis marker pfHb, indicating that a treatment, which minimizes trauma to the blood, is of immense importance for the outcome of ECMO treatment.
To examine whether these elevated values are primarily baseline-related or develop over time, Figure A.4 further details this analysis by comparing all data points from the first five days of ECMO support to those collected thereafter. The analysis indicates that elevated values for pfHb, bilirubin, LDH, and deoxyHb mainly emerge during the course of treatment, rather than being present at baseline. This trend supports the assumption that blood trauma accumulates with ongoing ECMO support.
From the literature it is known that the RBP of the ECMO circuit plays a key role in inducing hemolysis [165–167]. Therefore a final breakdown of the RBP parameters pressure head and flow rate is shown in Figure A.5 for the entire cohort, as well as the survivor and non-survivor subgroups. This helps to investigate whether different operating conditions in the subgroups survivor and non-survivor could explain the observed differences in blood markers.
In the Panels C and F it can be seen that median flow rate and pressure head values are slightly elevated in non-survivors. While a Mann–Whitney U test revealed a statistically significant difference ($p < 0.001$), the corresponding effect sizes (Cohen's $d = 0.18$ for flow rate and 0.28 for pressure head) are small. This suggests that the overall operating points during ECMO support were largely comparable between groups. Thus, it is unlikely that differences in hemolysis markers are driven solely by external mechanical conditions. Instead, this finding supports the conclusion from the main chapter that substantial inter-individual variability in blood damage susceptibility exists within the clinical cohort.

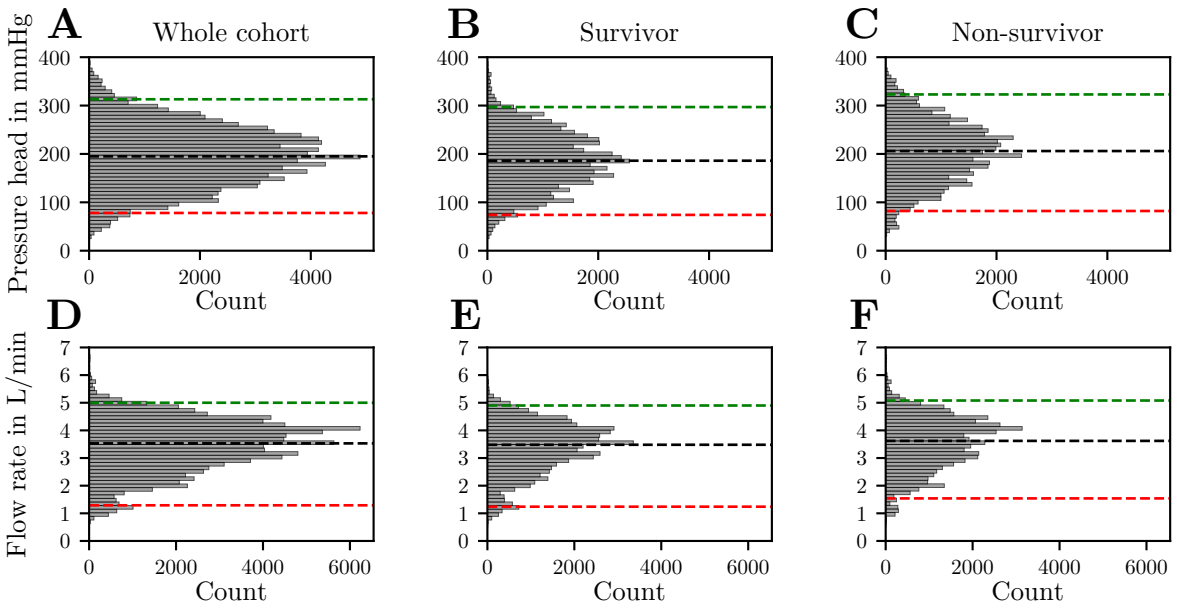

Figure A.5: Histograms of the RBP parameters pressure head in Panels A–C and flow rate in Panels D–F across the whole cohort (A and D) and the sub cohorts survivor (B and E) and non-survivor (C and F).

A.2 Linking in-vivo and in-silico data

As described in Chapter 4, the comparison between *in-vivo* blood parameter measurements and *in-silico* hemolysis predictions required linking two independent databases using patient ID and timestamp. Since the blood parameters and RBP operating metrics were recorded at different frequencies, defining a suitable time window for alignment was essential.
To determine the most appropriate interval, two aspects were analyzed: the frequency of measurements and the variability of operating point data over time.
Figure A.6 provides an overview of the sampling intervals across the entire patient cohort. As shown in panels A–C, pfHb, bilirubin, and LDH were most frequently measured every 24

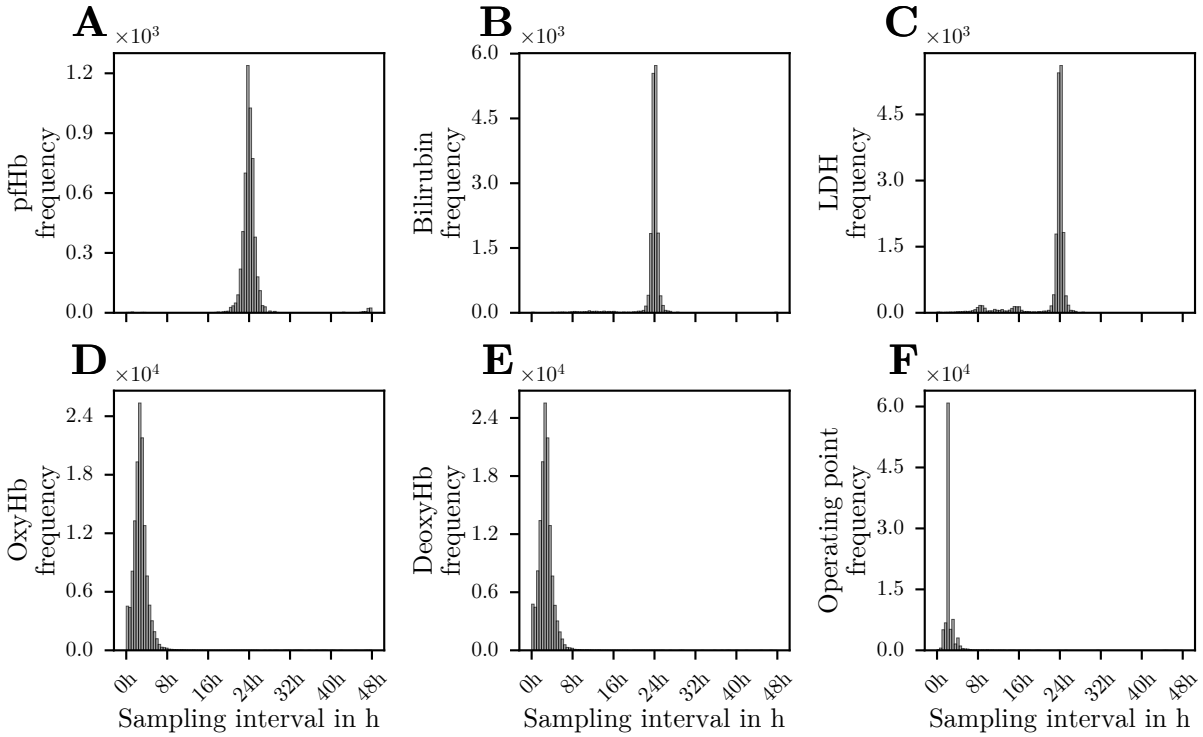

Figure A.6: Histogram of the measurement frequencies of plasma free hemoglobin (pfhb) (A), bilirubin (B), lactate dehydrogenase (LDH) (C), oxyhemoglobin (oxyHb) (D), deoxyhemoglobin (deoxyHb) (E) and RBP operating points (F) across the entire cohort.

hours, with the vast majority of entries tightly clustered around this interval. In contrast, oxyHb and deoxyHb (panels D and E) were generally measured every 2 hours, and the RBP operating points (panel F) also show a dominant peak at a 2-hour sampling interval. This implies that the RBP data offers higher temporal resolution than most blood parameter data, allowing for multiple RBP measurements to be associated with a single blood test result. Consequently, a time-averaging approach of the operating point data is appropriate as long as the resulting operating data reflect a sufficiently stable state around the averaging time window.
To quantify this temporal stability of the RBP operating data, Figure A.7 illustrates the variation in flow rate and pressure head across time intervals ranging from 1–48 hours. Panels

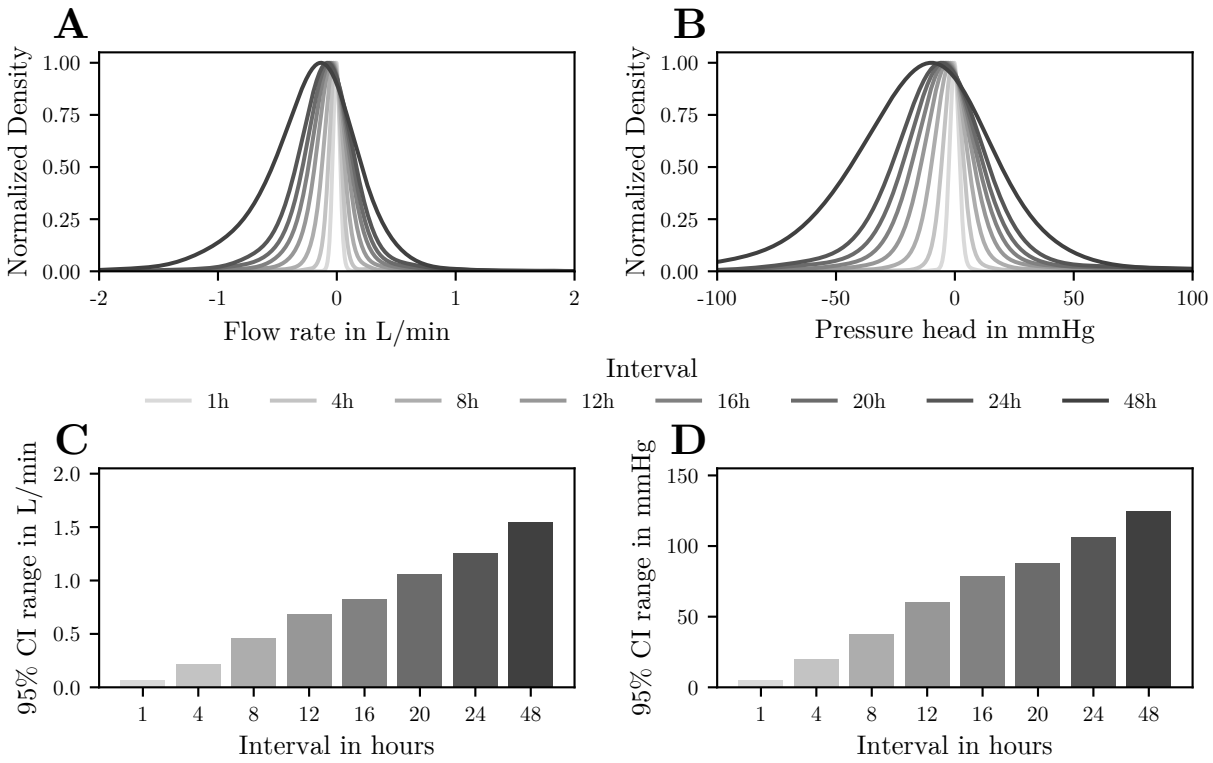

Figure A.7: Displays flow rate (A) and pressure head variation (B) for all cohort operating points over intervals of 1, 4, 8, 12, 16, 20, 24 and 48 hours. Panels C and D show the 95 % confidence intervals of variation range of all operating points during these specific time intervals for flow rate and pressure head, respectively.

