

# Spatio-Temporal Interface Reconstruction by Means of Glare Points and Deep Learning

Zur Erlangung des akademischen Grades eines  
**Doktors der Ingenieurwissenschaften (Dr.-Ing.)**  
von der KIT-Fakultät für Maschinenbau  
des Karlsruher Instituts für Technologie (KIT)

genehmigte

**Dissertation**

von

**M.Sc. Maximilian Dreisbach**

Tag der mündlichen Prüfung:  
Hauptreferentin:  
Korreferent:

24.02.2025  
Prof. Dr.-Ing. Bettina Frohnäpfel  
T.T.-Prof. Dr. Pascal Friederich





Dieses Werk ist lizenziert unter einer Creative Commons Namensnennung 4.0 International Lizenz (CC BY 4.0): <https://creativecommons.org/licenses/by/4.0/deed.de>



# Abstract

This thesis explores the potential of deep learning techniques for the spatio-temporal reconstruction of gas-liquid interfaces in two-phase flows based on optical measurements. The inherently three-dimensional (3D) nature of two-phase flows currently renders single-view reconstruction techniques impractical due to ambiguities in the projection from 3D interfaces onto 2D images, especially for deformed and self-occluding interfaces. To address this challenge, this thesis proposes a novel framework for the reconstruction of gas-liquid interfaces based on neural networks and a purposefully developed optical measurement technique. The canonical shadowgraphy setup is extended by glare points from lateral light sources to encode additional 3D information on the interface in the images. A convolutional neural network (CNN) architecture is employed to extract spatially-aligned, context-rich features from these images, which are then used for interface reconstruction. During the course of the present investigation, feature visualization techniques and a comparative experimental study revealed that the neural network successfully leverages the additional information in the images to achieve highly accurate interface reconstruction. Based on the insight that the reconstruction of the neural network closely adheres to the shadowgraphy contour, it can be concluded that the combination of the high spatial accuracy of the shadowgraphy technique and the spatially-aligned feature extraction by the CNN facilitates a precise in-plane reconstruction. Moreover, experiments involving substrates with unknown wetting properties uncovered that this interplay enables the generalization to unfamiliar interface geometries and varying fluid mechanical conditions in the experiments. These findings demonstrate that substantial synergies can be unlocked through the alignment and purposeful co-development of experimental measurement methods and data-driven algorithms. To address the ill-posed nature of monocular reconstruction, the data-driven optimization of the neural network is regularized by incorporating prior knowledge from numerical simulations and analytical models. Synthetic training data generation through physically-based rendering is employed to enable supervised learning from data obtained by numerical simulation. This strategy is found to provide an effective way to implicitly train the neural network to adhere to the physics of the two-phase flow, as demonstrated by the reliable reconstruction of hidden and self-obscured interfaces. To further enhance the framework, physics-informed neural networks (PINNs) are integrated to directly incorporate the knowledge of the underlying physics of two-phase flows in the data-driven reconstruction. This approach significantly improves the accuracy and temporal consistency of the interface reconstruction, while also providing the additional capability for the prediction of velocity and pressure fields in both phases. These findings demonstrate that additional prior knowledge from encoded governing equations facilitates the training of physically accurate models of two-phase flows. In conclusion, the insights uncovered in this thesis highlight the potential of deep learning techniques to seamlessly integrate experimental, numerical, and analytical domain knowledge, outlining promising strategies to address previously unsolved challenges in the reconstruction of dynamic interfaces in two-phase flows.



# Kurzfassung

In der vorliegenden Arbeit wird das Potenzial von Deep-Learning-Techniken für die raum-zeitliche Rekonstruktion von Gas-Flüssigkeits-Grenzflächen in Zweiphasenströmungen auf Grundlage optischer Messverfahren untersucht. Die dreidimensionale (3D) Natur von Zweiphasenströmungen stellt eine erhebliche Herausforderung für monokulare Rekonstruktionstechniken dar, da die Projektion der 3D-Grenzflächen auf 2D-Bilder zu Mehrdeutigkeiten führt, insbesondere bei deformierten oder selbstverdeckenden Grenzflächen. Um diese grundlegende Limitation zu überwinden, wird in der vorliegenden Arbeit ein innovativer Ansatz für die Rekonstruktion von Gas-Flüssigkeits-Grenzflächen vorgestellt, der neuronale Netzwerke mit einer speziell entwickelten optischen Messtechnik kombiniert. Das traditionelle Schattenbildverfahren wird hierbei durch Glanzpunkte seitlicher Lichtquellen erweitert, wodurch zusätzliche 3D-Informationen über die Grenzfläche in den aufgenommenen Bildern kodiert werden. Eine Netzwerkarchitektur auf Basis sogenannter Convolutional Neural Networks (CNNs) wird verwendet, um räumlich orientierte und kontextreiche Bildmerkmale zu extrahieren, die anschließend für die Grenzflächenrekonstruktion eingesetzt werden. Mittels Visualisierungstechniken und einer vergleichenden experimentellen Studie wird gezeigt, dass das neuronale Netzwerk diese zusätzlichen Bildinformationen erfolgreich für eine präzise Rekonstruktion der Grenzfläche nutzt. Die Erkenntnis, dass die Rekonstruktion des neuronalen Netzwerks exakt der Schattenbildkontur folgt, belegt, dass die Kombination der hohen räumlichen Genauigkeit der Schattenbildtechnik und der räumlich orientierten Merkmalsextraktion durch das CNN eine präzise Rekonstruktion in der Bildebene ermöglicht. Darüber hinaus zeigen Experimente mit Substraten unbekannter Benetzbarkeit, dass dieses Zusammenspiel der Methoden eine Generalisierung auf unbekannte Grenzflächengeometrien und variierende fluidmechanische Bedingungen in den Experimenten ermöglicht. Diese Ergebnisse verdeutlichen, dass durch die Abstimmung und gezielte gemeinsame Entwicklung von experimentellen Messmethoden und datengestützten Algorithmen erhebliche Synergien freigesetzt werden können. Das schwer lösbare Problem der monokularen Rekonstruktion wird zusätzlich durch die Einbindung von Domänenwissen aus numerischen Simulationen und analytischen Modellen adressiert, um die datengestützte Optimierung des neuronalen Netzwerks zu regulieren. Hierbei wird die Erzeugung synthetischer Trainingsdaten durch physikalisch-basiertes Rendering eingesetzt, um überwachtes Lernen mit Daten aus numerischen Simulationen zu ermöglichen. Diese Strategie erwies sich als effektive Möglichkeit, um das neuronale Netzwerk implizit auf die Einhaltung der Physik von Zweiphasenströmungen zu trainieren – wie durch die zuverlässige Rekonstruktion verborgener und selbstverdeckender Grenzflächen demonstriert werden kann. Zur weiteren Optimierung der Rekonstruktionsmethode werden Physik-informierte neuronale Netzwerke (PINNs) integriert, um die zugrundeliegende Erhaltungsgleichungen der Zweiphasenströmungen direkt in die datengestützte Rekonstruktion einzubeziehen. Dieser Ansatz führt zu einer signifikanten Verbesserung der Genauigkeit und zeitlichen Konsistenz der Grenzflächenrekonstruktion und ermöglicht darüber hinaus die Vorhersage von Geschwindigkeits- und Druckfeldern in beiden Phasen. Diese Ergebnisse verdeutlichen, dass die Einbeziehung von Domänenwissen aus den zugrundeliegenden Erhaltungsgleichungen das Training physikalisch präziser Modelle für Zweiphasenströmungen erheblich verbessert. Zusammenfassend veranschaulichen die in dieser Dissertation gewonnenen Erkenntnisse das Potenzial von Deep-Learning-Techniken für die nahtlose Integration von Domänenwissen aus Experimenten,

numerischer Simulation und analytischen Modellen, was erfolgversprechende Strategien zur Bewältigung ungelöster Herausforderungen bei der Rekonstruktion dynamischer Grenzflächen in Zweiphasenströmungen bietet.

# Acknowledgements

I would like to express my deepest gratitude to those who have supported, guided, and inspired me throughout my PhD research. First and foremost, I would like to thank my doctoral advisor, Bettina Frohnapfel, for providing an exceptional research environment that gave me the freedom to explore new ideas combined with steady guidance keeping the greater scope in check. I am sincerely thankful to Pascal Friederich for the examination of this thesis and for his insightful comments, which greatly enriched my work. I am immensely grateful to my supervisors, Jochen Kriegseis and Alexander Stroh, for their invaluable contributions and support. Your time and dedication throughout numerous insightful discussions over the years and your availability to answer my many questions have been integral to this work. Thank you, Jochen, for the patient assistance in setting up my initial experiments and sharing your expert knowledge, which was key in advancing the optical measurement methods in this thesis. And thank you, Alex, for your guidance on numerical methods for two-phase flows and your encouragement to pursue innovative ideas in machine learning. The inspiration I drew from your combined supervision was instrumental in advancing my research.

I am grateful to Sebastian Burgmann for our enjoyable and productive collaboration. Special thanks also go to Sepideh Khodaparast and Francis Dent for their warm hospitality and the fruitful collaboration during my research visit to their lab at the University of Leeds. I am deeply thankful to George Karniadakis and the members of the Crunch Group for their warm hospitality during my research stay at Brown University. Your openness, collaborative spirit, and expertise in physics-informed machine learning techniques were inspiring and made this experience invaluable. I am especially grateful to Elham Kiyani for the enriching discussions during and beyond this time. Special thanks to my colleagues in the workshop, Daniel Kissenkötter, Dominik Eikerling, Heiko Bendler, Marcel Stotz, Michael Abramov, and Otto Vorreiter. Thanks to my direct colleagues, Agnes, Andrea, Annika, Arjun, Carola, Christian, David, David, Davide, Dominic, Erik, Francesco, Franco, Gabriel, Georg, Heinrich, Jacopo, Jiasheng, Joachim, Jonathan, Jonathan, Karin, Kay, Lars, Marc, Marco, Moritz, Naseem, Nima, Patricia, Saskia, Sebastian, Sibela, Simon, Stefano, Robin, Tobias, Vishnu, and Wiebke for the joyful and inspiring discussions, as well as the welcoming atmosphere you created - I truly enjoyed my time at ISTM thanks to you! I gratefully acknowledge the financial support provided by the Friedrich and Elisabeth Boysen Foundation (BOY-160) and the Karlsruhe House of Young Scientists (KHYS), which made this work possible.

On a personal note, I want to thank my parents, Brigitte and Dietrich, for their unwavering support and for fostering my curiosity from an early age. To my aunt Heike and uncle Jörg, and my siblings Carl, Jan, and Julia, thank you for your guidance and encouragement. I must also mention my dog, Taquito, who brought joy to everyone in the laboratory and reminded me to take regular breaks. To everyone who contributed to this journey, directly or indirectly, thank you. This work would not have been possible without you.