A and B show the normalized density distributions of all flow rate and pressure head changes within a given time interval, respectively. The distributions become wider with increasing interval length, reflecting the growing variability over time. The corresponding 95 % confidence interval ranges of these distributions are plotted in panels C and D. Over an 8-hour interval, 95 % of the flow rate changes fall within overall change of 0.5 L/min, and pressure head changes are bounded within 50 mmHg. This represents approximately 7–10 % of the total operating point range defined in the RoI from Chapter 4.

This moderate level of variation over 8 hours was considered acceptable for averaging, especially given the subsequent correlation analysis between blood parameters and *in-silico* predictions, where the data was mapped onto a 20×20 grid of the RoI, corresponding to a spatial resolution of 5 %. The 95 % confidence interval for operating point changes across the full cohort spans 7–10 % of the RoI range, reflecting bidirectional variation. However, it is reasonable to assume that, for individual patients, changes within this time interval generally occur in one direction. Therefore, only about half of this range would typically apply per patient, which aligns well with the 5 % resolution of the analysis grid.

B Appendix to Chapter 5

This Appendix B provides additional experimental data, methodological details, and exploratory analyses to complement the results presented in Chapter 5, which is based on [179]. The contents support the development and validation of the refined shearing protocol, offer context for observed hemolysis trends, and explore a potential physiological factor contributing to variability in hemolysis outcomes.

B.1 Preliminary experiments

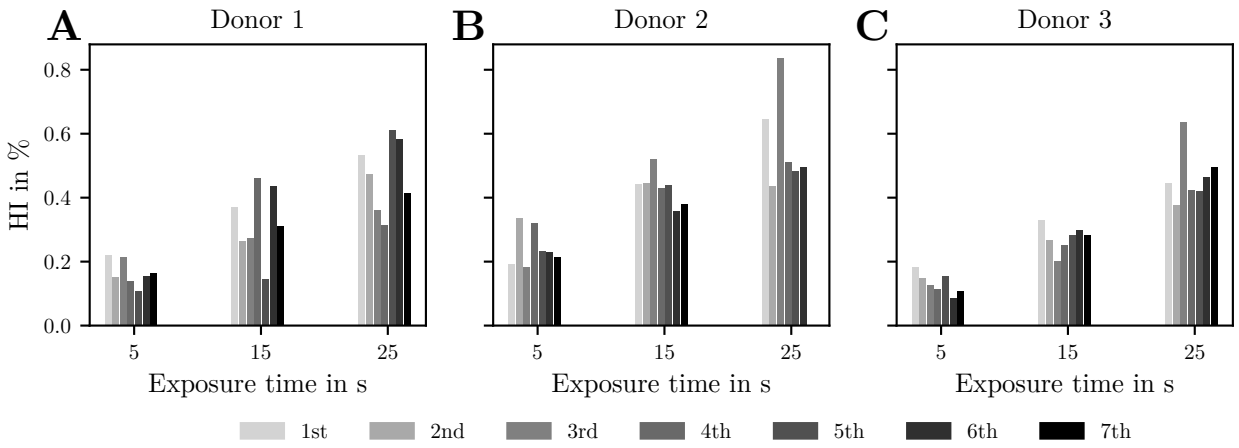

Figure B.1: Illustrates representative shear durations of 5, 15 and 25 seconds at a hematocrit of 30 % and a maximum rotational speed of 3000 rpm from three donor animals. Despite pronounced intra-sample variability, clear differences in hemolysis levels among the three shear durations are evident, with a repetition count of $n = 7$ for each duration.

To develop and test the novel shearing geometry (*see* Figure 5.1), multiple experimental runs were conducted using abattoir-derived porcine blood. In an initial test with a repetition count of $n = 6$, only the acceleration and deceleration triangular profile, totaling 0.48 seconds (*see* Figure 5.1), was evaluated. This experiment did not result in an increase in pfHb. The resulting HI values had an average of 0.08 %.
Subsequent experiments involved shear durations of 0.5, 1.0, 1.5, 5.0, 15.0 and 25.0 seconds at a maximum rotational speed of 3000 rpm, under varying hematocrit levels and different experimental protocols. These preliminary trials demonstrated that a shear duration of 15 seconds at a hematocrit of 30 % resulted in an HI of approximately 0.5 %. Moreover, these

extensive measurements informed the design and refinement of the experimental protocol presented in Chapter 5, which emphasizes reproducibility. Figure B.1 shows the results of a final experimental run validating the developed protocol using porcine blood subjected to maximum shear stress for 5, 15, and 25 seconds at a hematocrit of 30 %. An HI of 0.5 % indicates that 0.5 % of the total hemoglobin in the blood sample was released due to shear stress. Achieving this level of damage ensures two key outcomes: first, only a small fraction of RBC are lysed, minimizing alterations to the overall blood properties during the shear duration; second, this HI level achieves a good signal-to-noise ratio, accounting for potential measurement inaccuracies.

B.2 Baseline characteristics and temporal progression of blood parameters

To interpret hemolysis results in a physiologically meaningful context, baseline characteristics and temporal changes in blood parameters were closely monitored.
Table B.1 summarizes the mean and standard deviation of donor demographics, standard RBC metrics, and metabolic parameters (pH, glucose, lactate) for all five donors before and after dilution to a hematocrit of 25 ± 1 %. All baseline values fell within physiological norms,

Table B.1: Baseline mean values and standard deviations across all 5 human donors, including demographics (age, sex, and body mass index (BMI)), selected hematological RBC parameters (hematocrit (HCT), total hemoglobin (Hb), mean corpuscular volume (MCV), mean corpuscular hemoglobin concentration (MCHC), and red cell distribution width standard deviation (RDW-SD)) measured before hematocrit adjustment, and additional metabolic parameters (pH, glucose, and lactate) measured after hematocrit adjustment.

	Parameter	**Mean**	**Standard Deviation**
Demographics	Age/(years)	36.8	11.12
	Sex/(% female)	40	N/A
	BMI/(kg/m^2)	24.01	2.32
Before hematocrit adjustment	Hematocrit/(%)	39.72	2.27
	Total hemoglobin/(g/dL)	14.32	0.95
	MCV/(fL)	86.62	3.36
	MCHC/(g/dL)	36.08	0.85
	RDW-SD/(fL)	41.28	2.87
After hematocrit adjustment	Hematocrit/(%)	24.54	0.35
	Total hemoglobin/(g/dL)	8.92	0.16
	MCV/(fL)	85	3.39
	MCHC/(g/dL)	36.36	1.05
	RDW-SD/(fL)	40.04	3.02
	pH/(-)	7.28	0.04
	Glucose/(mmol/L)	5.02	0.74
	Lacate/(mmol/L)	1.94	0.42

confirming that the experimental blood samples were comparable and free from pre-existing abnormalities. This consistency is crucial for attributing observed hemolysis differences to

shear stress exposure rather than baseline variation.
Figure B.2 displays the temporal evolution of parameters such as pH, glucose, lactate, hematocrit, hemoglobin concentration, MCV, MCHC, and RDW. Donors 1–5 are distinguished by varying marker styles and shades of gray, progressing from light gray to black. The data indicate that all donors exhibited similar initial physiological values and consistent trends in blood parameters over the whole experimental duration. The slight increase in lactate levels and decrease in glucose levels had no detectable impact on RBC fragility throughout the experimental day. Furthermore, no association was observed between variations in MCV and RDW at the start of the experiment and an overall increased susceptibility to mechanical damage in RBCs.

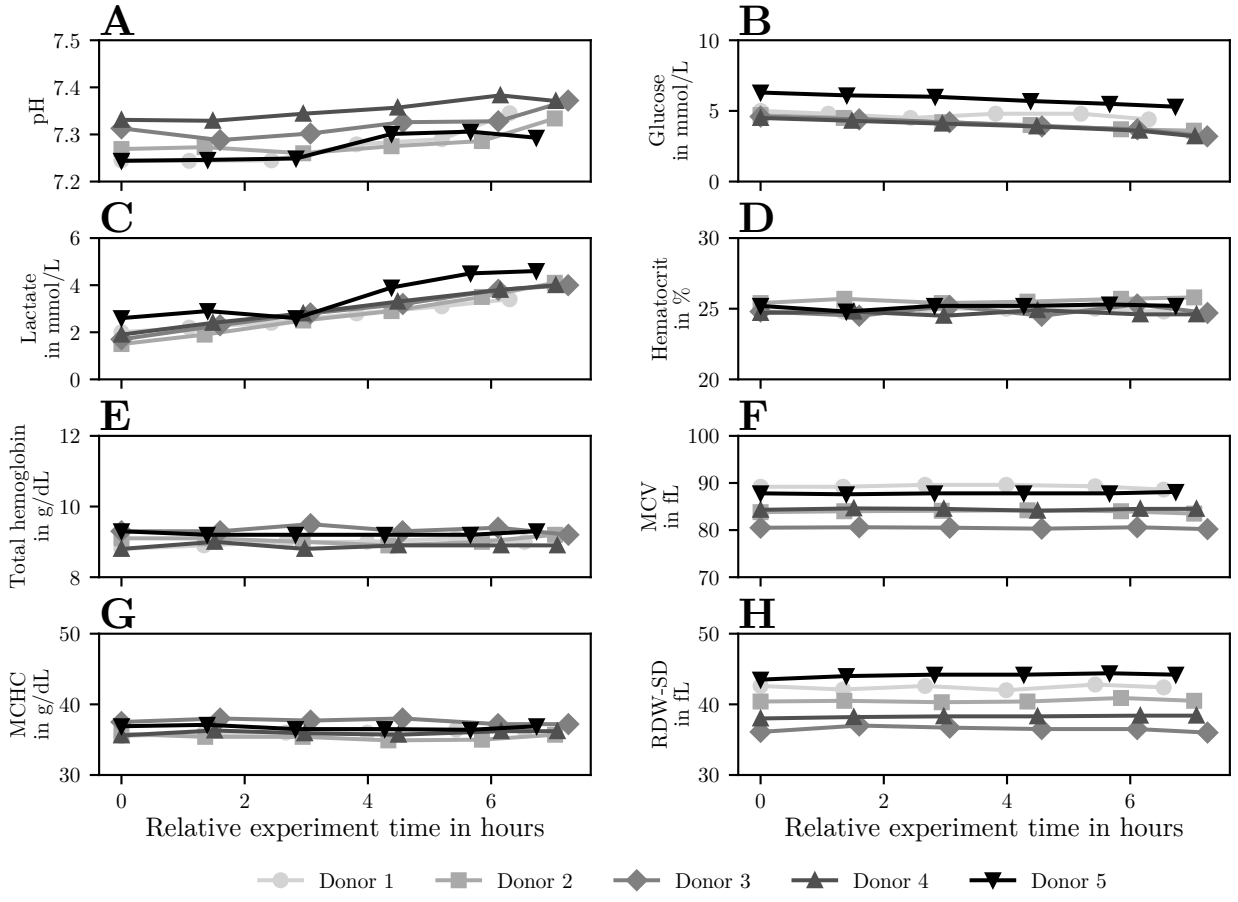

Figure B.2: Temporal progression of blood parameters [Potential of Hydrogen (pH), glucose, lactate, hematocrit, hemoglobin, Mean Corpuscular Volume (MCV), Mean Corpuscular Hemoglobin Concentration (MCHC), and Red Cell Distribution Width (RDW)] over the course of the experiment for five different donors. Each donor is represented by a unique marker and a shade of gray, ranging from light gray (Donor 1) to black (Donor 5).

B.3 Preliminary investigation into potential causes of increased intra-donor variability

Chapter 5 revealed substantial intra-donor variability in hemolysis index (HI) measurements. This observation prompted a follow-up investigation to explore potential contributing factors. The focus was placed on available blood count parameters, particularly RBC-related metrics

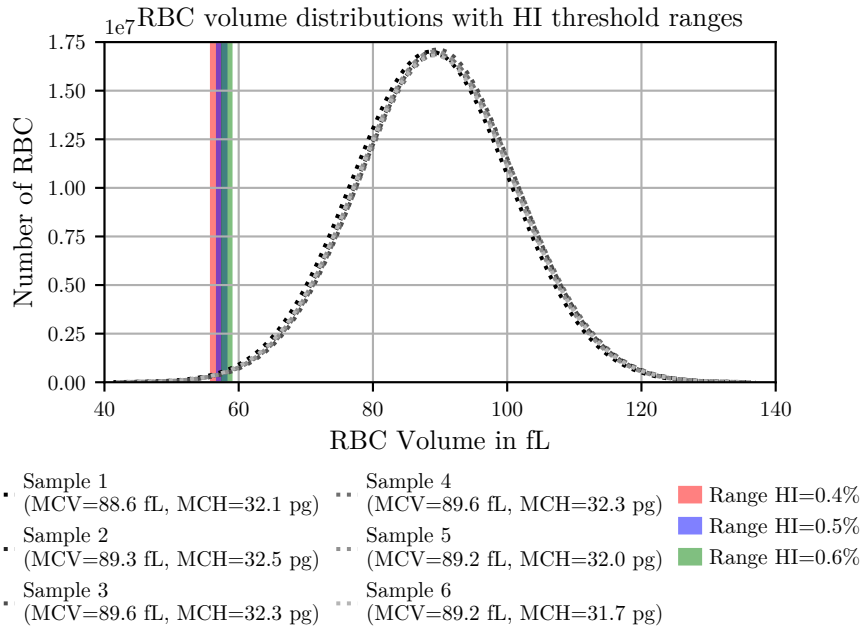

Figure B.3: Estimated RBC volume distributions from six blood samples collected from a single donor. Normal distributions are based on MCV and RDW-SD values for each sample. Colored bands indicate volume thresholds required to reach HI values of 0.4 % (red), 0.5 % (blue), and 0.6 % (green), assuming at a given shear rate first smaller, stiffer RBCs will rupture.

such as mean corpuscular volume (MCV), red cell distribution width (RDW-SD), and mean corpuscular hemoglobin (MCH).

As an initial step, it was analyzed whether any correlation existed between these RBC metrics and HI across all donors. While no statistically significant correlations were found, there was an indication that higher RDW-SD values (reflecting greater variation in RBC size) may be associated with increased intra-donor variability. Although the trend was not significant, it provided a rationale for further investigation.

To explore this observation in greater detail, six independent blood samples collected from a single donor during a day of the experiments were analyzed. For each sample, RBC volume distributions were approximated by assuming a normal distribution defined by the sample's MCV and RDW-SD values. At first glance, these volume distributions appeared very similar (*see* Figure B.3). However, under the assumption that smaller and potentially stiffer RBCs are more prone to shear-induced rupture, even minor shifts in the distribution could impact the HI. To quantify this effect, the number of the smallest RBCs that would need to rupture to produce an HI of 0.4, 0.5 or 0.6 % was estimated. Based on this, the corresponding volume

thresholds were determined. Because the distributions varied slightly between samples, the volume thresholds required to reach each HI value spanned a range, as shown in Figure B.3. The partial overlap of these threshold ranges suggests that even subtle fluctuations in RBC volume distribution may be sufficient to explain intra-donor HI variability of approximately $\pm 0.1\,\%$.

While this exploratory analysis does not offer a definitive explanation, it provides a plausible hypothesis and underscores the potential value of including routine blood count parameters when interpreting donor-specific hemolysis results.

C Appendix to Chapter 6

In this Appendix C, additional methodological details and supporting analyses related to the parameter estimation and reduced-order modeling presented in Chapter 6, which is based on [195], are provided, including MCMC diagnostics, ROM training, and parameter optimization.

C.1 MCMC

This section presents the trace plots of the four MCMC chains used in the parameter inference process (*see* Figure C.1). These plots are essential for assessing the mixing behavior and

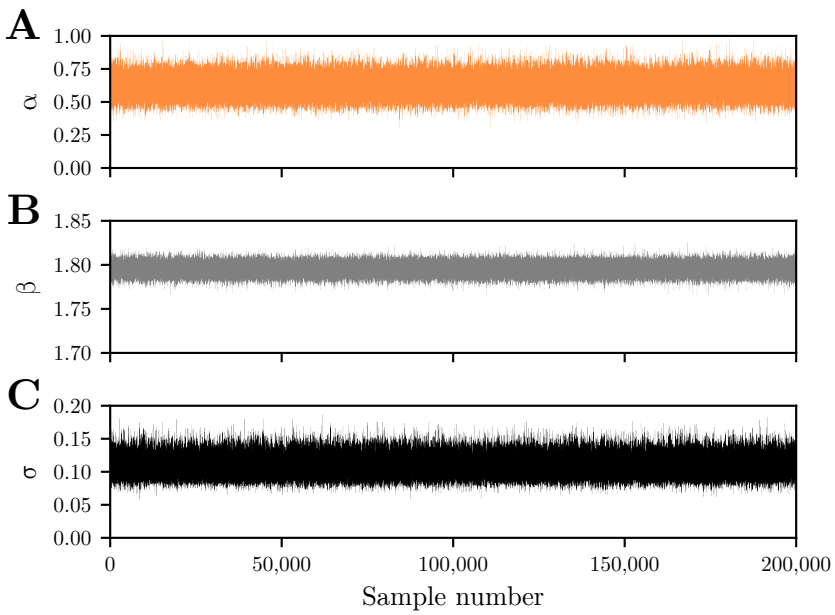

Figure C.1: Trace plots of the parameters α, β, and σ from the constant C Markov Chain Monte Carlo (MCMC) simulation, shown in panels A, B, and C respectively. These plots are used to assess convergence and mixing of the chains over 200,000 samples, showing no visible trends or drift, which indicates good mixing and convergence.

convergence of the sampling process. Each chain was run for 50,000 iterations, following a burn-in period of 1000 samples to allow initial transients to dissipate. Stable trace plots without visible trends are indicative of good stationarity and suggest that the chains adequately explore the target posterior distribution. In this case, the samples span the parameter space without

signs of correlation or divergence. Convergence was further evaluated using the Gelman-Rubin diagnostic (R-hat), which yielded values of 1 for all parameters. An R-hat value close to 1 confirms that the chains have converged to a common posterior distribution. These diagnostic results provide confidence in the reliability of the inferred parameter distributions, which serve as the basis for all subsequent surrogate modeling and uncertainty quantification tasks in Chapter 6.
To illustrate the impact of experimental variance on the results of the MCMC method, synthetic experimental data without variance were generated over the same interval as the experimental data of Zhang *et al.* [106]. These synthetic data are positioned exactly on the

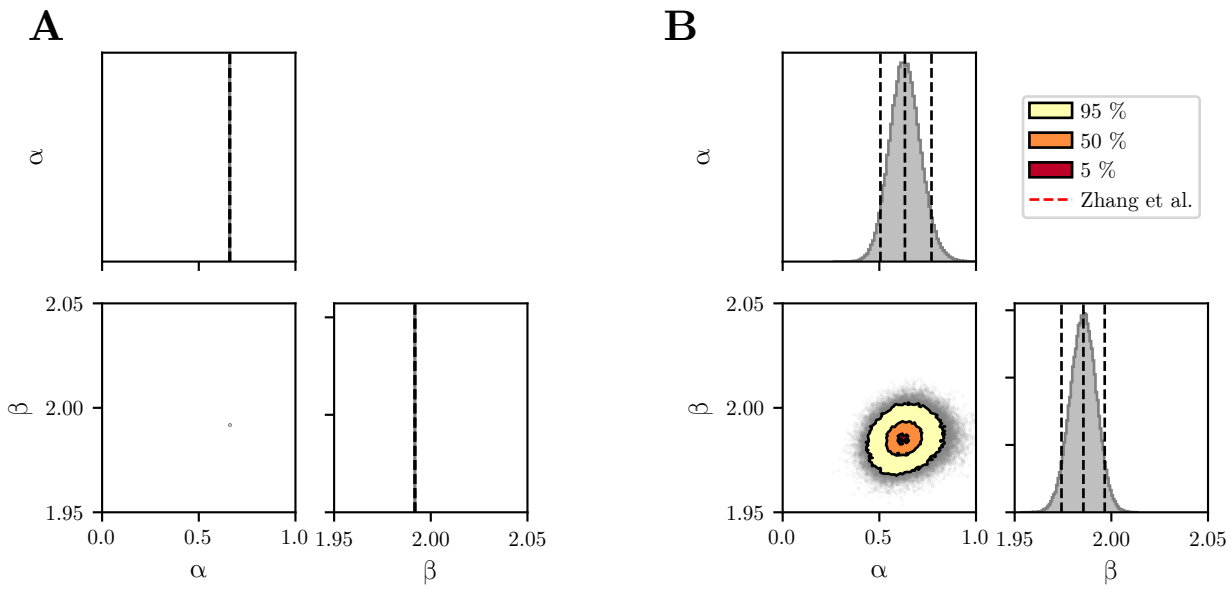

Figure C.2: Corner plot comparison of the MCMC results using a constant $C = 1.228 \times 10^{-5}$, based on (A) synthetic data without variance and (B) experimental data including measurement variance.

surface defined by Zhang's parameter values. As shown in Figure C.2 A, the MCMC method yields Zhang's parameter values precisely, without any distribution and σ being 0, when using the synthetic data. This demonstrates that the distributions of the model parameters resulting from the MCMC method and the actual experimental data directly originate from the variance of the experimental data (*see* Figure C.2B).