# Contents

<b>Abstract</b>	<b>i</b>
<b>Kurzfassung</b>	<b>iii</b>
<b>Acknowledgements</b>	<b>v</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Motivation	2
1.2 Research objective and outline of the study	3
<b>2 Fundamentals</b>	<b>5</b>
2.1 Governing equations of fluid mechanics	5
2.2 Two-phase flow phenomena	7
2.2.1 Wetting of smooth and structured surfaces	9
2.2.2 Droplet impingement	11
2.2.3 Adhering droplets in shear flows	12
2.2.4 Numerical modeling of two-phase flows	13
2.3 Optical measurement principles	18
2.3.1 Shadowgraphy technique	18
2.3.2 Glare points	19
2.4 Deep Learning	21
2.4.1 Neural networks	22
2.4.2 Optimization techniques	24
2.4.3 Convolutional neural networks	25
2.4.4 Physics-informed neural networks	28
<b>3 Development of the reconstruction framework</b>	<b>31</b>
3.1 Review of two-phase flow reconstruction	31
3.2 Review of deep learning techniques for volumetric reconstruction	33
3.3 Reconstruction framework	35
3.4 Neural network architecture and optimization procedure	40
<b>4 Glare-point shadowgraphy</b>	<b>43</b>
4.1 Experimental setup	43
4.2 Image processing	44
4.3 Validation	47
4.3.1 Spectral reconstruction of image data	47
4.3.2 Glare point prediction accuracy	49
4.3.3 Influence of the elevation angle	52
4.4 Discussion	53
4.4.1 Optimal position of the lateral lights	53
4.4.2 Three-dimensional encoding of the gas-liquid interface	54
4.4.3 Limitations of the method	56
4.5 Concluding remarks regarding glare point shadowgraphy	58

<b>5</b>	<b>Synthetic training data generation</b>	<b>61</b>
5.1	Problem statement	61
5.2	Synthetic image generation	62
5.3	Validation	64
5.4	Training datasets	67
5.5	Concluding remarks regarding synthetic training data generation	68
<b>6</b>	<b>Spatio-temporal reconstruction of gas-liquid interfaces in droplet flows</b>	<b>69</b>
6.1	Methodology	69
6.1.1	Evaluation procedure	69
6.1.2	Experimental setup	71
6.1.3	Data processing and network training	74
6.1.4	Evaluation metrics	75
6.2	Validation	76
6.2.1	Training and validation on synthetic image data	77
6.2.2	Feature visualization	79
6.3	Results and Discussion	81
6.3.1	Axisymmetrical droplet impingement	81
6.3.2	Non-axisymmetrical droplet impingement	83
6.3.3	Reconstruction of highly deformed gas-liquid interfaces	86
6.3.4	Reconstruction of obscured and hidden shapes	89
6.3.5	Adhering droplets in external shear flow	90
6.3.6	Extrapolation to unknown droplet dynamics	94
6.3.7	Surface quality of the reconstructed interfaces	97
6.4	Limitations of the method	99
6.5	Concluding remarks regarding the volumetric reconstruction of droplet flows by means of deep learning	100
<b>7</b>	<b>Physics-informed learning for spatio-temporal droplet reconstruction</b>	<b>105</b>
7.1	Physics-informed neural networks	105
7.1.1	PINNs for fluid mechanics	105
7.1.2	PINNs for two-phase flows	106
7.2	Methodology	107
7.2.1	Governing equations	109
7.2.2	Loss weighting	110
7.2.3	Sampling methods	111
7.2.4	Neural network architecture	112
7.2.5	Evaluation metrics	113
7.3	Results and Discussion	113
7.3.1	Validation on synthetic data	113
7.3.2	Reconstruction of experimental data	120
7.4	Concluding remarks regarding physics-informed learning of droplet dynamics	123
<b>8</b>	<b>Conclusions and Outlook</b>	<b>127</b>
<b>9</b>	<b>References</b>	<b>131</b>
	Bibliography	131
	List of Conference Contributions	147
	List of Journal Publications	148
	List of Preprints	148
	Code and Data Repositories	149
	Co-Supervised Theses	149
	<b>Acronyms and symbols</b>	<b>153</b>

---

<b>List of Figures</b> . . . . .	<b>161</b>
<b>List of Tables</b> . . . . .	<b>165</b>
<b>A Appendix: Glare-point shadowgraphy</b> . . . . .	<b>167</b>
A.1 Analytical glare-point derivation . . . . .	167
A.2 Time series of RGB images . . . . .	169
<b>B Appendix: Synthetic image rendering</b> . . . . .	<b>173</b>
B.1 Capability of the rendering setup as a digital twin of the experiment . . . . .	173
<b>C Appendix: Spatio-temporal reconstruction</b> . . . . .	<b>175</b>
C.1 Adhering droplets subjected to shear flows . . . . .	175
C.2 Droplet impingement on structured substrates . . . . .	177
C.3 Physics-informed neural networks . . . . .	180
<b>D Appendix: Code and data repositories</b> . . . . .	<b>183</b>



# 1 Introduction

Multiphase flows involve two or more thermodynamic phases and are ubiquitous in nature, ranging from tiny droplets of rain to rushing rivers and waterfalls that shape the face of the Earth. Moreover, multiphase flows form the underlying process in a large proportion of technical applications, ranging from the circulation of steam and water in nearly all thermal electric power stations [195, 32] to the sloshing of cryogenic fuels in the tanks of spacecraft [134]. The ever-increasing pace of technological development since the Industrial Revolution has led to prosperity and a widespread improvement in living conditions. However, it is becoming increasingly apparent that the unsustainable use of limited resources is likely to cause significant problems in the future. In particular, the use of fossil fuels entails the emission of greenhouse gases, which are the main driver of climate change that will lead to an increase in the occurrence and severity of extreme weather events, increasing sea levels, and a significant reduction in biodiversity [58, 209, 193]. Furthermore, the pollution from the combustion of fossil fuels causes a reduction in air quality with negative effects on health, and the depletion of finite resources leads to energy insecurity and rising geopolitical tensions [30]. In order to mitigate these negative effects, the use of fossil fuels needs to be reduced drastically and a transition towards sustainable technologies based on renewable energy sources is of utmost importance. In the development of more energy-efficient technologies, an in-depth understanding of the underlying physical processes is required. The research on fundamental phenomena lays the foundation for predictive models that are used for optimization and, more importantly, can lead to breakthrough innovations.

Two-phase flows play an important role in numerous key technologies involved in the transformation of different industry sectors, including energy, transportation, and chemicals. The involved two-phase flows are often governed by interfacial phenomena, such as the breakup or coalescence of droplets or bubbles [5], phase change [227], and complex three-dimensional (3D) deformations of the interface, especially in turbulent environments [137]. Due to a combination of these phenomena and the intricate non-linear interaction of the flow in both phases with the moving interfaces, the resulting two-phase flow dynamics are inherently 3D. To gain a deeper understanding of these dynamics, it is essential to extract three-dimensional data of the two-phase flows through experimental investigations. In particular, capturing the 3D interface shape is critical for the identification of process parameters such as the volume fraction in bubbly flows. The accurate measurement of the volume fraction is required for the prediction of different regimes that determine the interaction of the two phases and the overall dynamics of flow [38, 157]. Moreover, the interfacial area is a key determinant of heat and mass transfer rates, which is important for applications like chemical reactors, heat exchangers, and cooling systems [108, 102, 120]. In internal combustion engines, the complex interaction of fuel droplets with the walls of the engine determines the performance and pollutant emissions [155, 249]. The wetted area and wall contact times, which are key parameters in such processes, are directly governed by the dynamics of interfaces. In the context of energy transformation, two promising technologies in which two-phase flows involving liquid droplets play a crucial role are lithium-ion (Li-ion) batteries and proton-exchange membrane fuel cells (PEMFCs). Li-ion batteries have emerged as the preferred option for energy storage in battery electric vehicles due to their longevity, high energy density, and fast charging capability. However, their performance and lifespan are highly dependent on their operating temperature. Additionally, safety

concerns due to overheating and a potential runaway thermal reaction warrant a reliable thermal management system. Spray cooling has been proposed as a promising candidate for the cooling of Li-ion batteries in electric vehicles, as very high heat flux densities can be achieved with a compact design [196]. The efficiency of spray cooling is driven by multiple heat transfer mechanisms, involving the evaporation of droplets on the heated wall and the interaction of impinging droplets with liquid films that form on the surface [36]. A better understanding of impinging droplet dynamics is therefore indispensable for the further development of spray cooling techniques. For fuel cell electric vehicles, PEMFCs are considered to be the most viable option, due to their high power density and low operating temperatures [2]. However, to date, water management remains a major challenge for their application. In PEMFCs, the chemical energy of hydrogen and oxygen is directly converted into electricity, while water, as the main byproduct of the reaction, is typically released as vapor. The low operating temperatures in PEMFCs can lead to condensation and the emergence of water droplets in the gas channels of the fuel cell. The subsequent "flooding" of the fuel cell with liquid water blocks the supply of reactants and reduces active area, which leads to performance deterioration, unstable operating conditions, and faster degradation of the fuel cell [3]. Furthermore, the blockage of the gas channels by liquid droplets results in a pressure drop and flow maldistribution in the fuel cell manifold [151]. Consequently, the removal of liquid droplets from the gas channels is crucial for the efficient operation of PEMFC.

## 1.1 Motivation

The conditions for droplet removal in PEMFCs depend on the operating conditions, in particular the gas flow rate and the dynamic contact angles of the surfaces in the gas channels, among other surface properties [3]. However, the understanding of the fluid-mechanical processes and the parameters that influence droplet stability remains incomplete [3, 151]. The redirection of the flow around the adhering droplet eventually leads to periodic flow separation, which in turn induces an oscillation of the droplet. Due to the drag forces acting on the droplet, the gas-liquid interface is deformed and, furthermore, set in motion. The moving gas-liquid interface, in turn, induces a circulating flow inside the droplet. The complex interaction of the external and internal flow with the gas-liquid interface is assumed to have a significant influence on the stability of the droplet [26]. To date, little empirical data is available, so that further experiments are required to gain insight into the detachment of adhering droplets in external shear flows. Previously, particle image velocimetry (PIV) has been used to reveal the secondary flow inside of the droplet [24, 56, 151]. However, the refraction of light by the gas-liquid interface of the droplet leads to a distortion of the particle image positions and, consequently, the measured velocity field. For droplets with rotational symmetry and limited deformation, analytical methods for distortion correction have been developed [107, 150]. However, adhering and oscillating droplets in external shear flows exhibit significantly large and non-axisymmetrical deformation, which is, furthermore, temporally variable. Therefore, the instantaneous three-dimensional shape of the gas-liquid interface has to be obtained for an accurate measurement of the internal flow field of adhering droplets [26].

“For the reconstruction of the three-dimensional gas-liquid interfaces in two-phase flows various measurement techniques, based on different optical phenomena have been proposed, including refraction, fluorescence, light scattering, and structured light techniques.” from [DKS24, Section 1] The most common approaches employ the shadowgraphy technique due to its simplicity and high spatial accuracy. For gas-liquid interfaces with little deformation, single-camera techniques in combination with the assumption of symmetries and simple geometric shapes can provide an accurate reconstruction [226, 85]. For the reconstruction of highly deformed interfaces, however, methods based on multiple viewpoints and elaborate reconstruction algorithms are required [66,

143]. The increased complexity of these setups leads to higher costs and demands a greater effort in calibration, while the number of potential sources of error increases. Furthermore, optical access for multiple cameras and light sources is required, thus limiting the applicability of multi-camera techniques. Even with a large number of viewpoints, non-convex shapes with self-occlusion occasionally remain a difficulty and produce large errors. Therefore, the development of a single-view technique for the reconstruction of highly deformed and self-occluding gas-liquid interfaces is highly desirable.

Deep learning has proven to be a powerful and versatile approach for problems in which the system behavior is unknown or complicated to model, such as the reconstruction of three-dimensional shapes from two-dimensional images [149, 191, 251]. A high level of interest towards image-based volumetric reconstruction by means of neural networks has led to impressive results in recent years, thus rendering deep learning as a promising approach for interface reconstruction. The inclusion of prior knowledge, obtained by deep learning techniques from training data, has been shown to enhance the accuracy of the reconstruction, in particular in regions where little or no information can be extracted from the images directly. Recently, physics-informed machine learning [110] has emerged as a promising framework for the seamless integration of mathematical models, such as the governing equations of fluid flows, into data-driven algorithms, providing further prior knowledge for the regularization of ill-posed inverse problems.

## 1.2 Research objective and outline of the study

Due to the relevance of adhering droplet flow dynamics for the application in fuel cells, the research on such phenomena has attracted much attention over recent years. Consequently, a large body of knowledge and research data across different domains, including experimental and numerical investigations, as well as modeling efforts based on analytical and empirical approaches, is available for the development of data-driven techniques. This immediately raises the question of what potential lies within deep learning to bridge the gaps in the research through the integration of different domain-specific insights toward a more comprehensive model of the underlying physics in droplet flows. The focus of this work is the development of a predictive framework for gas-liquid interfaces in two-phase flows, which combines insights from experiments, numerical simulations, and analytical approaches to enable more in-depth and more efficient exploration of two-phase flows. This general tool is demonstrated by the example of adhering droplets but is ultimately intended to serve as a foundation for further research into the complex field of interface reconstruction in two-phase flows. In order to address this open issue, it is further broken down into specific research questions in the following.