C.2 NIPCE

To generate training data samples for the ROM, the posterior distributions of the model parameters (C, α, β, and σ) obtained from the MCMC sampling were first fitted using log-normal probability distributions. To ensure robustness and avoid the influence of extreme values, these distributions were truncated at the lower and upper ends of the 99 % confidence intervals (*see* Figure C.3). The individual fitted distributions were then combined to form a multivariate distribution. From this multivariate distribution, 200 samples were drawn using Latin Hypercube Sampling (LHS), a method that ensures a well-distributed sampling

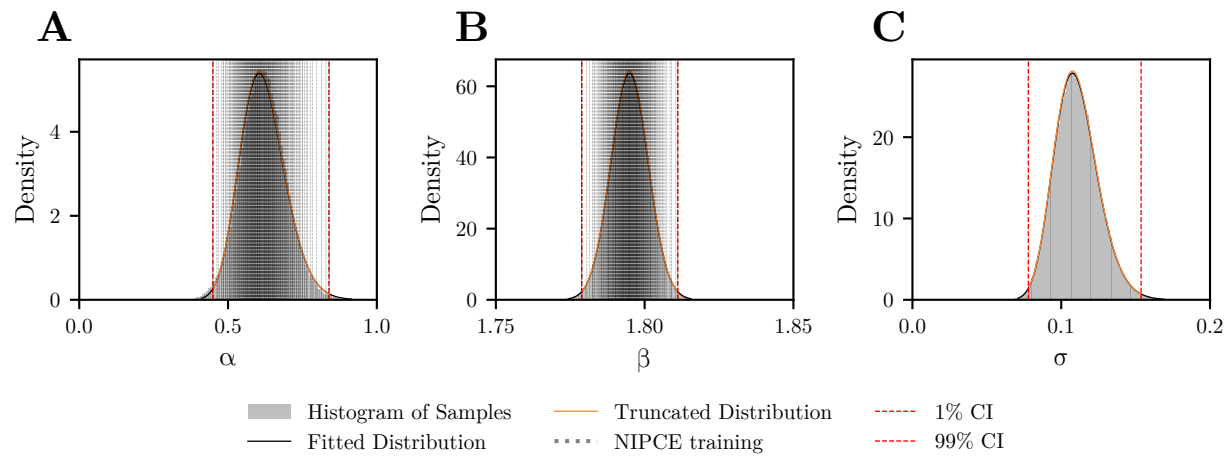

Figure C.3: Posterior Distributions of the model parameters α, β and σ are shown in Panels A, B and C respectively. Fitted log-Normal distributions are indicated by black lines and truncated distributions by orange lines. The 1 % and 99 % truncation criterions are indicated by dashed red lines. For α and β the Latin Hypercube training points are shown with black dotted lines.

across the parameter space defined by the 99 % confidence intervals. Additionally, to include potential boundary effects, combinations of the maximum and minimum values of the individual distributions were added, resulting in a total of 209 (for the C, α, β models) or 204 (for the α, β models) training points. Each of these parameter sets was used to compute hemolysis values via the CFD post-processing approach described in Chapter 6, thus providing the data needed to train the NIPCE model.

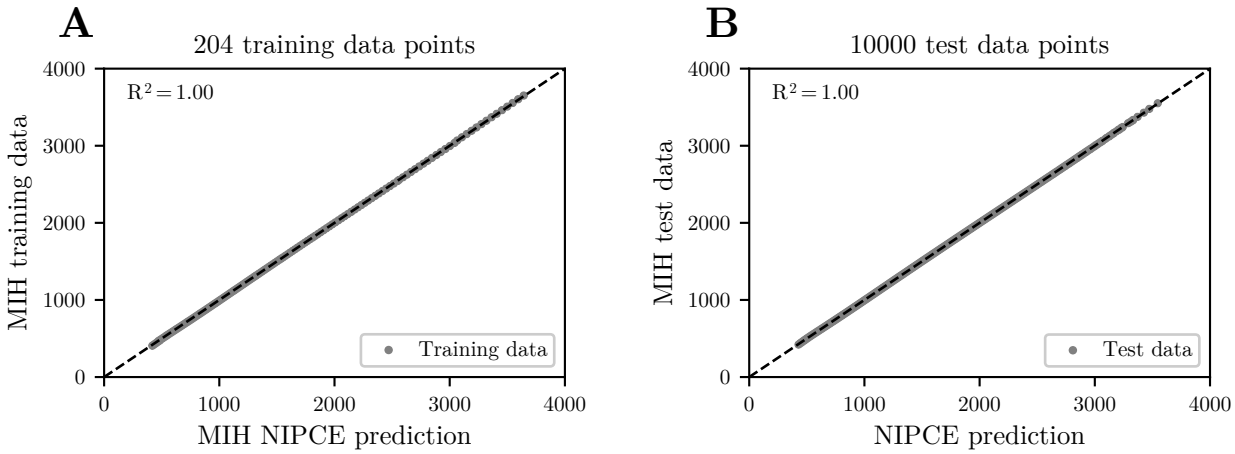

Figure C.4: Correlation plot of NIPCE prediction and CFD calculation for MIH across 204 training data points (A) and 10,000 test data points for Condition 5 of the FDA RBP benchmark.

A fourth-order polynomial expansion was selected for the NIPCE model, balancing complexity and computational efficiency. Once trained, the NIPCE model enables rapid transformation of large numbers of MCMC samples into corresponding hemolysis value distributions within seconds.

The predictive accuracy of the NIPCE surrogate model was validated by comparing its outputs with those from the full CFD-based pipeline. While this validation was performed

for all benchmark conditions of the FDA pump, Figure C.4 presents a representative case for Condition 5. In Panel (A), the NIPCE model predictions closely match the original CFD results for all training points, indicating excellent prediction on the traing data set. Panel (B) demonstrates the model's generalization capability: predictions for 10,000 additional test samples, drawn randomly within the truncated posterior space, also show strong agreement with the CFD baseline. This confirms that the surrogate model reliably approximates hemolysis values across the full range of parameter uncertainty.

C.3 Parameter optimization

Bayesian optimization is an iterative process that uses a surrogate model to approximate the objective function and an acquisition function to decide where to sample next. In Chapter 6, this method was applied to find the optimal C value by minimizing the median of the σ distribution. The optimization process used a Gaussian Process (GP) as the surrogate model to approximate the relationship between C values and the median of the σ distribution. Additionally, an acquisition function guided the selection of the next C value to evaluate. The Bayesian optimization was performed on the parameter σ, with C values ranging from 1×10^{-6} to 1×10^{-3}. For each C value, the posterior distributions of α, β and σ were evaluated. As can be seen from Figure C.5 the optimization process identified 3.515×10^{-3} as the optimal C value after 100 Iterations.

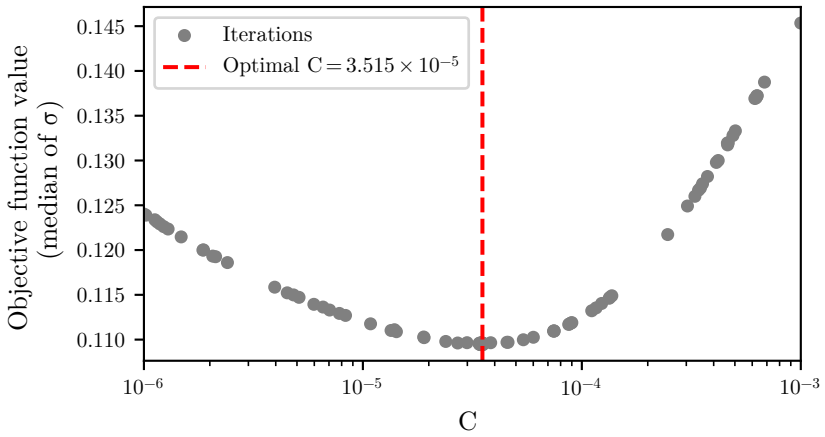

Figure C.5: Optimization history of 100 iterations finding the smallest median corresponding to an optimal $C = 3.515 \times 10^{-3}$ as indicated by a red vertical line. Grey points indicate the individual objective function realizations.

List of Figures

List of Tables

References

[1] Extracorporeal Life Support Organization. *ECLS International Summary of Statistics*. URL: https://www.elso.org/registry/internationalsummaryandreports/internationalsummary.aspx (visited on 06/2025).

[2] Nunez, J. I. et al. "Bleeding and thrombotic events in adults supported with venovenous extracorporeal membrane oxygenation: an ELSO registry analysis". *Intensive care medicine* 48.2 (2022), pp. 213–224. DOI: 10.1007/s00134-021-06593-x.

[3] Willers, A. et al. "Longitudinal Trends in Bleeding Complications on Extracorporeal Life Support Over the Past Two Decades-Extracorporeal Life Support Organization Registry Analysis". *Critical care medicine* 50.6 (2022), e569–e580. DOI: 10.1097/CCM.0000000000005466.

[4] Valladolid, C. et al. "von Willebrand Factor, Free Hemoglobin and Thrombosis in ECMO". *Frontiers in medicine* 5 (2018), p. 228. ISSN: 2296-858X. DOI: 10.3389/fmed.2018.00228.

[5] Sun, W. et al. "Impact of high mechanical shear stress and oxygenator membrane surface on blood damage relevant to thrombosis and bleeding in a pediatric ECMO circuit". *Artificial organs* 44.7 (2020), pp. 717–726. DOI: 10.1111/aor.13646.

[6] Reichel, F. et al. "High-Throughput Microfluidic Characterization of Erythrocyte Shapes and Mechanical Variability". *Biophysical Journal* 117.1 (2019), pp. 14–24. DOI: 10.1016/j.bpj.2019.05.022.

[7] Hadjesfandiari, N. et al. "Current Understanding of the Relationship between Blood Donor Variability and Blood Component Quality". *International journal of molecular sciences* 22.8 (2021). DOI: 10.3390/ijms22083943.

[8] Ponnaluri, S. V. et al. "Results of the Interlaboratory Computational Fluid Dynamics Study of the FDA Benchmark Blood Pump". *Annals of biomedical engineering* 51.1 (2023), pp. 253–269. DOI: 10.1007/s10439-022-03105-w.

[9] Karagiannidis, C. et al. "Physiological and Technical Considerations of Extracorporeal CO2 Removal". *Critical Care* 23.1 (2019), p. 75. DOI: 10.1186/s13054-019-2367-z.

[10] ELSO. "Guidelines for Cardiopulmonary Extracorporeal Life Support Extracorporeal Life Support Organization, Version 1.4 August 2017 Ann Arbor, MI, USA". URL: www.elso.org.

[11] Strengers, P. F. W. "Evidence-based clinical indications of plasma products and future prospects". *Annals of Blood* 2 (2017), p. 20. DOI: 10.21037/aob.2017.12.03.

[12] Lee, S. et al. “Narrowed reference intervals for complete blood count in a multiethnic population”. *Clinical chemistry and laboratory medicine* 57.9 (2019), pp. 1382–1387. DOI: 10.1515/cclm-2018-1263.

[13] Richardson, K. J. et al. “Beyond oxygen transport: active role of erythrocytes in the regulation of blood flow”. *American journal of physiology. Heart and circulatory physiology* 319.4 (2020), H866–H872. DOI: 10.1152/ajpheart.00441.2020.

[14] Elahi, S. “Neglected Cells: Immunomodulatory Roles of CD71+ Erythroid Cells”. *Trends in immunology* 40.3 (2019), pp. 181–185. DOI: 10.1016/j.it.2019.01.003.

[15] Jensen, F. B. “The dual roles of red blood cells in tissue oxygen delivery: oxygen carriers and regulators of local blood flow”. *The Journal of experimental biology* 212.Pt 21 (2009), pp. 3387–3393. DOI: 10.1242/jeb.023697.

[16] Chen, G. et al. “Exploring ethnic differences in the distribution of blood test results in healthy adult populations to inform earlier cancer detection: a systematic review”. *Family practice* 41.5 (2024), pp. 638–648. DOI: 10.1093/fampra/cmae021.

[17] Kanias, T. et al. “Ethnicity, sex, and age are determinants of red blood cell storage and stress hemolysis: results of the REDS-III RBC-Omics study”. *Blood Advances* 1.15 (2017), pp. 1132–1141. URL: 10.1182/bloodadvances.2017004820.

[18] Turkbey, E. B. et al. “Determinants and normal values of ascending aortic diameter by age, gender, and race/ethnicity in the Multi-Ethnic Study of Atherosclerosis (MESA)”. *Journal of magnetic resonance imaging : JMRI* 39.2 (2014), pp. 360–368. DOI: 10.1002/jmri.24183.

[19] Potter, R. F. and Groom, A. C. “Capillary Diameter and Geometry in Cardiac and Skeletal Muscle Studied by Means of Corrosion Casts”. *Microvasc Res.* 25.1 (1983), pp. 68–84. URL: doi:%2010.1016/0026-2862(83)90044-4.

[20] Callaghan, F. M. and Grieve, S. M. “Normal patterns of thoracic aortic wall shear stress measured using four-dimensional flow MRI in a large population”. *American journal of physiology. Heart and circulatory physiology* 315.5 (2018), H1174–H1181. DOI: 10.1152/ajpheart.00017.2018.

[21] Popel, A. S. and Johnson, P. C. “Microcirculation and Hemorheology”. *Annual review of fluid mechanics* 37 (2005), pp. 43–69. ISSN: 0066-4189. DOI: 10.1146/annurev.fluid.37.042604.133933.

[22] Horner, J. S. et al. “Measurements of human blood viscoelasticity and thixotropy under steady and transient shear and constitutive modeling thereof”. *Journal of Rheology* 63.5 (2019), pp. 799–813. ISSN: 0148-6055. DOI: 10.1122/1.5108737.

[23] Nader, E. et al. “Blood Rheology: Key Parameters, Impact on Blood Flow, Role in Sickle Cell Disease and Effects of Exercise”. *Frontiers in physiology* 10 (2019), p. 1329. ISSN: 1664-042X. DOI: 10.3389/fphys.2019.01329.

[24] Lanotte, L. et al. “Red cells’ dynamic morphologies govern blood shear thinning under microcirculatory flow conditions”. *Proceedings of the National Academy of Sciences of the United States of America* 113.47 (2016), pp. 13289–13294. DOI: 10.1073/pnas.1608074113.

[25] Zhang, W. et al. "Viscoelasticity reduces the dynamic stresses and strains in the vessel wall: implications for vessel fatigue". *American journal of physiology. Heart and circulatory physiology* 293.4 (2007), H2355–60. DOI: 10.1152/ajpheart.00423.2007.

[26] Thurston, G. B. "Viscoelasticity of Human Blood". *Biophysical Journal* 12.9 (1972), pp. 1205–1217. URL: https://doi.org/10.1016/S0006-3495(72)86156-3.

[27] Brust, M. et al. "Rheology of human blood plasma: viscoelastic versus Newtonian behavior". *Physical review letters* 110.7 (2013), p. 078305. DOI: 10.1103/PhysRevLett.110.078305.

[28] Baskurt, O. K. and Meiselman, H. J. "Erythrocyte aggregation: basic aspects and clinical importance". *Clinical hemorheology and microcirculation* 53.1-2 (2013), pp. 23–37. DOI: 10.3233/CH-2012-1573.

[29] Fung, Y. C. et al. "High-resolution data on the geometry of red blood cells." *Biorheology* 18.3-6 (1981), pp. 369–385. URL: doi:%2010.3233/bir-1981-183-606.%20PMID:%207326382..

[30] Himbert, S. et al. "The Molecular Structure of Human Red Blood Cell Membranes from Highly Oriented, Solid Supported Multi-Lamellar Membranes". *Scientific reports* 7 (2017), p. 39661. DOI: 10.1038/srep39661.

[31] Skalak, R. et al. "Biofluid Mechanics". *Annual review of fluid mechanics* 21 (1989), pp. 167–200. ISSN: 0066-4189. URL: https://doi.org/10.1146/annurev.fl.21.010189.

[32] Damiani, E. et al. "Plasma free hemoglobin and microcirculatory response to fresh or old blood transfusions in sepsis". *PloS one* 10.5 (2015), e0122655. DOI: 10.1371/journal.pone.0122655.

[33] Chai, S. et al. "Novel nonsense mutation p. Gln264Ter in the ANK1 confirms causative role for hereditary spherocytosis: a case report". *BMC medical genetics* 21.1 (2020), p. 223. DOI: 10.1186/s12881-020-01161-4.

[34] Pretini, V. et al. "Red Blood Cells: Chasing Interactions". *Frontiers in physiology* 10 (2019), p. 945. ISSN: 1664-042X. DOI: 10.3389/fphys.2019.00945.

[35] Mohandas, N. and Gallagher, P. "Red cell membrane: past, present, and future". *Blood* 112.10 (2008), pp. 3939–3948. ISSN: 0006-4971. URL: doi:%2010.1182/blood-2008-07-161166.

[36] Singel, D. J. and Stamler, J. S. "Chemical physiology of blood flow regulation by red blood cells: the role of nitric oxide and S-nitrosohemoglobin". *Annual review of physiology* 67 (2005), pp. 99–145. ISSN: 0066-4278. DOI: 10.1146/annurev.physiol.67.060603.090918.

[37] Kuck, L. et al. "Active modulation of human erythrocyte mechanics". *American journal of physiology. Cell physiology* 319.2 (2020), pp. C250–C257. DOI: 10.1152/ajpcell.00210.2020.

[38] Horobin, J. T. et al. "Repetitive Supra-Physiological Shear Stress Impairs Red Blood Cell Deformability and Induces Hemolysis". *Artificial organs* 41.11 (2017), pp. 1017–1025. DOI: 10.1111/aor.12890.

[39] Delvasto-Nuñez, L. et al. “It takes two to thrombosis: Hemolysis and complement”. *Blood reviews* 50 (2021), p. 100834. DOI: 10.1016/j.blre.2021.100834.

[40] Woollard, K. J. et al. “Erythrocyte hemolysis and hemoglobin oxidation promote ferric chloride-induced vascular injury”. *The Journal of biological chemistry* 284.19 (2009), pp. 13110–13118. ISSN: 0021-9258. DOI: 10.1074/jbc.M809095200.

[41] van Avondt, K. et al. “Mechanisms of haemolysis-induced kidney injury”. *Nature reviews. Nephrology* 15.11 (2019), pp. 671–692. DOI: 10.1038/s41581-019-0181-0.

[42] Schlagenhauf, A. et al. “Acquired von Willebrand Syndrome and Platelet Function Defects during Extracorporeal Life Support (Mechanical Circulatory Support)”. *Hamostaseologie* 40.2 (2020), pp. 221–225. DOI: 10.1055/a-1150-2016.

[43] Goh, T. et al. “Thrombotic response to mechanical circulatory support devices”. *Journal of thrombosis and haemostasis : JTH* (2025). DOI: 10.1016/j.jtha.2025.02.037.

[44] Ploppa, A. et al. “Mechanisms of leukocyte distribution during sepsis: an experimental study on the interdependence of cell activation, shear stress and endothelial injury”. *Critical Care* 14.6 (2010), R201. URL: https://doi.org/10.1186/cc9322.

[45] Brogan, T. et al. *Extracorporeal life support: the ELSO Red book*. 5th edition. Ann Arbor, Michigan: Extracorporeal Life Support Organization, 2017. ISBN: 978-0-9656756-5-9.

[46] Navas-Blanco, J. R. et al. “Parallel veno-venous and veno-arterial extracorporeal membrane circuits for coexisting refractory hypoxemia and cardiovascular failure: a case report”. *BMC anesthesiology* 21.1 (2021), p. 77. DOI: 10.1186/s12871-021-01299-5.

[47] Squiers, J. J. et al. “Contemporary extracorporeal membrane oxygenation therapy in adults: Fundamental principles and systematic review of the evidence”. *The Journal of thoracic and cardiovascular surgery* 152.1 (2016), pp. 20–32. DOI: 10.1016/j.jtcvs.2016.02.067.

[48] Erdil, T. et al. “Extracorporeal membrane oxygenation support in pediatrics”. *Annals of cardiothoracic surgery* 8.1 (2019), pp. 109–115. ISSN: 2225-319X. DOI: 10.21037/acs.2018.09.08.

[49] MacLaren, G. et al. *Extracorporeal life support: the ELSO Red book*. 6th edition. Ann Arbor, Michigan: Extracorporeal Life Support Organization, 2022. ISBN: 978-0-9656756-8-0.

[50] Kohler, K. et al. “ECMO cannula review”. *Perfusion* 28.2 (2013), pp. 114–124. DOI: 10.1177/0267659112468014.

[51] Fuchs, G. et al. “Flow-induced platelet activation in components of the extracorporeal membrane oxygenation circuit”. *Scientific reports* 8.1 (2018), p. 13985. DOI: 10.1038/s41598-018-32247-y.

[52] Williams, D. C. et al. “Circuit oxygenator contributes to extracorporeal membrane oxygenation-induced hemolysis”. *ASAIO journal (American Society for Artificial Internal Organs : 1992)* 61.2 (2015), pp. 190–195. DOI: 10.1097/MAT.0000000000000173.

[53] Doyle, A. J. and Hunt, B. J. "Current Understanding of How Extracorporeal Membrane Oxygenators Activate Haemostasis and Other Blood Components". *Frontiers in medicine* 5 (2018), p. 352. ISSN: 2296-858X. DOI: 10.3389/fmed.2018.00352.

[54] Gartner, M. J. et al. "Modeling flow effects on thrombotic deposition in a membrane oxygenator". *Artificial organs* 24.1 (2000), pp. 29–36. DOI: 10.1046/j.1525-1594.2000.06384.x.

[55] Kato, G. J. et al. "Lactate dehydrogenase as a biomarker of hemolysis-associated nitric oxide resistance, priapism, leg ulceration, pulmonary hypertension, and death in patients with sickle cell disease". *Blood* 107.6 (2006), pp. 2279–2285. ISSN: 0006-4971. DOI: 10.1182/blood-2005-06-2373.

[56] Barcellini, W. and Fattizzo, B. "Clinical Applications of Hemolytic Markers in the Differential Diagnosis and Management of Hemolytic Anemia". *Disease markers* 2015 (2015), p. 635670. DOI: 10.1155/2015/635670.

[57] Winnersbach, P. et al. "Clinical Practice of Pre-Assembling and Storing of Extracorporeal Membrane Oxygenation Systems". *ASAIO journal (American Society for Artificial Internal Organs : 1992)* 70.11 (2024), pp. 979–986. DOI: 10.1097/MAT.0000000000002232.

[58] Martucci, G. et al. "Transfusion practice in patients receiving VV ECMO (PROTECMO): a prospective, multicentre, observational study". *The Lancet. Respiratory medicine* 11.3 (2023), pp. 245–255. DOI: 10.1016/S2213-2600(22)00353-8.

[59] Adkins, K. L. "Sedation Strategies for Extracorporeal Membrane Oxygenation Support". *ASAIO journal (American Society for Artificial Internal Organs : 1992)* 63.2 (2017), pp. 113–114. DOI: 10.1097/MAT.0000000000000536.

[60] Vajter, J. and Volod, O. "Anticoagulation Management During ECMO: Narrative Review". *JHLT open* 8 (2025), p. 100216. DOI: 10.1016/j.jhlto.2025.100216.