### Research questions:

- RQ1 Can deep learning techniques be used to reconstruct the spatio-temporal gas-liquid interface dynamics in two-phase flows from optical experiments?
- RQ2 How should a deep learning framework for the instantaneous reconstruction from single-view optical experiments be designed?
- RQ3 Does the information in the images sufficiently represent the three-dimensional droplet dynamics to allow for an accurate volumetric reconstruction? How can the experiments be modified in order to enhance the amount of usable information in the images?
- RQ4 How and to what extent can available knowledge, *e.g.*, from numerical simulation and analytical approaches be used to improve the accuracy of the reconstruction?

RQ5 How accurate and how reliable is the data-driven reconstruction? What are the factors that influence the quality of the reconstruction and determine the limits for the application?

The first question (RQ1) states the overarching problem definition of the present work, which is aspired to be answered over the course of this thesis through a review of current reconstruction methods, theoretical considerations, and extensive experimental investigations. The second question (RQ2) is addressed in Chapter 3. To lay the groundwork for an informed decision on the design of the reconstruction framework, first, a literature review on the current state-of-the-art in two-phase flow reconstruction and volumetric reconstruction by means of deep learning is conducted. This is followed by the conception of a reconstruction framework, in which the suitability of different deep learning techniques for the spatio-temporal reconstruction of the gas-liquid interface, as well as the availability of resources, in particular existing experimental facilities and databases from previous numerical and experimental investigations, are considered. The most suitable reconstruction method is adapted and integrated with a data processing pipeline that allows for the generation of appropriate training data and provides ground truth data for the quantitative evaluation of the reconstruction quality. A part of the findings regarding RQ2 have been published in [DKS24]. The third question (RQ3) is addressed by the comparative evaluation of a purposefully developed novel measurement technique with the canonical shadowgraphy method. Based on the theoretical consideration of light scattering on spherical droplets, an optical measurement method that employs color-coded glare points to embed additional three-dimensional information of the gas-liquid interface in the images is developed in Chapter 4. The effectiveness of glare points for encoding additional three-dimensional information is demonstrated first by the reconstruction of a simplified droplet geometry by analytical means. The foregoing results for RQ3 have been published in [DBB<sup>+</sup>23]. In order to shed light on the impact of glare points for the reconstruction of complex deformed gas-liquid interfaces, a comparative study with the same reconstruction method that receives only shadowgraph images as input is conducted in Chapter 6. The results of this study were previously published in [DKS24]. The fourth question (RQ4) is addressed in two differentiated evaluation steps. The first part of the question is addressed by the incorporation of results from numerical simulation into the data generation and training process during the conception stage of the reconstruction framework in Chapter 5. The effect of numerical training data is subsequently analyzed on the basis of the reconstruction results obtained by experimental investigation in Chapter 6. The results for this first part of RQ4 have been published in [DKS24] and [DHK<sup>+</sup>25]. In the second step, physics-informed learning techniques are explored and their effect on the reconstruction performance is measured in a comparative experimental investigation to the previously developed data-driven baseline model in Chapter 7. The results of this study were previously published in [DKK<sup>+</sup>24]. In order to address the last question (RQ5), an experimental investigation of the reconstruction performance for droplet flows with gradually increasing complexity of droplet deformation is conducted in Chapter 6. The accuracy of the reconstruction methods is quantified under various conditions and for different reconstruction tasks, starting from axisymmetrical droplet impingement, over substrates with different structures and anisotropic wettability, and finally adhering and oscillating droplets in external shear flows. Thereby, possible sources of uncertainty and their influence on the reliability of the novel reconstruction technique, as well as the quality of the obtained reconstruction results are determined. Part of the results regarding RQ5 have been published in [DKS24] and [DHK<sup>+</sup>25]. In the present thesis, two different styles of citation will be used to differentiate references to published works that include the contribution of the author of this thesis from the work of others. Self-citations will be referenced in an alphabetical style, as seen in the previous paragraph, while the work of others will be referenced in a numeric style.



## 2 Fundamentals

In this chapter, first, the fundamentals of two-phase flow phenomena are reviewed, including a brief derivation of the governing equations for two-phase flows. Subsequently, the basic principles of the experimental techniques employed for the observation of two-phase flows in the present work are introduced. This is followed by an overview of the deep learning techniques that provide the foundation for the development of data-driven techniques for the reconstruction of two-phase flows. Parts of this chapter have previously been published in the articles *Three-dimensional encoding of a gas-liquid interface by means of color-coded glare points* by Dreisbach *et al.* [DBB<sup>+</sup>23], *Spatio-temporal reconstruction of droplet impingement dynamics by means of color-coded glare points and deep learning* by Dreisbach *et al.* [DKS24], and *Interface reconstruction of adhering droplets for distortion correction using glare points and deep learning* by Dreisbach *et al.* [DHK<sup>+</sup>25]. Where findings have already been published in one of the referenced works, this is clearly marked and identified as related quotations.

### 2.1 Governing equations of fluid mechanics

In the following, the governing equations that describe the motion of fluids, which are used throughout this work, are introduced. The behavior of a fluid flow can be described by the interpretation of the fundamental equations of continuum mechanics in the context of fluids. The consideration of fluids as continua allows for the description of the flow field at any spacial point  $\mathbf{x} = x_i$  and time  $t$  by the differential form of the conservation laws for mass and momentum.

**Conservation of Mass** The differential form of the equation for mass conservation, which in the following will be referred to as the continuity equation, reads [210, p. 37]

$$\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x_i} (\rho u_i) = 0, \quad (2.1)$$

where  $\rho$  is the density of the fluid and  $u_i$  the velocity components of the flow with  $i = 1, 2, 3$ . Through the application of the material derivative  $\frac{D(\dots)}{Dt} = \frac{\partial(\dots)}{\partial t} + u_i \frac{\partial(\dots)}{\partial x_i}$ , the continuity equation can be expressed in the following form

$$\frac{D\rho}{Dt} + \rho \frac{\partial u_i}{\partial x_i} = 0. \quad (2.2)$$

If the fluid can be considered incompressible, the density of a material particle remains constant along its track, even under varying pressure, which results in  $\frac{D\rho}{Dt} = 0$ . Equation 2.2 is then reduced to the conservation of volume and becomes the continuity equation for incompressible flow

$$\frac{\partial u_i}{\partial x_i} = \nabla \cdot \mathbf{u} = 0. \quad (2.3)$$

It should be noted, that incompressibility does not imply uniform density, as two different material particles can retain different constant densities in an incompressible flow [210, p. 38].

**Conservation of Momentum** The differential form of the linear momentum conservation equation is given by [210, p. 37]

$$\rho \frac{Du_j}{Dt} = \rho k_j + \frac{\partial \sigma_{ij}}{\partial x_i}, \quad (2.4)$$

where  $k_j$  represents a specific force per volume, such as the gravitational acceleration and  $\sigma_{ij}$  the stress tensor. The differential form of the momentum conservation equation is valid for any continuum and therefore has to be combined with a material law for the considered fluid that relates the stress tensor to the motion of the continuum to obtain the governing equations that allow for the description of the specific fluid flow. The stress tensor  $\sigma_{ij} = \sigma_{ji}$  is symmetrical [210, p. 49] and can be decomposed into the fluid pressure  $p$  and viscous stress tensor  $\tau_{ij}$  for the case of fluids [199, pp. 56–57]

$$\sigma_{ij} = -p\delta_{ij} + \tau_{ij}, \quad (2.5)$$

where  $\delta_{ij}$  is the Kronecker delta. For a Newtonian fluid the viscous stress tensor  $\tau_{ij}$  is linearly dependent on the strain rate tensor  $\varepsilon_{ij}$ , which is defined as [199, pp. 62–64]

$$\varepsilon_{ij} = \frac{1}{2} \left( \frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right). \quad (2.6)$$

The material law for an isotropic Newtonian fluid can be expressed as [199, p. 65]

$$\tau_{ij} = -\frac{2}{3}\mu\varepsilon_{kk} + 2\mu\varepsilon_{ij}, \quad (2.7)$$

where  $\mu$  is the dynamic viscosity and  $\varepsilon_{kk} = \frac{\partial u_k}{\partial x_k}$ . The substitution of the material law for a Newtonian fluid into the momentum conservation equation (Equation 2.4) yields the Navier-Stokes equations (NSE)

$$\rho \left( \frac{\partial u_j}{\partial t} + u_i \frac{\partial u_j}{\partial x_i} \right) = \rho k_j - \frac{\partial}{\partial x_j} \left( p + \frac{2}{3}\mu \frac{\partial u_k}{\partial x_k} \right) + \frac{\partial}{\partial x_i} \left( \mu \left( \frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) \right). \quad (2.8)$$

For incompressible fluids Equation 2.8 is simplified to

$$\rho \left( \frac{\partial u_j}{\partial t} + u_i \frac{\partial u_j}{\partial x_i} \right) = \rho k_j - \frac{\partial p}{\partial x_j} + \frac{\partial}{\partial x_i} \left( \mu \left( \frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) \right), \quad (2.9)$$

under consideration of Equation 2.3, or in tensor notation

$$\rho \left( \frac{\mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} \right) = \rho \mathbf{k} - \nabla p + \nabla \cdot (\mu (\nabla \mathbf{u} + (\nabla \mathbf{u})^\top)). \quad (2.10)$$

The dynamic viscosity does not necessarily need to be uniform across the flow field in incompressible flows. In general, the viscosity is a function of the temperature [9, p. 59] [see also 210, p. 6], so

that for isothermal flows of one thermodynamical phase a uniform distribution of the viscosity can be assumed [210, p. 104]. However, for two-phase flows, the viscosity generally has a non-uniform distribution even in isothermal conditions, due to the discontinuity at the interface.

## 2.2 Two-phase flow phenomena

Multiphase flows involve two or more different thermodynamic phases, which can be gaseous, liquid, or solid. Gas-liquid or immiscible liquid-liquid two-phase flows are characterized by a sharp interface that separates the two phases. This interface is deformable and can undergo topological changes due to break-up or coalescence. The physical properties of the bulk phases, such as density and viscosity change across a very small transition region at the interface, which is typically only a few molecules wide. In the framework of continuum mechanics, the interface thickness vanishes and instead, a sharp interface, *i.e.* a discontinuity at the phase boundary is assumed. Two phase-flows may involve phase change phenomena, such as boiling at high temperatures or cavitation at low pressure. However, the present work focuses on the two-phase flow of water and air at room temperature and atmospheric pressure, and therefore only evaporation and condensation require consideration. Additionally, the flow velocities in this work are low enough that compressibility effects can be neglected. Two-phase flows at the scale of a few millimeters, as considered in this work, are affected by capillary forces that arise due to surface tension, which will be further elucidated in the following section. Subsequently, static and dynamic wetting of flat and structured solid surfaces will be discussed, followed by a brief review of the research on droplet impingement phenomena and adhering droplets subjected to external shear flows. This section closes with an overview of mathematical models employed for the description of two-phase flows in numerical simulation.

**Surface tension** Surface tension originates from molecular interactions and thermal effects on the micro-scale and is represented in continuum mechanics by its macroscopic effects, as a stress on the interface. The molecular interactions are determined by long-range attractive and short-range repulsive forces, which are represented by the Lennard-Jones intermolecular potential [129]. Consequently, at large distances attractive forces dominate. In the bulk of the phases the attractive and repulsive intermolecular forces generally even out. However, at the interface, the attractive forces parallel to the surface are stronger than the repulsive forces, due to a larger intermolecular distance that results from a lower density in the transition region at the interface. This force anisotropy localized close to the interface gives rise to surface tension. For a more extensive explanation of the formation mechanisms that lead to the occurrence of surface tension, the reader is referred to the publication of Marchand et al. [142]. The resulting surface tension can be considered from a thermodynamic or mechanical point of view. From the thermodynamic standpoint, surface tension is defined as the excess free energy at the interface. The molecules at the interface are in an energetically unfavorable state and, as a consequence, the creation of surface area requires energy.

$$dE = \gamma dA, \quad (2.11)$$

where  $dA$  is an infinitesimal surface area and  $\gamma$  the coefficient of surface tension. In order to reach the minimal possible free surface energy, a fluid system tends to minimize its surface area. Surface tension from the mechanical point of view is defined as a force per length unit  $l$  that acts parallel to the surface and normal to the contour of the surface

$$F_\gamma = \gamma l. \quad (2.12)$$

The equivalence of both points of view can be demonstrated by the principle of virtual work. For the creation of the surface  $dA = l \cdot dx$ , where  $dx$  is an infinitesimal length, the work  $F_\gamma \cdot dx$  has to be expended, which increases the free interfacial energy by  $dE = \gamma l \cdot dx = \gamma dA$ . The surface tension coefficient is a function of temperature and concentration, where a small concentration of surfactants can significantly alter surface tension already [188]. If no surface tension arises at the interface of two fluids, they are considered miscible; otherwise, they are classified as immiscible or partially miscible.