[61] Hayes, K. et al. "Rehabilitation of adult patients on extracorporeal membrane oxygenation: A scoping review". *Australian critical care : official journal of the Confederation of Australian Critical Care Nurses* 35.5 (2022), pp. 575–582. ISSN: 1036-7314. DOI: 10.1016/j.aucc.2021.08.009.

[62] Vasques, F. et al. "How I wean patients from veno-venous extra-corporeal membrane oxygenation". *Critical care (London, England)* 23.1 (2019), p. 316. DOI: 10.1186/s13054-019-2592-5.

[63] Charbonneau, F. et al. "Parameters associated with successful weaning of veno-arterial extracorporeal membrane oxygenation: a systematic review". *Critical care (London, England)* 26.1 (2022), p. 375. DOI: 10.1186/s13054-022-04249-w.

[64] Reul, H. M. and Akdis Mustafa. "Blood pumps for circulatory support". *Perfusion* 15.4 (2000), pp. 295–311. URL: doi:%2010.1177/026765910001500404.

[65] Goldstein, D. J. et al. “Third Annual Report From the ISHLT Mechanically Assisted Circulatory Support Registry: A comparison of centrifugal and axial continuous-flow left ventricular assist devices”. *The Journal of heart and lung transplantation : the official publication of the International Society for Heart Transplantation* 38.4 (2019), pp. 352–363. DOI: 10.1016/j.healun.2019.02.004.

[66] MacLaren, G. et al. “Contemporary extracorporeal membrane oxygenation for adult respiratory failure: life support in the new era”. *Intensive care medicine* 38.2 (2012), pp. 210–220. DOI: 10.1007/s00134-011-2439-2.

[67] Johnson, K. N. et al. “Switching to centrifugal pumps may decrease hemolysis rates among pediatric ECMO patients”. *Perfusion* 37.2 (2022), pp. 123–127. DOI: 10.1177/0267659120982572.

[68] Gülich, J. F. *Kreiselpumpen*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013. ISBN: 978-3-642-40031-5. DOI: 10.1007/978-3-642-40032-2.

[69] Escher, A. et al. “Linking Hydraulic Properties to Hemolytic Performance of Rotodynamic Blood Pumps”. *Advanced Theory and Simulations* 5.9 (2022). ISSN: 2513-0390. DOI: 10.1002/adts.202200117.

[70] Fischer, L. et al. “Impact of extracorporeal blood pump gap sizes on the performance and hemocompatibility under off-design operation”. *Artificial organs* 49.3 (2025), pp. 401–409. DOI: 10.1111/aor.14875.

[71] Yazdanpanah-Ardakani, K. et al. “Optimization of a centrifugal blood pump designed using an industrial method through experimental and numerical study”. *Scientific reports* 14.1 (2024), p. 7443. DOI: 10.1038/s41598-024-57019-9.

[72] Simmonds, M. J. et al. “Will blood-informed design signal the fourth generation of cardiac assist devices?” *The Journal of heart and lung transplantation : the official publication of the International Society for Heart Transplantation* 43.11 (2024), pp. 1767–1770. DOI: 10.1016/j.healun.2024.08.011.

[73] International Medical Device Regulators Forum. *Terms of Reference IMDRF/MC/N1FINAL: 2025(Edition 8)*. URL: https://www.imdrf.org/documents/imdrf-terms-reference (visited on 06/2025).

[74] International Organization for Standardization. *ISO 13485:2016. Medical devices - Quality management systems - Requirements for regulatory purposes*. 2016.

[75] International Organization for Standardization. *ISO 14971:2019. Medical devices — Application of risk management to medical devices*. 2019.

[76] Gregory, S. D. et al., eds. *Mechanical Circulatory and Respiratory Support: Regulatory approvals*. Vol. 2nd. Academic Press, 2024. ISBN: 9780443221477.

[77] Amaral, C. et al. “Global Regulatory Challenges for Medical Devices: Impact on Innovation and Market Access”. *Applied Sciences* 14.20 (2024), p. 9304. DOI: 10.3390/app14209304.

[78] Bretthauer, M. et al. “The New European Medical Device Regulation: Balancing Innovation and Patient Safety”. *Annals of internal medicine* 176.6 (2023), pp. 844–848. DOI: 10.7326/M23-0454.

[79] American Society for Testing and Materials. *Standard Practice for Assessment of Hemolysis in Continuous Flow Blood Pumps*. West Conshohocken, PA. DOI: 10.1520/F1841-19E01.

[80] van Norman, G. A. "Drugs, Devices, and the FDA: Part 2: An Overview of Approval Processes: FDA Approval of Medical Devices". *JACC. Basic to translational science* 1.4 (2016), pp. 277–287. DOI: 10.1016/j.jacbts.2016.03.009.

[81] Viceconti, M. et al. "Possible Contexts of Use for In Silico Trials Methodologies: A Consensus-Based Review". *IEEE journal of biomedical and health informatics* 25.10 (2021), pp. 3977–3982. DOI: 10.1109/JBHI.2021.3090469.

[82] Jose, J. et al. "In Silico Trial Approach for Biomedical Products: A Regulatory Perspective". *Combinatorial chemistry & high throughput screening* 25.12 (2022), pp. 1991–2000. DOI: 10.2174/1386207325666220105150147.

[83] Pappalardo, F. et al. "Toward A Regulatory Pathway for the Use of in Silico Trials in the CE Marking of Medical Devices". *IEEE journal of biomedical and health informatics* 26.11 (2022), pp. 5282–5286. DOI: 10.1109/JBHI.2022.3198145.

[84] American Society of Mechanical Engineers. *Assessing Credibility of Computational Modeling Through Verification and Validation: Application to Medical Devices*. New York, 2018.

[85] Morrison, T. M. et al. "Assessing Computational Model Credibility Using a Risk-Based Framework: Application to Hemolysis in Centrifugal Blood Pumps". *ASAIO journal (American Society for Artificial Internal Organs : 1992)* 65.4 (2019), pp. 349–360. DOI: 10.1097/MAT.0000000000000996.

[86] Gacic, M. et al. "Cost-effectiveness analysis of in silico clinical trials of vascular stents". *2021 IEEE 21st International Conference on Bioinformatics and Bioengineering (BIBE)* (2021), pp. 1–5. URL: 10.1109/BIBE52308.2021.9635321.

[87] Badano, A. "In silico imaging clinical trials: cheaper, faster, better, safer, and more scalable". *Trials* 22.1 (2021), p. 64. DOI: 10.1186/s13063-020-05002-w.

[88] Musuamba, F. T. et al. "Scientific and regulatory evaluation of mechanistic in silico drug and disease models in drug development: Building model credibility". *CPT Pharmacometrics Syst Pharmacol.* 10.8 (2021), pp. 804–825. DOI: 10.1002/psp4.12669.

[89] Segalman, D. J. et al. "Epistemic and Aleatoric Uncertainty in Modeling". *Proceedings of the ASME 2013 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference* Volume 8: 22nd Reliability, Stress Analysis, and Failure Prevention Conference; 25th Conference on Mechanical Vibration and Noise (2013).

[90] Tang, Y.-H. et al. "OpenRBC: A Fast Simulator of Red Blood Cells at Protein Resolution". *Biophysical Journal* 112.10 (2017), pp. 2030–2037. DOI: 10.1016/j.bpj.2017.04.020.

[91] Závodszky, G. et al. "Cellular Level In-silico Modeling of Blood Rheology with An Improved Material Model for Red Blood Cells". *Frontiers in physiology* 8 (2017), p. 563. ISSN: 1664-042X. DOI: 10.3389/fphys.2017.00563.

[92] Malinauskas, R. A. et al. "FDA Benchmark Medical Device Flow Models for CFD Validation". *ASAIO journal* 63.2 (2017), pp. 150–160. DOI: 10.1097/MAT.0000000000000499.

[93] Li, J. et al. "Cytoskeletal dynamics of human erythrocyte". *Proceedings of the National Academy of Sciences* 104.12 (2007), pp. 4937–4942. URL: 10.1073/pnas.0700257104.

[94] Xu, Z. et al. "The erythrocyte destruction mechanism in non-physiological shear mechanical hemolysis". *Physics of Fluids* 34.11 (2022). ISSN: 1070-6631. DOI: 10.1063/5.0112967.

[95] Li, H. et al. "Mechanics of diseased red blood cells in human spleen and consequences for hereditary blood disorders". *Proceedings of the National Academy of Sciences of the United States of America* 115.38 (2018), pp. 9574–9579. DOI: 10.1073/pnas.1806501115.

[96] Kuchel, P. W. et al. "Surface model of the human red blood cell simulating changes in membrane curvature under strain". *Scientific reports* 11.1 (2021), p. 13712. DOI: 10.1038/s41598-021-92699-7.

[97] Zavodszky, G. et al. "Hemocell: a high-performance microscopic cellular library". *Procedia Computer Science* 108 (2017), pp. 159–165. ISSN: 18770509. DOI: 10.1016/j.procs.2017.05.084.

[98] Balachandran Nair, A. N. et al. "Resolved CFD-DEM simulation of blood flow with a reduced-order RBC model". *Computational Particle Mechanics* 9.4 (2022), pp. 759–774. ISSN: 2196-4378. DOI: 10.1007/s40571-021-00441-x.

[99] Giersiepen, M. et al. "Estimation of Shear Stress-related Blood Damage in Heart Valve Prostheses - in Vitro Comparison of 25 Aortic Valves". *Int J Artif Organs* (1990), 13:300–306.

[100] Grigioni, M. et al. "A novel formulation for blood trauma prediction by a modified power-law mathematical model". *Biomechanics and modeling in mechanobiology* 4.4 (2005), pp. 249–260. DOI: 10.1007/s10237-005-0005-y.

[101] Arwatz, G. and Smits, A. J. "A viscoelastic model of shear-induced hemolysis in laminar flow". *Biorheology* 50.1-2 (2013), pp. 45–55. DOI: 10.3233/BIR-130626.

[102] Nikfar, M. et al. "Prediction of mechanical hemolysis in medical devices via a Lagrangian strain-based multiscale model". *Artificial organs* 44.8 (2020), E348–E368. DOI: 10.1111/aor.13663.

[103] Jędrzejczak, K. et al. "Hemolysis of red blood cells in blood vessels modeled via computational fluid dynamics". *International journal for numerical methods in biomedical engineering* 39.11 (2023), e3699. DOI: 10.1002/cnm.3699.

[104] Froese, V. et al. "Experimental validation of the power law hemolysis model using a Couette shearing device". *Artificial organs* 48.5 (2024), pp. 495–503. DOI: 10.1111/aor.14702.

[105] Mei, X. et al. "In vitro study of red blood cell and VWF damage in mechanical circulatory support devices based on blood-shearing platform". *Journal of engineering in medicine* 236.6 (2022), pp. 860–866. DOI: 10.1177/09544119221088420.

[106] Zhang, T. et al. "Study of flow-induced hemolysis using novel Couette-type blood-shearing devices". *Artificial organs* 35.12 (2011), pp. 1180–1186. DOI: 10.1111/j.1525-1594.2011.01243.x.

[107] Ding, J. et al. "Shear-Induced Hemolysis: Species Differences". *Artificial organs* 39.9 (2015), pp. 795–802. DOI: 10.1111/aor.12459.

[108] Heuser, G. and Opitz, R. "A Couette viscometer for short time shearing of blood." *Biorheology* 17:17–24 (1980).