**Laplace pressure** A consequence of surface tension is the pressure difference over curved interfaces. The force balance over a curved interface in the equilibrium state reveals that the resultant force of the surface tension requires an opposing pressure force of equal magnitude. Consequently, the pressure on the side of the curved interface at the center of the radius is greater than the outside pressure. The pressure difference is given by the Young-Laplace equation [44, pp. 6–8]

$$\Delta p = \gamma \left( \frac{1}{r_1} + \frac{1}{r_2} \right), \quad (2.13)$$

where  $r_1$  and  $r_2$  are the principal radii of curvature in two orthogonal directions. This pressure caused by surface tension is called Laplace pressure.

**Dimensionless numbers** The two-phase flows considered in the following work can be regarded as a physical system that is dependent on six physical variables, in particular the density  $\rho$ , the dynamic viscosity  $\mu$ , the surface tension coefficient  $\gamma$  and the gravitational acceleration  $g$ , as well as the characteristic length scale  $L$  and velocity  $U$ . Furthermore, the system is defined by the three fundamental units mass, length, and time. According to Buckingham's theorem [20] this physical system can be represented by three dimensionless groups:

- Bond number  $Bo = \frac{\rho g L^2}{\gamma} = \frac{\text{gravitational forces}}{\text{interfacial forces}}$
- Reynolds number  $Re = \frac{\rho U L}{\mu} = \frac{\text{inertial forces}}{\text{viscous forces}}$
- Weber number  $We = \frac{\rho U^2 L}{\gamma} = \frac{\text{inertial forces}}{\text{interfacial forces}}$


The Weber and Reynolds numbers are crucial for the characterization of various interfacial phenomena, including the splashing of liquid droplets impacting on solid surfaces, the break-up of droplets during atomization, or the transition between different regimes in bubble flows. For  $Bo \approx 1$  the gravitational and interfacial forces are comparable in magnitude, which yields the definition of the capillary length  $l_c = \sqrt{\frac{\gamma}{\rho g}}$ . For characteristic length scales smaller than  $l_c$  two-phase flow problems are typically dominated by capillary effects. For air-water flows the capillary length is  $l_c \approx 2$  mm, computed with typical values of  $\gamma = 70$  mN/m,  $\rho_w = 997$  kg/m<sup>3</sup> and  $g = 9.81$  m/s<sup>2</sup>. Without the influence of gravity, a free-floating droplet takes the shape of a sphere, which still holds approximately true for a falling droplet at  $Bo < 1$ .

Further relevant dimensionless numbers in the context of two-phase flows are the Froude number  $Fr = \frac{U}{\sqrt{gL}}$ , the Capillary number  $Ca = \frac{\mu U}{\gamma}$  and the dimensionless time  $t^* = \frac{tU}{L}$ , which relates the physical time  $t$  to the characteristic length and velocity. The Froude number represents the ratio of inertial and gravitational forces and is used for the characterization of two-phase flows

where gravity is important, including vertical, stratified, and free-surface flows. The Capillary number represents the ratio of the viscous and interfacial forces and is crucial for representing the dimensionless contact line velocity in dynamic wetting phenomena.

### 2.2.1 Wetting of smooth and structured surfaces

**Static wetting** A relevant two-phase flow phenomenon is the wetting of a solid surface by a liquid phase in the presence of a gaseous phase. The adhesive forces acting between the liquid phase and the solid cause the liquid to spread on the surface, while the cohesive forces in the liquid phase simultaneously induce the liquid to take a spherical form. The particular force balance results in the formation of a liquid droplet or film on the surface, depending on the wettability of the surface.



**Figure 2.1:** Static wetting of an ideal solid surface by a liquid droplet. The balance of the interfacial energies between the liquid (l), gaseous (g) and solid (s) phases,  $\gamma$ ,  $\gamma_{ls}$  and  $\gamma_{gs}$ , determines the equilibrium contact angle  $\Theta_{eq}$ .

The Young equation [258] describes the equilibrium state for a droplet wetting an ideal surface, which is characterized as a flat, smooth, rigid, and chemically homogeneous surface in horizontal orientation

$$\gamma_{gs} = \gamma_{ls} + \gamma \cos(\Theta_{eq}), \quad (2.14)$$

where  $\gamma$ ,  $\gamma_{ls}$  and  $\gamma_{gs}$  are the interfacial energies at the gas-liquid, the liquid-solid, and the gas-solid interface, respectively and  $\Theta_{eq}$  is the equilibrium contact angle. The contact angle is defined as the angle between the tangent of the gas-liquid interface and the tangent at the solid-liquid interface at the point where the three phases meet, *i.e.* the contact line, as illustrated in Figure 2.1. By convention, the contact angle is defined as inward facing to the liquid phase. The balance of the three interfacial energies determines a unique solution for the equilibrium contact angle  $\Theta_{eq}$ , which is a measure of the degree of wettability: In particular, total wetting occurs for  $\Theta_{eq} = 0^\circ$  and leads to the formation of a liquid film, whereas partial wetting occurs for  $\Theta_{eq} > 0^\circ$  and leads to the formation of a droplet. Furthermore,  $\Theta_{eq} < 90^\circ$  characterizes a wetting surface, which in the case of water is referred to as hydrophilic, while  $\Theta_{eq} > 90^\circ$  characterizes a weakly wetting or hydrophobic surface and  $\Theta_{eq} > 150^\circ$  superhydrophobic surfaces.

**Dynamic wetting** Dynamic wetting occurs during the motion of a liquid droplet surrounded by gas on a solid surface. It has been found that in reality, the contact angle has a hysteresis, *i.e.* that no movement of the interface occurs for  $\Theta_{adv} < \Theta < \Theta_{rec}$ , where  $\Theta_{adv}$  and  $\Theta_{rec}$  are defined as the advancing and receding contact angles, respectively [12, pp. 255ff.]. The contact angle hysteresis is defined as  $\Delta\Theta = \Theta_{adv} - \Theta_{rec}$  and is a measure for the "stickiness" of a surface, *i.e.* how much external force would be required to move the droplet on the surface. The cause for the contact

angle hysteresis lies in the pinning of the contact line on surface heterogeneities, such as roughness or chemical heterogeneities, which leads to further equilibrium states between  $\Theta_{\text{adv}}$  and  $\Theta_{\text{rec}}$ . The contact line then has to overcome these local energy barriers on the heterogeneous surface in order to start moving. Furthermore, it has been found that the contact angles of moving contact lines, in the following referred to as dynamic contact angles, depend on the velocity of the contact line  $u_{\text{CL}}$ . The advancing contact angle increases with contact line velocity, while the receding contact angle decreases. This dependency is represented by the non-dimensional capillary number  $Ca = \frac{\mu u_{\text{CL}}}{\gamma}$ , for which the relationship  $\Theta = f(\Theta_{\text{eq}}, Ca)$  has been determined experimentally [89]. The monographs of Berg [12] and De Gennes et al. [44] are recommended for a comprehensive review of capillarity effects and wetting phenomena. For the mathematical modeling in the framework of continuum mechanics, it has to be considered that the no-slip boundary condition for the velocity is not appropriate for the moving contact line in the context of a sharp interface. A velocity boundary condition with a tangential slip of the interfacial velocity at the solid surface is required instead, such as the Navier slip condition. Furthermore, more complex boundary conditions for the dynamic contact angle have to be considered.

**Wetting of structured surfaces** Surface structures can exist on a wide range of scales, and relevant lengths in the context of two-phase flows typically range from a few nanometers to hundreds of micrometers. Furthermore, the structures can be regular or irregular and thereby include roughness. The wetting of structured surfaces is characterized by two distinguished regimes, namely the Wenzel state [247] and the Cassie-Baxter state [31]. The wetting state has a significant influence on the wettability and the "stickiness" of the wetted substrates, which are expressed by varying static contact angles and contact angle hysteresis, respectively. In the Wenzel state, the surface is completely wetted, as liquid enters into the surface structure. The heterogeneity of the surface enhances the hydrophobicity of already hydrophobic materials but also makes hydrophilic surfaces more hydrophilic, which is reflected in the Wenzel equation

$$\cos(\Theta_{\text{W}}) = R_{\text{W}} \cos(\Theta_{\text{eq}}), \quad (2.15)$$

where  $R_{\text{W}}$  is the ratio of actual surface area to normally projected surface area. Furthermore, the Wenzel state is associated with an increase in the contact angle hysteresis due to an increased pinning of the liquid interface at the surface heterogeneities. In the Cassie-Baxter state, the surface is partially wetted by the liquid, because gas is trapped in the surface structure so that the liquid is only in contact with the peaks of the structure. This composite wetting state leads to a reduced wettability, which is expressed by the Cassie-Baxter equation


$$\cos(\Theta_{\text{CB}}) = R_1 \cos(\Theta_{\text{eq}}) - R_2, \quad (2.16)$$

where  $R_1$  is the area proportion of the wetted area and  $R_2 = 1 - R_1$  is the area proportion of the gas-liquid interface in the composite surface. Structured surfaces of hydrophobic materials in the Cassie-Baxter state can become superhydrophobic. Typically the Cassie-Baxter state is associated with a low contact angle hysteresis due to a reduced solid-liquid interaction. However, technically engineered surfaces can combine hydrophobicity with "stickiness", as the combination of roughness and surface structure can result in a combined Cassie-Baxter/Wenzel state [183]. Both regimes represent separate energy minima, however, the Cassie-Baxter state is meta-stable, which can lead to a transition from the Cassie-Baxter to the Wenzel state. For this transition, an energy barrier has to be overcome and external influences, such as gravity, vibration, thermal perturbation, compression, or mechanical forces can induce the transition [158]. Anisotropy in the

surface structure, such as parallel grooves, stripes of roughness or chemical functionality, as well as gradients in wettability can induce anisotropic wetting, which introduces a directionality in the distance that the liquid spreads on the surface, as well as static and dynamic contact angles [253]. Anisotropic structured surfaces are ubiquitous in nature and often possess intricate hierarchical structures with excellent water-repellent properties. Furthermore, gradients in wettability induce a force acting on the liquid, which is exploited in technically engineered surfaces to guide the fluid flow *e.g.* for droplet removal.

### 2.2.2 Droplet impingement

“The impingement of liquid droplets on wet or dry surfaces is an important phenomenon in a multitude of technical applications, such as spray cooling [154], spray coating [4, 43], inkjet printing [135] and combustion [155].” from [DKS24, Section 1] The outcome of droplet impingement is dependent on the kinematic conditions, in particular the impact velocity, the angle of incidence, and the volume of the droplet, as well as the fluid properties – most importantly the surface tension, the viscosity, and the density of the liquid droplet. If the impact on a solid substrate is considered, the wettability, structure, and roughness of the substrate influence whether a deposition, splashing, a partial or complete rebound of the droplet occurs. The most relevant dimensionless groups that determine the outcome of droplet impingement are the Weber and Reynolds numbers, defined as  $We = \frac{\rho u_0^2 d_0}{\gamma}$  and  $Re = \frac{\rho u_0 d_0}{\mu}$ , with the initial spherical droplet diameter  $d_0$ , the impact velocity  $u_0$ , as well as the density  $\rho$  and dynamic viscosity  $\mu$  of the liquid phase. The different outcomes of droplet impingement on solid substrates, as classified by Rioboo et al. [180], are illustrated in Figure 2.2.