[109] Taskin, M. E. et al. "Evaluation of Eulerian and Lagrangian models for hemolysis estimation". *ASAIO journal* 58.4 (2012), pp. 363–372. DOI: 10.1097/MAT.0b013e318254833b.

[110] Craven, B. A. et al. "A CFD-based Kriging surrogate modeling approach for predicting device-specific hemolysis power law coefficients in blood-contacting medical devices". *Biomechanics and modeling in mechanobiology* 18.4 (2019), pp. 1005–1030. DOI: 10.1007/s10237-019-01126-4.

[111] Torner, B. et al. "Flow simulation-based particle swarm optimization for developing improved hemolysis models". *Biomechanics and modeling in mechanobiology* 22.2 (2023), pp. 401–416. DOI: 10.1007/s10237-022-01653-7.

[112] Kameneva, M. V. et al. "Effects of turbulent stresses upon mechanical hemolysis: experimental and computational analysis". *ASAIO journal* 50(5) (2004), pp. 418–423. DOI: 10.1097/01.mat.0000136512.36370.b5.

[113] Hariharan, P. et al. "Multilaboratory Particle Image Velocimetry Analysis of the FDA Benchmark Nozzle Model to Support Validation of Computational Fluid Dynamics Simulations". *ASME. J Biomech Eng* 133 (2011), 133(4): 041002. URL: https://doi.org/10.1115/1.4003440.

[114] Garon, A. and Farinas, M.-I. "Fast Three-dimensional Numerical Hemolysis Approximation". *Artificial organs* 28(11):1016–1025 (2004).

[115] Monaghan, J. J. "Smoothed particle hydrodynamics". *Reports on Progress in Physics* 68.8 (2005), pp. 1703–1759. ISSN: 0034-4885. DOI: 10.1088/0034-4885/68/8/R01.

[116] Faghih, M. M. et al. "Practical implications of the erroneous treatment of exposure time in the Eulerian hemolysis power law model". *Artificial organs* 47.9 (2023), pp. 1531–1538. DOI: 10.1111/aor.14543.

[117] Faghih, M. M. and Sharp, M. K. "Modeling and prediction of flow-induced hemolysis: a review". *Biomechanics and modeling in mechanobiology* 18.4 (2019), pp. 845–881. DOI: 10.1007/s10237-019-01137-1.

[118] Goubergrits, L. and Affeld, K. "Numerical estimation of blood damage in artificial organs". *Artificial organs* 28.5 (2004), pp. 499–507. DOI: 10.1111/j.1525-1594.2004.07265.x.

[119] Sharp, M. K. and Mohammad, S. F. "Scaling of Hemolysis in Needles and Catheters". *Annals of biomedical engineering* 26.5 (1998), pp. 788–797. URL: 10.1114/1.65.

[120] Yu, H. et al. “A Review of Hemolysis Prediction Models for Computational Fluid Dynamics”. *Artificial organs* 41.7 (2017), pp. 603–621. DOI: 10.1111/aor.12871.

[121] Faghih, M. M. and Keith Sharp, M. “Extending the Power-Law Hemolysis Model to Complex Flows”. *Journal of biomechanical engineering* 138.12 (2016). DOI: 10.1115/1.4034786.

[122] Arora, D. et al. “A tensor-based measure for estimating blood damage”. *Artificial organs* 28.11 (2004), pp. 1002–1015. DOI: 10.1111/j.1525-1594.2004.00072.x.

[123] Pauli, L. et al. “Transient stress-based and strain-based hemolysis estimation in a simplified blood pump”. *International journal for numerical methods in biomedical engineering* 29.10 (2013), pp. 1148–1160. DOI: 10.1002/cnm.2576.

[124] Dirkes, N. et al. “Eulerian formulation of the tensor-based morphology equations for strain-based blood damage modeling”. *Computer Methods in Applied Mechanics and Engineering* 426 (2024), p. 116979. ISSN: 00457825. DOI: 10.1016/j.cma.2024.116979.

[125] Lommel, M. et al. “Experimental Investigation of the Applicability of the Stress-Based and Strain-Based Hemolysis Models for Short-Term Stress Peaks Typical for Rotary Blood Pumps”. *Artificial organs* (2025). DOI: 10.1111/aor.15002.

[126] Wu, P. et al. “An energy-dissipation-based power-law formulation for estimating hemolysis”. *Biomechanics and modeling in mechanobiology* 19.2 (2020), pp. 591–602. DOI: 10.1007/s10237-019-01232-3.

[127] Faghih, M. M. and Sharp, M. K. “Evaluation of energy dissipation rate as a predictor of mechanical blood damage”. *Artificial organs* 43.7 (2019), pp. 666–676. DOI: 10.1111/aor.13418.

[128] Yu, J. et al. “Non-intrusive reduced-order modeling for fluid problems: A brief review”. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering* 233.16 (2019), pp. 5896–5912. ISSN: 0954-4100. DOI: 10.1177/0954410019890721.

[129] Padula, G. et al. “A brief review of reduced order models using intrusive and non–intrusive techniques”. *PAMM* 24.4 (2024). ISSN: 1617-7061. DOI: 10.1002/pamm.202400210.

[130] Song, C. and Heuveline, V. “Uncertainty Assessment of the Blood Damage in a FDA Blood Pump”. *Proc. 3rd ECCOMAS Thematic Conference on Uncertainty Quantification in Computational Sciences and Engineering* (2019). URL: doi:%2010.7712/120219.6366.18426.

[131] Hosder, S. et al. “Efficient Sampling for Non-Intrusive Polynomial Chaos Applications with Multiple Uncertain Input Variables”. *Proc. 48th AIAA/ASME/ASCE/AHS/ASC Struct. Structural Dyn. Mater. Conf.* (2007). URL: 10.2514/6.2007-1939.

[132] Friedman, J. H. “Multivariate Adaptive Regression Splines”. *The Annals of Statistics* 19.1 (1991), pp. 1–141.

[133] Wiener, N. “The Homogeneous Chaos”. *American Journal of Mathematics* 60.4 (1938), pp. 897–936. URL: https://doi.org/10.2307/2371268.

[134] Wang, J. "An Intuitive Tutorial to Gaussian Process Regression". *Computing in Science & Engineering* 25.4 (2023), pp. 4–11. ISSN: 1521-9615. DOI: 10.1109/MCSE.2023.3342149.

[135] Arora, G. et al. "A review of radial basis function with applications explored". *Journal of the Egyptian Mathematical Society* 31.1 (2023). DOI: 10.1186/s42787-023-00164-3.

[136] Smola, A. J. and Schölkopf, B. "A tutorial on support vector regression". *Statistics and Computing* 14 (2004), pp. 199–222.

[137] Rumelhart, D. E. et al. "Learning representations by back-propagating errors". *Nature* Vol 323 (1986), pp. 533–536.

[138] Hinton, G. E. and Salakhutdinov, R. R. "Reducing the Dimensionality of Data with Neural Networks". *Science (New York, N.Y.)* 313.5786 (2006), pp. 504–507. DOI: 10.1126/science.1127647.

[139] Cheng, K. et al. "Surrogate-assisted global sensitivity analysis: an overview". *Struct Multidisc Optim* 61.3 (2020), pp. 1187–1213. DOI: 10.1007/s00158-019-02413-5.

[140] Herbertson, L. H. et al. "Multilaboratory study of flow-induced hemolysis using the FDA benchmark nozzle model". *Artificial organs* 39.3 (2015), pp. 237–248. DOI: 10.1111/aor.12368.

[141] Magni, P. et al. "Bayesian function learning using MCMC methods". *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20.12 (1998), pp. 1319–1331. ISSN: 01628828. DOI: 10.1109/34.735805.

[142] Murphy, K. P. *Machine learning: A probabilistic perspective*. Adaptive computation and machine learning series. Cambridge MA: MIT Press, 2012. ISBN: 9780262018029.

[143] Chib, S. and Greenberg, E. "Understanding the Metropolis-Hastings Algorithm". *The American Statistician* 49.4 (1995), pp. 327–335. URL: https://doi.org/10.2307/2684568.

[144] Hoffman, M. and Gelman, A. "The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo". *Journal of Machine Learning Research* 15 (2014), pp. 1593–1623.

[145] Blum, C. et al. "Systematic analysis of non-intrusive polynomial chaos expansion to determine rotary blood pump performance over the entire operating range". *Computers in biology and medicine* 168 (2024), p. 107772. DOI: 10.1016/j.compbiomed.2023.107772.

[146] Hao, M. et al. "Uncertainty Quantification on the Cooling Performance of a Transonic Turbine Vane With Upstream Endwall Misalignment". *Journal of Turbomachinery* 144.12 (2022). ISSN: 0889-504X. DOI: 10.1115/1.4055269.

[147] Tao, Z. et al. "Uncertainty Quantification of Aero-Thermal Performance of a Blade Endwall Considering Slot Geometry Deviation and Mainstream Fluctuation". *Journal of Turbomachinery* 143.11 (2021). ISSN: 0889-504X. DOI: 10.1115/1.4051416.

[148] Salehi, S. et al. "On the flow field and performance of a centrifugal pump under operational and geometrical uncertainties". *Applied Mathematical Modelling* 61 (2018), pp. 540–560. ISSN: 0307904X. DOI: 10.1016/j.apm.2018.05.008.

[149] Karimi, M. S. et al. “Probabilistic CFD computations of gas turbine vane under uncertain operational conditions”. *Applied Thermal Engineering* 148 (2019), pp. 754–767. ISSN: 13594311. DOI: 10.1016/j.applthermaleng.2018.11.072.

[150] Mohammadi, R. et al. “Probabilistic CFD analysis on the flow field and performance of the FDA centrifugal blood pump”. *Applied Mathematical Modelling* 109 (2022), pp. 555–577. ISSN: 0307904X. DOI: 10.1016/j.apm.2022.05.016.

[151] Hariharan, P. et al. “Inter-Laboratory Characterization of the Velocity Field in the FDA Blood Pump Model Using Particle Image Velocimetry (PIV)”. *Cardiovascular engineering and technology* 9.4 (2018), pp. 623–640. DOI: 10.1007/s13239-018-00378-y.

[152] Gross-Hardt, S. et al. “Low-flow assessment of current ECMO/ECCO2R rotary blood pumps and the potential effect on hemocompatibility”. *Critical Care* 23.1 (2019), p. 348. DOI: 10.1186/s13054-019-2622-3.

[153] McKay, M. D. et al. “Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code”. *Technometrics* 21.2 (1979), pp. 239–245. ISSN: 0040-1706. DOI: 10.1080/00401706.1979.10489755.

[154] Metropolis, N. and Ulam, S. “The Monte Carlo Method”. *Journal of the American Statistical Association* 44 (1949), pp. 335–341. ISSN: 0162-1459. URL: 10.2307/2280232.

[155] Xiu, D. and Karniadakis, G. E. “The Wiener–Askey PolynomialChaos for Stochastic Differential Equations”. *SIAM J. SCI. COMPUT* Vol. 24, No. 2 (2002), pp. 619–644.

[156] Feinberg, J. and Langtangen, H. P. “Chaospy: An open source tool for designing methods of uncertainty quantification”. *Journal of Computational Science* 11 (2015), pp. 46–57. ISSN: 18777503. DOI: 10.1016/j.jocs.2015.08.008.

[157] Harris, C. R. et al. “Array programming with NumPy”. *Nature* 585.7825 (2020), pp. 357–362. DOI: 10.1038/s41586-020-2649-2.