**Figure 2.2:** Regimes of droplet impingement; time series of an impinging droplet until maximum spreading (a) to (c), resulting in deposition (e<sub>1</sub>) or the formation of a jet. Depending on the kinetic energy of the impact, a thin jet (d<sub>1</sub>) emerges, resulting in secondary droplet ejection (e<sub>2</sub>), or a thick jet, leading to partial rebound (e<sub>3</sub>) or rebound (e<sub>4</sub>). Impingement with high kinetic energy resulting in the splashing of the droplet (f) to (h).

A time series of recorded images from droplet impingement experiments that resulted in the deposition of the droplet is illustrated in Appendix Figure A.3. In the deposition regime, the droplet remains on the surface of the substrate from the moment of initial impact until it reaches its



equilibrium state. Due to its high kinetic energy upon impact, the droplet deforms and dynamically wets the substrate. On horizontal flat surfaces the droplet spreads radially and forms a thin lamella that is bounded by an elevated rim of almost circular cross-section [257]. In this case, the droplet deformation is axisymmetrical. During this first, so-called spreading phase, the droplet's kinetic energy is partly dissipated by viscous effects and partly transformed into surface energy as the surface area increases significantly. The droplet reaches its maximum extent of spreading, indicated by the maximum spreading diameter when the inertial and surface tension forces balance out. The maximum spreading diameter is governed by the kinetic energy of the impacting droplet and, therefore, increases with higher impact velocities. After reaching the maximum spreading diameter, the behavior of the droplet is primarily influenced by the wettability of the substrate. For wettable substrates, the droplet will continue to spread on the surface until equilibrium conditions are met. On hydrophobic surfaces, the droplet retracts while minimizing the surface energy, which is partly transformed back into kinetic energy and partly dissipated. This phase of droplet impingement is referred to as the receding phase [181]. The droplet might spread again after the receding phase, depending on the remaining kinetic energy. Consequently, a repeating cycle of spreading and receding follows, which induces a dampened oscillation of the droplet until an equilibrium state is reached. At higher impact velocities the droplet might splash upon impact, ejecting numerous secondary droplets radially. On hydrophobic substrates, receding break-up can occur, where secondary droplets are split off during the receding phase. If the remaining kinetic energy at the end of the receding phase is sufficiently high, the radially converging flow produces an upward jet. Depending on the wettability of the substrate and the remaining kinetic energy, the height and diameter of this jet vary and small secondary droplets might be ejected due to the instability of thin jets [180]. Larger jets result in a partial rebound of the droplet, where a significant portion of the droplet volume is ejected or in a complete rebound, characterized by the complete detachment of the droplet from the substrate [256]. In particular, superhydrophobic surfaces are associated with a rebound of the droplet, rendering such substrates interesting for technical applications such as self-cleaning or anti-icing surfaces [105, 112]. As introduced in Subsection 2.2.1 structured hydrophobic substrates can attain superhydrophobicity, which has instigated research into manufacturing processes and properties of structured hydrophobic surfaces [174, 64]. “Droplet impingement is an inherently three-dimensional (3D) process due to the presence of outside perturbation from the environment and heterogeneity of the substrate originating from *e.g.*, variation in chemical properties or surface roughness [69]. In particular, angled impacts and the anisotropic wetting of structured surfaces lead to a complex three-dimensional deformation of the gas-liquid interface [153].” from [DKS24, Section 1]

### 2.2.3 Adhering droplets in shear flows



**Figure 2.3:** Visualization of an adhering droplet subjected to an external shear flow; time-averaged flow topology inside the droplet and in the external flow (left), droplet and flow oscillation (right). Figure adapted from Burgmann et al. [25].

“Adhering droplets in external flows occur in many technical applications, for instance in cleaning and drying processes [224, 200], oil recovery [223, 198, 76, 140], heat exchangers [120, 106, 246], airfoil icing prevention [109] and in fuel cells, where a removal of the droplets is crucial for an efficient operation [222, 123, 60, 22]. The flow around the droplet eventually leads to a detachment



of the droplet from the solid surface at critical conditions, which occur when the external forces acting on the droplet become greater than the adhesion of the liquid droplet to the surface [63, 65, 8, 25]. Several investigations have revealed an intricate interaction of the droplet with the surrounding flow. The droplet deforms [76, 200, 7] and additionally oscillates, when subjected to external shear flows [265, 23]. It was shown that the contact angle hysteresis of the substrate and the droplet volume have an influence on the critical velocity for droplet detachment [53, 123, 8].

Previous investigations based on particle image velocimetry (PIV) measurements revealed that the internal flow within the droplet is relevant to understanding the mechanism of droplet detachment. It was found that the internal flow follows the main flow direction of the external flow [55] and exhibits a clockwise rotational flow pattern [151]. Minor et al. [151] deduced that the secondary flow inside a droplet is induced by shear forces on the interface resulting from the external flow. Burgmann et al. [24] confirmed the findings of a clockwise vortex at low external flow velocities but also found an additional counterclockwise rotation at higher flow velocity that becomes dominant when approaching the critical velocity. Furthermore, the authors found an oscillation of the droplet with increasing amplitude towards higher external flow velocities. Burgmann et al. [25] proposed that the droplet detachment is associated with a self-excitation process resulting from the oscillation of the gas-liquid interface, as well as the oscillation of the inner flow structure and the flow field around the droplet. The authors demonstrated that an emerging backflow region at higher external flow velocities induces the change in the flow pattern within the droplet and that the wake flow oscillates with the same frequency as the droplet. Later, Burgmann et al. [26] found that the coupling of the internal and external flow leads to increased flow separation in comparison to rigid bodies and, consequently, a higher pressure force driving droplet detachment. Bilsing et al. [13] measured the internal 3D flow field by means of microscopic particle tracking velocimetry (PTV) through a transparent substrate from below and thereby revealed that the internal flow topology consists of multiple 3D vortical structures that change drastically depending on the external flow velocity. In experiments with rigid droplet models, Zhang et al. [264] found that the external flow topology is significantly influenced by the shape of the adhering droplet. Consequently, a better understanding of the droplet deformation is required for any future investigation of the external flow. A visualization of the flow phenomena inside and around the droplet is shown in Figure 2.3. The mutual dependence of wake flow, contour oscillation, and internal flow structure influences the stability of the droplet and eventually leads to the detachment of the droplet at critical conditions. While there has been significant scientific progress in recent years, the exact mechanism of droplet detachment still remains to be fully understood. Therefore, further experiments are required in order to obtain a better analytical model of droplet stability that accounts for aeroelastic effects and, consequently, allows for a more accurate prediction of the onset of droplet motion.” from [DHK<sup>+</sup>25, Section 1]

## 2.2.4 Numerical modeling of two-phase flows

A multitude of numerical methods for the simulation of two-phase flows have been developed, including particle-base models [212, 54] and continuum methods, encompassing two-fluid models [190, 87] and one-fluid models. In two-fluid models, each phase is simulated in a separate subdomain, coupled by boundary conditions at the interface. Conversely, one-fluid models treat both phases as a single fluid with locally varying properties. One-fluid models can be further categorized into interface-tracking approaches [79, 230] and interface-capturing approaches. The latter includes level set [214, 68], volume-of-fluid (VoF) [164, 86, 259], and the phase-field [100, 261] methods. A comprehensive description of the aforementioned techniques can be found in the reviews by Wörner [252] and Mirjalili et al. [152]. The current work focuses on one-fluid models, particularly VoF and

the phase-field approach. Furthermore, this subsection is deliberately kept brief to focus on the fundamental concepts that are required for the development of physics-informed neural networks and, underlie the numerical datasets that are used in this work to train the neural networks. Continuum methods for the numerical simulation of two-phase flows require a method that describes the spatio-temporal interface evolution and can be coupled with the momentum equation. The techniques for the representation of interface evolution are generally classified as either sharp or diffuse interface methods. In sharp interface methods, the interface thickness is assumed to be zero and the physical quantities of the fluids are discontinuous at the interface, whereas in diffuse interface methods, the interface has a finite thickness and physical quantities change continuously across the interface.

**Volume-of-fluid method** The Volume-of-fluid method [164, 86, 259] is a sharp interface technique based on a scalar indicator function  $\chi$ , which is used to follow the interface evolution and is defined as

$$\chi(\mathbf{x}, t) = \begin{cases} 1 & \text{for } \mathbf{x} \in \Omega_l \\ 0 & \text{for } \mathbf{x} \notin \Omega_l \end{cases} \quad (2.17)$$

on fixed Eulerian grid, with  $\Omega_l$  representing the subdomain of the liquid phase (more generally one of the phases) in the domain of the two-phase flow  $\Omega$ . If no phase change occurs, *i.e.* the value of  $\chi$  for a fluid particle does not change along its track, the following transport equation for  $\chi$  holds

$$\frac{D\chi}{Dt} = \frac{\partial\chi}{\partial t} + \mathbf{u} \cdot \nabla\chi = 0. \quad (2.18)$$

The volume fraction  $\alpha$  is defined as the volumetric average of  $\chi$  over the control volume  $V$

$$\alpha = \frac{1}{|V|} \int_V \chi dV \quad (2.19)$$

and indicates whether a computational cell is occupied by the liquid ( $\alpha = 1$ ), the gaseous phase ( $\alpha = 0$ ), or both ( $0 < \alpha < 1$ ). The integration of Equation 2.18 over  $V$  and the application of the Gauss divergence theorem yields

$$\frac{\partial\alpha}{\partial t} + \frac{1}{|V|} \int_{\partial V} (\mathbf{n}_e \cdot \mathbf{u}) \chi dA = 0, \quad (2.20)$$

with the unit normal vector  $\mathbf{n}_e$  on the control volume boundary  $\partial V$ . Two common approaches exist for the evaluation of the surface integrals in Equation 2.20, namely geometric VoF and algebraic VoF. In geometric VoF approaches the local orientation and position of the interface is reconstructed by geometric means. A commonly employed method is the piece-wise linear interface calculation (PLIC) [259], in which the interface is represented by a plane in each computational cell on the interface. The reconstructed interface is then further used to approximate the distribution of  $\chi$  in the cell and, subsequently, to calculate the numerical fluxes in Equation 2.20. An advantage of geometric VoF is inherent mass conservation. However, accurate interface reconstruction schemes are typically complicated and computationally expensive.

In algebraic VoF approaches [86] Equation 2.20 is approximated by the following transport equation for the volume fraction  $\alpha$

$$\frac{\partial \alpha}{\partial t} + (\mathbf{u} \cdot \nabla) \alpha = 0, \quad (2.21)$$

which can be solved by numerical schemes without the requirement of a geometric interface reconstruction. Consequently, algebraic VoF in its basic form offers a simplified and thus computationally cheap model for the simulation of two-phase flows. While algebraic VoF is derived under the assumption of zero interface thickness, in practice the numerical implementation of Equation 2.21 results in a finite interface thickness. The interface location is appropriately indicated by  $\alpha = 0.5$ . As the interface is smeared across multiple computational cells, algebraic VoF approaches are usually sensible to mesh resolution.

In both VoF approaches the transport equation for the volume fraction is solved alongside the continuity and momentum equations in order to obtain the velocity field information. An advantage of the VoF method is that topological changes, such as break-up or coalescence can be represented on a fixed grid, without a dedicated handling of the dynamics interface.

**Single-field Navier-Stokes equations** The discontinuity at the interface has to be considered in the momentum equation. For a two-phase flow in the domain  $\Omega$ , which is occupied by two immiscible incompressible Newtonian fluids in the respective subdomains  $\Omega_l$  and  $\Omega_g$  that are separated by the interface, the local NSE and continuity equation (Equations 2.10 and 2.3) are valid on  $\Omega_l$  and  $\Omega_g$ . In order to couple the conservation equations in the respective subdomains, jump conditions for the velocity and stresses at the interface are required. In each subdomain, the density and the viscosity of the fluids are assumed to be piece-wise constant, with  $\rho_l$  and  $\mu_l$  on  $\Omega_l$  and  $\rho_g$  and  $\mu_g$  on  $\Omega_g$ . The local NSE in each phase can be combined into the single-field NSE

$$\rho_M \left( \frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} \right) = -\nabla p + \nabla \cdot (\mu_M (\nabla \mathbf{u} + (\nabla \mathbf{u})^T)) + \rho_M \mathbf{g} + \mathbf{f}_\gamma, \quad (2.22)$$

which are valid across the whole domain  $\Omega$ , even if  $\rho$  and  $\mu$  change discontinuously [231, p. 42]. In the single-field approach, the immiscible fluids of both phases are treated as one effective fluid with varying density and viscosity  $\rho_M = \rho_M(\mathbf{x}, t)$  and  $\mu_M = \mu_M(\mathbf{x}, t)$ , which are piece-wise constant in each phase and have discontinuities at the interface. The physical properties are calculated using weighted averages of the properties of the two respective fluids as a function of the volume fraction  $\alpha$ . Commonly, the arithmetic mean is used as a mixture rule for both  $\rho_M$  and  $\mu_M$

$$\xi_M = \alpha \xi_l + (1 - \alpha) \xi_g \quad \text{with} \quad \xi \in \{\rho, \mu\}. \quad (2.23)$$

However, in some approaches the viscosity is weighted using the harmonic [252]. Consequently, the instantaneous phase field distribution is required to solve single-field NSE and, thereby, the interface evolution equation is coupled with the momentum equation. Furthermore, Equation 2.23 reveals that an accurate computation of the volume fraction is crucial, as small errors in the volume fraction translate into errors of the mixture density  $\rho_M$ , which in turn introduces an error to the momentum balance, in particular for high density ratios.

**Modeling of interfacial forces** The accurate representation of interfacial forces is a critical issue for the simulation of two-phase flows. The continuum surface force (CSF) model [16] is commonly employed in algebraic VoF methods. In CSF the surface tension  $\mathbf{f}_\gamma$  is modeled as a localized body force within a transition region of finite thickness at the interface

$$\mathbf{f}_\gamma = -\gamma\kappa\nabla\alpha, \quad (2.24)$$

with the surface tension coefficient  $\gamma$  and the cell-averaged curvature of the interface  $\kappa$ . The curvature of the interface is approximated by  $\kappa = -\nabla \cdot \mathbf{n}_i$  with the outwards pointing normal vector of the liquid interface  $\mathbf{n}_i$ , which is represented by the gradient of the volume fraction  $\mathbf{n}_i = \frac{\nabla\alpha}{\|\nabla\alpha\|}$ . Consequently, this results in

$$\mathbf{f}_\gamma = \gamma\nabla \cdot \left( \frac{\nabla\alpha}{\|\nabla\alpha\|} \right) \nabla\alpha. \quad (2.25)$$

Spurious currents, *i.e.* artificial velocities that distort the solution of the velocity field even if both fluids are at rest are a common issue with most surface tension models. This issue arises if the numerical implementation of surface tension does not balance with the pressure gradient forces when discretized or if the curvature of the interface is not accurately calculated. In the CSF model the correct approximation of the surface curvature and, consequently, the surface tension force are dependent on the accurate distribution of the volume fraction.

**Phase field methods** The phase field method [100, 261] is a diffuse interface method, in which the finite thickness of the interface is physically motivated by long-range van der Waals forces. The evaluation of the phase distribution is based on the conserved order parameter  $C$  that represents both phases and takes the value of  $C_l = 1$  in the liquid phase and  $C_g = -1$  in the gaseous phase. The order parameter  $C$  varies rapidly but smoothly at the interface and defines the interface location at  $C = 0$ . In the phase-field approach, the material properties in the bulk phases are considered to be homogeneous, while the components of the bulk phases are mixed in the interface region with a smooth and continuous transition. Consequently, the interface region contains mixing energy and the interface thickness is determined by molecular dynamics. The mixing energy density describes the molecular interaction of the two phases and is defined as

$$f_{\text{mix}}(C, \nabla C) = \frac{1}{2}\lambda_{\text{mix}}\|\nabla C\|^2 + \frac{\lambda_{\text{mix}}}{4\epsilon^2}(C^2 - 1)^2, \quad (2.26)$$

where  $\lambda_{\text{mix}}$  is the magnitude of the mixing energy and  $\epsilon$  is the capillary width, which is proportional to the interface thickness. The first term in  $f_{\text{mix}}$  is the free energy of the interface and represents weakly non-local interactions of the chemical components that drive the mixing of the components at the interface. The second term represents the bulk energy density that drives the separation of the two phases. the mixing energy  $F_{\text{mix}}$  is determined by integration of  $f_{\text{mix}}$  in the control volume  $V$

$$F_{\text{mix}} = \int_V f_{\text{mix}} dV. \quad (2.27)$$

According to the diffuse-interface theory [234] the phase separation and diffusion in two-phase flows are driven by the chemical potential at the interface. The chemical potential  $\psi$  can be obtained as the variational derivative of the mixing energy with respect to the order parameter  $C$

$$\psi = \frac{\delta F_{\text{mix}}}{\delta C} = \frac{\lambda_{\text{mix}}}{\epsilon^2}C(C^2 - 1) - \lambda_{\text{mix}}\nabla^2 C. \quad (2.28)$$

Following the generalized Fick's law, the diffusive mass flux of the two-phase system can be considered proportional to the gradient of the chemical potential, which leads to the convective Cahn-Hilliard equation [27]

$$\frac{\partial C}{\partial t} + (\mathbf{u} \cdot \nabla) C = M \nabla^2 \psi, \quad (2.29)$$

with the mobility parameter  $M$  that determines the relaxation time of the interface. The Cahn-Hilliard equation represents the dynamics of the interface and is conservative with respect to  $C$  for incompressible fluids. The mixing energy density represents the interfacial tension and, therefore, can be related to the surface tension coefficient  $\gamma$  at the sharp interface limit  $\epsilon \rightarrow 0$  for an equilibrium state of the two-phase system. Equating the surface energy and the mixing energy in the interface region yields [261]

$$\gamma = \frac{2\sqrt{2}}{3} \frac{\lambda_{\text{mix}}}{\epsilon}. \quad (2.30)$$

As the surface tension coefficient  $\gamma$  can be measured by experiments, Equation 2.30 can be used to determine the mixing energy magnitude  $\lambda_{\text{mix}}$ . The value for  $\epsilon$ , however, needs to be chosen in the numerical simulation and is typically defined in relation to the characteristic macroscopic length scale of the flow. A sufficiently high resolution of the computational mesh is required in order to represent the physics at the interface, which usually causes the numerical  $\epsilon$  to be much larger than the physical  $\epsilon$ . At the sharp interface limit, the diffusive term of Equation 2.29 vanishes and the mobility parameter  $M$  is scaled accordingly with  $M = \zeta \epsilon^2$ . The value for  $\zeta$  is obtained by fitting to experimental data.

In the phase field method the two-phase flow dynamics are modeled by the coupled Cahn-Hilliard and single-field Navier-Stokes equations (see Equation 2.22). The material properties of both phases are determined by the mixture rule

$$\xi_M = \frac{1+C}{2} \xi_l + \frac{1-C}{2} \xi_g \quad \text{with} \quad \xi \in \{\rho, \mu\}. \quad (2.31)$$

The surface tension term in the NSE is represented by a formulation based on the chemical potential, *i.e.* the continuum surface tension in the potential form [100]

$$\mathbf{f}_\gamma = \psi \nabla C. \quad (2.32)$$


The phase field method is advantageous for the simulation of phenomena involving contact line motion, as the stress singularity at the contact line, which occurs for sharp interface methods such as VoF, is avoided by the representation of the phase evolution as a diffusive process. Furthermore, the energy-based formulation of the phase distribution handles topological changes natively, which makes the phase field method suitable for the simulation of coalescence and break-up. However, mass conservation is not ensured, which can result in unphysical solutions at small scales and high density ratios.

## 2.3 Optical measurement principles

In this section, the fundamental optical measurement principle of the shadowgraphy technique and the theoretical framework for glare points that result from the scattering of light will be discussed, which provides the basis for the development of experimental methods in this work.

### 2.3.1 Shadowgraphy technique

“The most common optical measurement method for the observation of droplet dynamics, due to its simplicity and high spatial accuracy, is the shadowgraphy technique, in which a droplet or bubble is illuminated in parallel backlight that accurately maps the contour of the gas-liquid interface onto an image plane [237, 162].” from [DKS24, Section 1] The shadowgraphy technique is a density-based optical measurement technique that exploits the refraction or reflection of light due to density changes in the fluid to visualize its flow. In the case of two-phase flows, the primary mechanism is the reflection of light at the gas-liquid interface. In the simplest form of shadowgraphy, only a light source and a screen are required to illuminate the object in backlight, which casts a shadow on the screen. However, in practice a more elaborate experimental setup, as illustrated in Figure 2.4, is commonly employed in order to achieve a higher accuracy of the projected contour.

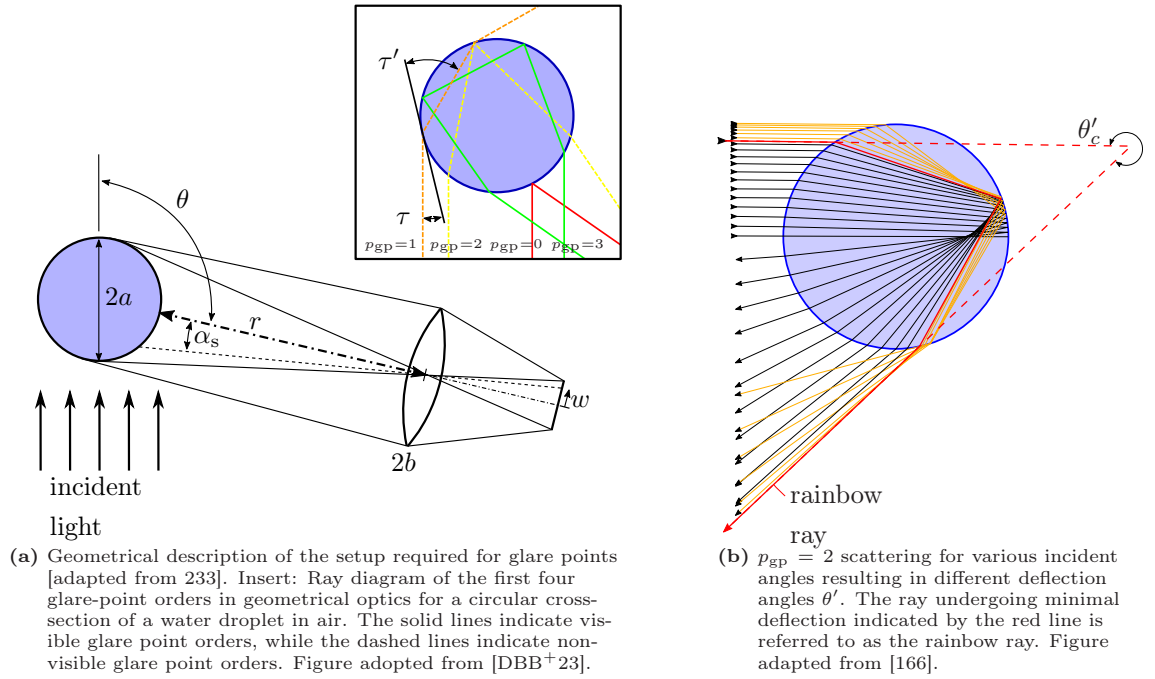


**Figure 2.4:** Sketch of the parallel light focused shadowgraphy; a collimated light beam illuminates the droplet in the test section, casting a shadow of its contour onto the recording plane. The shadowgraph image is focused by the objective lens.