[158] McKinney, W. “Data Structures for Statistical Computing in Python”. *Proceedings of the 9th Python in Science Conf.* (2010), pp. 56–61. URL: 10.25080/Majora-92bf1922-00a.

[159] Virtanen, P. et al. “SciPy 1.0: fundamental algorithms for scientific computing in Python”. *Nature methods* 17.3 (2020), pp. 261–272. DOI: 10.1038/s41592-019-0686-2.

[160] Hunter, J. D. “Matplotlib: A 2D graphics environment”. *Computing in Science & Engineering* 3 (2017), pp. 90–95. URL: 10.1109/MCSE.2007.55.

[161] Blum, C. et al. “Blood trauma in veno-venous extracorporeal membrane oxygenation: low pump pressures and low circuit resistance matter”. *Critical care (London, England)* 28.1 (2024), p. 330. DOI: 10.1186/s13054-024-05121-9.

[162] Combes, A. et al. “ECMO for severe ARDS: systematic review and individual patient data meta-analysis”. *Intensive care medicine* 46.11 (2020), pp. 2048–2057. DOI: 10.1007/s00134-020-06248-3.

[163] Grasselli, G. et al. “ESICM guidelines on acute respiratory distress syndrome: definition, phenotyping and respiratory support strategies”. *Intensive care medicine* 49.7 (2023), pp. 727–759. DOI: 10.1007/s00134-023-07050-7.

[164] Qadir, N. et al. "An Update on Management of Adult Patients with Acute Respiratory Distress Syndrome: An Official American Thoracic Society Clinical Practice Guideline". *American journal of respiratory and critical care medicine* 209.1 (2024), pp. 24–36. DOI: 10.1164/rccm.202311-2011ST.

[165] Halaweish, I. et al. "Roller and Centrifugal Pumps: A Retrospective Comparison of Bleeding Complications in Extracorporeal Membrane Oxygenation". *ASAIO journal* 61.5 (2015), pp. 496–501. DOI: 10.1097/MAT.0000000000000243.

[166] Costantini, S. et al. "A retrospective analysis of the hemolysis occurrence during extracorporeal membrane oxygenation in a single center". *Perfusion* 38.3 (2023), pp. 609–621. DOI: 10.1177/02676591211073768.

[167] Sniderman, J. et al. "Hematologic concerns in extracorporeal membrane oxygenation". *Research and practice in thrombosis and haemostasis* 4.4 (2020), pp. 455–468. DOI: 10.1002/rth2.12346.

[168] Omar, H. R. et al. "Plasma Free Hemoglobin Is an Independent Predictor of Mortality among Patients on Extracorporeal Membrane Oxygenation Support". *PloS one* 10.4 (2015), e0124034. DOI: 10.1371/journal.pone.0124034.

[169] Gross-Hardt, S. H. et al. "Crucial Aspects for Using Computational Fluid Dynamics as a Predictive Evaluation Tool for Blood Pumps". *ASAIO journal* 65.8 (2019), pp. 864–873. DOI: 10.1097/MAT.0000000000001023.

[170] Schöps, M. et al. "Hemolysis at low blood flow rates: in-vitro and in-silico evaluation of a centrifugal blood pump". *Journal of translational medicine* 19.1 (2021), p. 2. DOI: 10.1186/s12967-020-02599-z.

[171] Ki, K. K. et al. "Low flow rate alters haemostatic parameters in an ex-vivo extracorporeal membrane oxygenation circuit". *Intensive care medicine experimental* 7.1 (2019), p. 51. ISSN: 2197-425X. DOI: 10.1186/s40635-019-0264-z.

[172] Davies, L. and Gather, U. "The Identification of Multiple Outliers". *Journal of the American Statistical Association* 88.423 (1993), pp. 782–792. ISSN: 0162-1459. DOI: 10.1080/01621459.1993.10476339.

[173] Fraser, K. H. et al. "A quantitative comparison of mechanical blood damage parameters in rotary ventricular assist devices: shear stress, exposure time and hemolysis index". *Journal of biomechanical engineering* 134.8 (2012), p. 081002. DOI: 10.1115/1.4007092.

[174] Raposo, P. "Scale-specific automated line simplification by vertex clustering on a hexagonal tessellation". *Cartography and Geographic Information Science* 40.5 (2013), pp. 427–443. ISSN: 1523-0406. DOI: 10.1080/15230406.2013.803707.

[175] Barbič, B. et al. "The Failure of Extracorporeal Carbon Dioxide Removal May Be a Failure of Technology". *American journal of respiratory and critical care medicine* 209.7 (2024), pp. 884–887. DOI: 10.1164/rccm.202309-1628LE.

[176] Materne, L. A. et al. "Hemolysis in patients with Extracorporeal Membrane Oxygenation therapy for severe Acute Respiratory Distress Syndrome - a systematic review of the literature". *International journal of medical sciences* 18.8 (2021), pp. 1730–1738. DOI: 10.7150/ijms.50217.

[177] Kutleša, M. et al. “Extracorporeal membrane oxygenation treatment for H1N1-induced acute respiratory distress syndrome (ARDS): results of the Croatian Referral Center for Respiratory ECMO”. *The International journal of artificial organs* 37.10 (2014), pp. 748–752. DOI: 10.5301/ijao.5000356.

[178] Lehle, K. et al. “Flow dynamics of different adult ECMO systems: a clinical evaluation”. *Artificial organs* 38.5 (2014), pp. 391–398. DOI: 10.1111/aor.12180.

[179] Blum, C. et al. “Quantifying Experimental Variability in Shear-Induced Hemolysis to Support Uncertainty-Aware Hemolysis Models”. *Annals of biomedical engineering* (2025). DOI: 10.1007/s10439-025-03786-z.

[180] Lenormand, G. et al. “Direct Measurement of the Area Expansion and Shear Moduli of the Human Red Blood Cell Membrane Skeleton”. *Biophysical Journal* Vol. 81 (2001), pp. 43–56.

[181] Liu, X. et al. “The measurement of shear modulus and membrane surface viscosity of RBC membrane with Ektacytometry: a new technique”. *Mathematical biosciences* 209.1 (2007), pp. 190–204. ISSN: 0025-5564. DOI: 10.1016/j.mbs.2006.09.026.

[182] Hénon, S. et al. “A New Determination of the Shear Modulus of the Human Erythrocyte Membrane Using Optical Tweezers”. *Biophysical Journal* Vol. 76 (1999), pp. 1145–1151.

[183] Chan, C. H. H. et al. “Discrete responses of erythrocytes, platelets, and von Willebrand factor to shear”. *Journal of biomechanics* 130 (2022), p. 110898. DOI: 10.1016/j.jbiomech.2021.110898.

[184] Yen, J.-H. et al. “The effect of turbulent viscous shear stress on red blood cell hemolysis”. *Journal of artificial organs* 17.2 (2014), pp. 178–185. DOI: 10.1007/s10047-014-0755-3.

[185] Zhao, R. et al. “Microscopic investigation of erythrocyte deformation dynamics”. *Biorheology* 43.6 (2006), pp. 747–765. DOI: 10.1177/0006355X2006043006003.

[186] Woelke, E. et al. “Miniaturized Test Loop for the Assessment of Blood Damage by Continuous-Flow Left-Ventricular Assist Devices”. *Annals of biomedical engineering* 48.2 (2020), pp. 768–779. DOI: 10.1007/s10439-019-02404-z.

[187] Thamsen, B. et al. “Assessment of the Flow Field in the HeartMate 3 Using Three-Dimensional Particle Tracking Velocimetry and Comparison to Computational Fluid Dynamics”. *ASAIO journal* 66.2 (2020), pp. 173–182. DOI: 10.1097/MAT.0000000000000987.

[188] Chan, C. H. H. et al. “Shear Stress-Induced Total Blood Trauma in Multiple Species”. *Artificial organs* 41.10 (2017), pp. 934–947. DOI: 10.1111/aor.12932.

[189] International Council for Standardization in Haematology: Expert Panelon Cytometry. “Recommendations of the International Council for Standardization in Haematology for Ethylenediaminetetraacetic Acid Antico-agulation of Blood for Blood Cell Counting and Sizing”. *Am. J. Clin. Pathol.* 100 (1993), pp. 371–372.

[190] Chien, S. “Shear Dependence of Effective Cell Volume as a Determinant of Blood Viscosity”. 168(3934) (1970), pp. 977–979. DOI: 10.1126/science.168.3934.977.

[191] McMillan, D. E. et al. “Rapidly recovered transient flow resistance: a newly discovered property of blood”. *Am J Physiol.* 4 Pt 2 (1987), H919–926. DOI: 10.1152/ajpheart.1987.253.4.H919..

[192] Tapernon, K. et al. “Quality control of hemolysis rate of erythrocyte concentrates: a proficiency test for determination of free hemoglobin.” *Anasthesiol Intensivmed Notfallmed Schmerzther* 36 (Suppl 1) (2001), pp. 45–50. URL: 10.1055/s-2001-11838.

[193] Petersdorff-Campen, K. von and Schmid Daners, M. “Hemolysis Testing In Vitro: A Review of Challenges and Potential Improvements”. *ASAIO journal* 68.1 (2022), pp. 3–13. DOI: 10.1097/MAT.0000000000001454.

[194] Braune, S. et al. “Evaluation of platelet adhesion and activation on polymers: Round-robin study to assess inter-center variability”. *Colloids and surfaces. B, Biointerfaces* 158 (2017), pp. 416–422. DOI: 10.1016/j.colsurfb.2017.06.053.

[195] Blum, C. et al. “Toward Uncertainty-Aware Hemolysis Modeling: A Universal Approach to Address Experimental Variance”. *International journal for numerical methods in biomedical engineering* 41.5 (2025), e70040. DOI: 10.1002/cnm.70040.

[196] Gesenhues, L. et al. “Strain-based blood damage estimation for computational design of ventricular assist devices”. *The International journal of artificial organs* 39.4 (2016), pp. 166–170. DOI: 10.5301/ijao.5000484.

[197] Chen, Y. and Sharp, M. K. “A strain-based flow-induced hemolysis prediction model calibrated by in vitro erythrocyte deformation measurements”. *Artificial organs* 35.2 (2011), pp. 145–156. DOI: 10.1111/j.1525-1594.2010.01050.x.

[198] Speagle, J. S. “A Conceptual Introduction to Markov Chain Monte Carlo Methods” (2020). URL: http://arxiv.org/pdf/1909.12313v2.

[199] Hartig, F. et al. “Statistical inference for stochastic simulation models–theory and application”. *Ecology letters* 14.8 (2011), pp. 816–827. DOI: 10.1111/j.1461-0248.2011.01640.x.

[200] Abril-Pla, O. et al. “PyMC: a modern, and comprehensive probabilistic programming framework in Python”. *PeerJ. Computer science* 9 (2023), e1516. DOI: 10.7717/peerj-cs.1516.

[201] Pedregosa, F. et al. “Scikit-learn: Machine learning in Python”. *Journal of Machine Learning Research* 12.Oct (2011), pp. 2825–2830.

[202] Mantegazza, A. et al. “Examining the universality of the hemolysis power law model from simulations of the FDA nozzle using calibrated model coefficients”. *Biomechanics and modeling in mechanobiology* 22.2 (2023), pp. 433–451. DOI: 10.1007/s10237-022-01655-5.