As shown in Figure 2.4, the light emitted by an illumination source is focused through a condenser lens and subsequently passes a pinhole aperture, after which the light is collimated into a parallel beam by a second lens. Parallel light provides a significantly higher contrast in comparison to the divergent light of an uncollimated light source due to a more localized illumination. Therefore, parallel light shadowgraphy is the preferred choice for all applications in which collimation of the light is feasible [201]. The illumination setup depicted in Figure 2.4 is designed to enhance the accuracy of the projected contour by managing the adverse influences of geometric blur and diffraction blur. Geometric blur arises due to the divergence of the light beam that results from the finite size of the light source. Consequently, a point light source would be ideal from the viewpoint of geometrical optics. However, diffraction sets a lower limit for the size of the light source, beyond which visible diffraction fringes appear in the shadowgraph images that complicate a differentiation of the object contour. Furthermore, a bright illumination is required to attain a sufficient contrast between the shadowgraph and the background. Satisfactory results with respect to both sharpness and contrast are achieved through the aforementioned setup consisting of a light source that is focused by a condenser lens, with a pinhole aperture placed in the focal point of the condenser lens to obtain a small, yet bright light source, which is subsequently collimated by a

second lens. Various light sources, including light-emitting diode (LED) lights, are permissible for shadowgraphy. However, light with a broad spectrum offers the advantage of a reduced diffraction blur [201]. In order to obtain a more homogeneous background illumination an optical diffuser can be placed between the condenser lens and the pinhole. The imaging optics, comprising an objective lens and a recording plane as illustrated in Figure 2.4, enable the recording of shadowgraph images at different focal planes within the test section and at varying magnifications. In particular, an objective lens with variable focal length facilitates the calibration of the optical setup, as the resolution of the shadowgraph contour can be adjusted while observing the recorded images.

### 2.3.2 Glare points



**Figure 2.5:** Scattering of light at a spherical water droplet in air in the framework of geometrical optics.

“The formation and characteristics of glare points in both the theory of geometrical optics (GO) and wave optics [15] will be briefly recapitulated here in the context of droplets. Glare points are the bright spots on the surface of the droplet that result from the scattering of a wide beam of parallel light according to van de Hulst [232]. In GO, glare points are described as the exit points of light rays that are either reflected or refracted on the surface in the direction of scattering  $\theta$  [111], as can be seen in Figure 2.5a. This results in the emergence of different orders of glare points, defined by the chords traveled within the droplet  $p_{gp}$  [45], where  $p_{gp} = 0$  indicates the externally reflected light,  $p_{gp} = 1$  the transmitted and twice refracted light rays and  $p_{gp} = 2$  the refracted and internally reflected light rays.

As indicated in Figure 2.5a by the parallel  $p_{gp} = 0$  (red) and  $p_{gp} = 3$  (green) rays, multiple orders of glare points can be visible simultaneously for the same scattering angle. Conversely, there exist ranges of scattering angles at which certain orders of glare points are not visible at all, [as indicated by the absence of any scattered  $p_{gp} = 2$  rays beyond the red ray in Figure 2.5b]. According to van de Hulst and Wang [233] all relevant rays for the formation of glare points at a perfect sphere propagate within the scattering plane, defined by the direction of the incident light and the scatter direction. Furthermore, the GO theory of glare points is applicable for volumes with a plane of



circular cross-section and, therefore, to spheroids and cylinders as well. The description of glare points in GO is valid for size parameters  $x_s = (2\pi a)/\lambda \gg 1$ , with  $2a$  describing the droplet diameter and  $\lambda$  the wavelength of the incident light. This holds true for  $2a \approx 2$  mm and  $\lambda = 400 - 700$  nm in the experiments conducted in the present work. The required equations for the calculation of the glare point positions in GO will be summarized in the following. The angle between an incident light ray and the surface of the droplet  $\tau$  can be set in relation to the angle between the refracted ray and the surface  $\tau'$  according to

$$n_{r,1} \cos(\tau) = n_{r,2} \cos(\tau') \quad (2.33)$$

considering Snell's law [see *e.g.* 83] with the refractive indices  $n_{r,1}$  of the surrounding gas and  $n_{r,2}$  inside the droplet. Considering the possible cases of external reflection,  $p_{gp} - 1$  internal reflections, and the refraction of a light ray, the deflection angle  $\theta'$  can be determined. Subsequently, the scattering angle  $\theta$  can be calculated to be

$$\theta' = 2(\tau - p\tau') = 2\pi k + q\theta, \quad (2.34)$$

where  $k$  is an integer, depending on the number of internal reflections [232, pp.228ff] and  $q = \pm 1$ , depending on the side of the projected droplet relative to the optical axis, which can be described by the variable  $w = \alpha_s l/a$ . This dimensionless quantity  $w$  describes the position on the image plane relative to the droplet's projection as a function of the droplet diameter  $2a$ , the distance between the lens and the droplet  $l$ , and the angle  $\alpha_s$  between the optical axis and the position on the droplet, as seen from the lens perspective (see Figure 2.5a). The position of the glare point in the dimensionless coordinate  $w$  follows as

$$w = q \cos(\tau). \quad (2.35)$$

Note that Equations (2.34) and (2.35) only hold for a spherical droplet. Note, however, that a GO description for non-spherical droplets, which are symmetric around the optical axis, was developed by [94].

The scattering angle is non-monotonically dependent on the incident angle for  $p_{gp} \geq 2$ , as illustrated in Figure 2.5b]. Therefore,  $\theta(\tau)$  possesses an extremum at the rainbow angle  $\theta_c$  [238], which can be determined by  $d\theta/dw = 0$  [232]. The deflection of the incident light rays  $\theta'$  is minimal at the rainbow angle and increases towards both larger and smaller scattering angles. As such, the density of the light rays exiting the droplet reaches a maximum at  $\theta_c$ , which results in the highest scattered intensity at the rainbow angle. The scattering angle  $\theta$  can either be maximal or minimal at  $\theta_c$ , depending on the order of the rays  $p_{gp}$ . A further consequence is that no scattered rays beyond the rainbow angle are possible. Note that the two visible rainbows in nature can be explained accordingly – rays of the order  $p_{gp} = 2$  create the primary rainbow and  $p_{gp} = 3$  rays produce the weaker secondary rainbow above the primary rainbow. Higher orders of rainbows become increasingly dim as their incident angle increases such that less droplet surface is illuminated. Furthermore, the light is partly refracted at every reflection. Due to dispersion, the polychromatic light of the sun is scattered at slightly different angles around the rainbow angle, creating the well-known color gradients. For  $p_{gp} = 2$  the rainbow angle is the minimal possible scattering angle and for  $p_{gp} = 3$  it is the maximum scattering angle, which results in inverted color gradients and also explains the appearance of Alexander's dark band between the rainbows in which no light is scattered [166].

The complete description of glare points for arbitrary-sized spheres is given in wave optics through Lorenz-Mie theory, which is the exact solution for the problem of light scattering on a spherical



particle [148]. van de Hulst and Wang [233] derive the equation for the amplitude of the glare points with the following assumptions. Firstly, a well-defined scattering plane, that contains the direction of incidence and the direction of scattering is assumed. Consequently, the scattering plane excludes the  $\theta = 0^\circ$  and  $180^\circ$  scattering angles. Secondly, the authors assume that the lens both follows the thin lens equation [83] and is placed in the far field. The glare point equation accordingly follows as

$$A_{\parallel\perp}(w) = \int_{\theta_0-b/l}^{\theta_0+b/l} S_{\parallel\perp}(\theta) \exp[-ix_s w (\theta - \theta_0)] d\theta, \quad (2.36)$$


in which  $S_{\parallel\perp}(\theta)$  is the electrical field amplitude of the scattered light,  $\theta_0$  is the scattering angle and  $2b$  is the diameter of the lens.  $S_{\parallel\perp}(\theta)$  represents either parallelly  $S_{\parallel}$  or perpendicularly  $S_{\perp}$  polarized light. Consequently, the size of glare points arising from a plane wave impinging on the droplet will depend on the size (and shape) of the receiving aperture. A comprehensive derivation of the scattering on spherical particles is provided in the treatise of [232].” from [DBB<sup>+</sup>23, Section 2.1]

## 2.4 Deep Learning

In this section, first, the fundamentals of neural networks and deep learning will be briefly introduced in Subsections 2.4.1 and 2.4.2, which lay the theoretical foundations for convolutional neural networks (CNN), a type of neural network that is especially useful for processing image data, which are described in Subsection 2.4.3. Afterward, a brief introduction to physics-informed neural networks in Subsection 2.4.4 is presented. The following section contains a revised and extended version of the literature review in the Master thesis by the author [51] and features directly adopted and paraphrased passages from Chapter 3 of the aforementioned work.


In the context of deep learning an algorithm improves its ability to make a correct prediction by the data-driven optimization of a neural network on a specific set of data. As described in [74], this approach makes it possible to solve complicated tasks that would otherwise require complicated manually designed algorithms or are even practically impossible. Recently, deep learning techniques have increasingly become commonplace in science, industry, and everyday life. Large language models, such as ChatGPT [168], Gemini [70], or Llama [229] and generative artificial intelligence models, like DALL-E [178] and Stable Diffusion [186] are rapidly transforming the workflow in a multitude of different industries. Moreover, an increasing number of scientific discoveries are achieved with the assistance of deep learning methods, including breakthroughs in the likes of AlphaFold [104] for the prediction of protein structures. These recent successes can be attributed to a strong research interest driven by breakthroughs in various challenges in the field of computer science, particularly computer vision [121], speech recognition [78], machine translation [215], and natural language processing [235]. Deep learning methods are increasingly adopted in the scientific study of fluid mechanical problems, including turbulence modeling, PIV processing, and flow field super-resolution. For an in-depth review regarding the application of Machine Learning methods in fluid mechanical research, the reader is referred to the review papers of Brunton et al. [18] and Vinuesa and Brunton [236]. The success of deep learning techniques is related to their data-driven nature, as the performance of neural networks scales well with the amount of training data used in their optimization [213]. The aforementioned high research interest has led to the continuous development of efficient optimization algorithms as well as effective architectures for neural networks that have further facilitated the practical applicability of deep learning techniques.

### 2.4.1 Neural networks



**Figure 2.6:** Schematic of a fully-connected feedforward neural network consisting of an input layer, two hidden layers with four neurons each, and an output layer. The input layer receives the inputs  $X_i$ , which are processed by the neuron in the hidden layers, yielding the prediction  $\hat{Y}$  in the output neuron.

Neural networks consist of a multitude of small non-linear units, the so-called neurons that transform an input into a more abstract output. These neurons are interconnected and stacked in multiple consecutive layers, as illustrated in Figure 2.6, which allows the neural network to learn various non-linear functions.



**Figure 2.7:** Schematic representation of the process inside a neuron. The input  $\mathbf{X}$  is multiplied with the weights  $\mathbf{W}$ , a bias term  $B$  is added to the weighted sum, after which non-linear activation  $\hat{Y} = a(Z)$  is applied.

A neuron, shown schematically in Figure 2.7, processes an arbitrary amount of input values  $\mathbf{X}$  into one output  $\hat{Y}$  using the affine transformation  $Z = \mathbf{W} \cdot \mathbf{X} + B$ , followed by a nonlinear activation function  $\hat{Y} = a(Z)$ . The weights  $\mathbf{W}$  and bias  $B$  comprise the learnable parameters of a neuron and are, therefore, subject to change during the optimization of the neural network. In the following, all learnable parameters of the network are referred to as weights, implicitly including the biases. As illustrated in Figure 2.6 the neurons are ordered in successive layers, where information is passed in a consecutive order. The flow of information inside a feedforward neural network is unidirectional and any amount of the outputs from neurons in the previous layer can be considered as inputs for a neuron in the current layer. In a fully-connected layer (FC) all neurons between two layers are interconnected, which extends to all layers for a fully-connected neural network. A fully-connected feedforward neural network with multiple hidden layers is referred to as an multi-layer perceptron (MLP). The parameters of the hidden and output neurons are not defined in the network's design, but instead are randomly initialized and then changed during the training of a neural network by an optimization algorithm. By combining multiple non-linear units into a network, arbitrarily complicated non-linear functions can be learned. Specifically, neural networks are considered universal function approximators, as any continuous function can be approximated at arbitrarily close accuracy by an MLP with just one hidden layer and a finite number of neurons, according to the universal approximation theorem by Cybenko [42] and Hornik et al. [93]. Thereby,

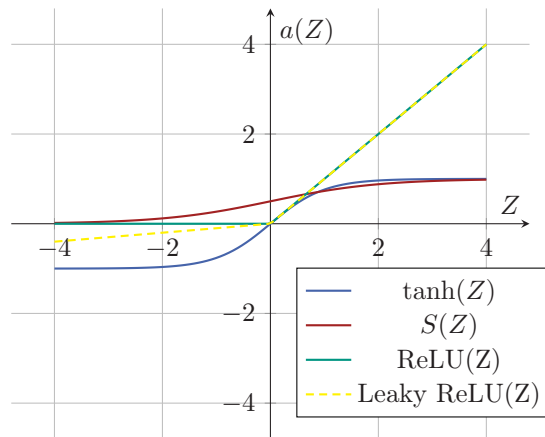
neural networks offer a high expressivity, rendering deep learning approaches useful for scientific computation. Furthermore, neural networks are able to accurately represent the underlying non-linear behavior in various physical systems [41].

**Activation functions** The activation function  $a(Z)$  introduces non-linearity to the neural network, which is necessary to learn functions of higher complexity, as the chained affine transformations of the stacked layers without any non-linearity would collapse into a single affine transformation. The most frequently used activation functions are exponential functions, such as the hyperbolic tangent (eq. 2.37) and sigmoid (eq. 2.38) functions and piecewise linear functions, such as the rectified linear unit (ReLU) (eq. 2.39) and leaky-ReLU (eq. 2.40). Furthermore, sine and exponential functions are employed as activations.

$$\tanh(Z) = \frac{e^Z - e^{-Z}}{e^Z + e^{-Z}} \quad (2.37) \quad S(Z) = \frac{1}{1 + e^{-Z}} \quad (2.38)$$

$$\text{ReLU}(Z) = \begin{cases} 0 & \text{if } Z \leq 0 \\ Z & \text{if } Z > 0 \end{cases} \quad (2.39) \quad \text{leaky ReLU}(Z) = \begin{cases} \Gamma Z & \text{if } Z \leq 0 \\ Z & \text{if } Z > 0 \end{cases} \quad (2.40)$$

The hyperbolic tangent function is commonly employed for hidden units, while the sigmoid function is mostly used as an activation function for output layers, where a bounded activation between zero and one is expected. Exponential activation functions may encounter the issue of saturation for high values of  $Z$ , which leads to a diminished sensibility to the inputs and, consequently, hinders the training through gradient-based optimization techniques. The ReLU activation function circumvents this saturation effect due to its (piecewise) linear nature, which has made ReLU the most common activation function for hidden units in contemporary neural networks. Furthermore, it has been shown that ReLU significantly enhances the convergence rate of neural network optimization [159]. The disregard of negative information, however, can lead to the so-called dying ReLU problem, in which the inputs of the neuron consistently yield negative values for  $Z$ , leading to output and gradient values of zero. Consequently, the weight parameters cannot change anymore during backpropagation and the neuron stops learning. This problem is addressed by the leaky ReLU function, which introduces a small constant  $\Gamma$  that ensures the existence of a non-zero gradient for negative values of  $Z$  [139]. The aforementioned activation functions are plotted for values of  $Z$  close to the origin in Figure 2.8.



**Figure 2.8:** Comparison of the hyperbolic tangent, sigmoid, rectified linear unit (ReLU), and leaky ReLU activation functions in proximity to the origin.

## 2.4.2 Optimization techniques

In deep learning, two paradigms for the optimization of a neural network are distinguished, namely supervised and unsupervised learning. In the approach of supervised learning the network learns to make correct predictions by training on labeled data, that is input data  $\mathbf{X}$  which has the expected output values  $Y$  associated with it. In unsupervised learning, the network only receives unlabeled input data and learns to detect the underlying structure or patterns in the data by extracting semantic features from it. Both approaches, however, employ gradient-based optimization schemes for the training of the neural networks, most commonly stochastic gradient descent (SGD) [184] or variants of it.

The data-driven learning of a transfer function is achieved by iteratively updating the weights of the network in two steps, described by the forward and backward propagation on the computational graph of the neural network [74, pp. 164ff.]. The calculations performed by the neural network can be expressed through a computational graph that details the dependence and order of the elemental operations in the neurons. The forward propagation on the computational graph yields the prediction  $\hat{Y}$  on the basis of the given inputs  $\mathbf{X}$ . The deviation of the network's prediction from the ground truth label  $Y$  is expressed by the loss function  $\mathcal{L}(Y, \hat{Y})$  that serves as the objective function for the network's optimization. The most common loss function is the mean squared error (MSE) loss

$$\mathcal{L}_{\text{MSE}}(Y, \hat{Y}) = \frac{1}{n} \sum_{i=1}^n (Y_i - \hat{Y}_i)^2, \quad (2.41)$$

where  $n$  is the number of samples for which the loss is computed simultaneously. The neural network is trained for correct predictions on a given dataset by minimizing the loss function, which results in a high-dimensional optimization problem. During backward propagation, the loss function is differentiated with respect to the weights of the neural network through the application of the chain rule on the computational graph. The respective gradient calculated for each weight parameter of each neuron expresses by which amount the objective function changes for an infinitesimal increase of the particular weight and, thereby, indicates the necessary change of the weight for achieving an improved prediction. The updated weight  $W^* = W - \alpha_{\text{lr}} \frac{d\mathcal{L}}{dW}$  is calculated by scaling the gradient  $\frac{d\mathcal{L}}{dW}$  with the learning rate  $\alpha_{\text{lr}}$  and subtraction from the weight  $W$  of the previous iteration in the optimization. In this fashion, updated weights for all neurons are calculated resulting in a more optimized network in each training iteration. The process is repeated until no significant change in loss can be observed anymore, *i.e.* convergence of the optimization process. Here, the properties of different activation functions become obvious. The saturation issue of the hyperbolic tangent activation function leads to vanishing gradients, which slows down the optimization. Furthermore, the evaluation of the gradient for the piece-wise linear activation function ReLU is greatly simplified in comparison to the exponential activation functions, which further accelerates the training. However, the hyperbolic tangent activation function can be differentiated infinitely many times, which is a required property for some applications, such as physics-informed neural networks, which employ tanh instead of ReLU.

The minimization of the objective function is a non-convex optimization problem, which requires a fine tuning the learning rate  $\alpha_{\text{lr}}$  for each problem. An insufficiently low learning rate results in a slow convergence and can cause the optimization algorithm to get trapped in a local minimum, while an excessively high learning rate can result in oscillation or divergence due to overshooting the minimum in every step. During training with stochastic gradient descent (SGD), batches of training data are used to compute the loss function instead of single examples, in order to reduce oscillations


in the training due to outliers in the dataset and improve convergence. Furthermore, extensions of SGD, such as stochastic gradient descent with momentum (SGDM) [171], root mean square propagation (RMSProp) [225], and adaptive moment estimation (Adam) [113] were developed to enhance the training dynamics of neural networks with the aim of a faster convergence. Further details on the training dynamics of neural networks can be found in [74, pp.442ff] and [126, pp.16ff].

During the training of neural networks by means of the aforementioned gradient-based optimization schemes two major sources of errors are considered, namely the optimization error and the generalization error [205]. The optimization of the loss function is a high-dimensional non-convex problem and, thus, no guarantees exist that the global minimum will be reached. In practice, the optimization will likely settle on a local minimum, which gives rise to an optimization error, which is dependent on the initialization of the network weights and the hyperparameters of the optimization scheme. The generalization error is a measure of the accuracy of the prediction for unseen data and consists of two contributing error sources, the approximation error (bias) and the estimation error (variance) [6]. In the theoretical limit of infinitely wide layers, neural networks are universal function approximators [42, 93]. However, in practice, the limited size and capacity of neural networks give rise to an approximation error. The estimation error results from the finite amount of available training data, which is a random sample of the complete data distribution. Consequently, the transfer function learned by the neural network from the training data is an estimation of the expected function, which introduces an estimation error [163]. Furthermore, the approximation and estimation errors are linked in the data-driven estimation, as sufficient amounts of data are required to optimize a large number of variable parameters in the neural network, which leads to a trade-off between bias and variance. Consequently, the choice of the network's capacity depends not only on the given problem but also on the amount and quality of available training data. An insufficient amount of training data, or an excessive capacity of the network, can lead to overfitting, which is characterized by the memorization of the training data by the network and a loss of its generalization ability. Regularization methods, such as Dropout [211],  $L_2$ -regularization, and Batch Normalization [98], can be used to alleviate this problem. A further technique that is employed to improve the optimization process is the so-called learning rate decay, in which the learning rate is reduced after a certain amount of training iterations. Learning rate decay effectively combines the benefits of an initially high learning rate that prevents the optimization scheme from getting stuck in a suboptimal local minimum early, with the benefit of reducing oscillations around a local minimum later in the optimization process through lower learning rates. Thereby, the convergence to an optimal local minimum is supported.

### 2.4.3 Convolutional neural networks

Convolutional neural networks are a type of neural network that has been proven to be especially useful for multidimensional data with an underlying structure, such as images, videos, or sequential data. The two main properties of CNNs are the preservation of these implicit structures and the sparsity of connections that result from the use of convolutional layers instead of fully-connected layers. In the following, the processing of image data characterized as matrices of pixel values in one (grayscale) or multiple channels (RGB) is considered.

**Spatial convolution** Convolutional layers preserve the spatial information of the image data by simultaneously processing small regions of neighboring pixels (commonly 3x3 px, 5x5 px, or 7x7 px) in a sliding-window manner. Particularly, the values of the output layer are calculated by the spatial convolution of learnable filters with the input image, as illustrated in Figure 2.9.



**Figure 2.9:** Spatial convolution of a 3x3 filter (gray cells) over a 5x5 input (blue cells) with padding (white dashed cells) to conserve the spatial dimensions in the output feature map (green cells). Figure adopted from [52].

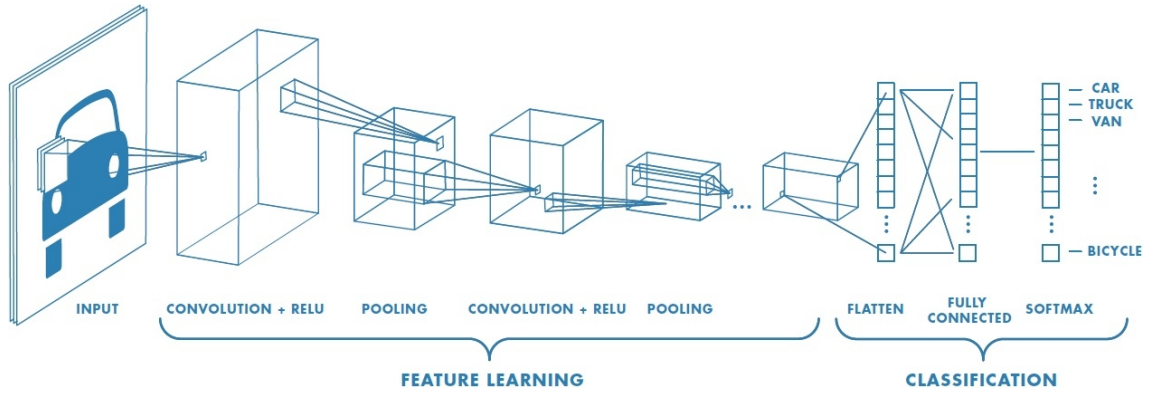
$$\mathcal{F}[i, j] = \sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} H[m, n] \cdot X[i - m, j - n] \quad (2.42)$$

The convolution operation (Equation 2.42) calculates an element-wise multiplication of the filter  $H$  with the spatial dimensions  $m$  and  $n$  and an input region  $X$  of the same size, producing an output feature map  $\mathcal{F}$ . This can be interpreted as local pattern recognition; a high similarity between the filter and the input image yields a high activation of the neuron. The parameters of the filter  $H$  are learnable weights, which are optimized through backpropagation. Analogously to fully-connected (FC) layers a bias term is added to the weighted sum of the input, and the result is further transformed by a non-linear activation function. A convolutional layer is comprised of multiple convolutional filters, which are trained for the recognition of different features, such as edges, patterns, textures, or shapes. The sliding-window approach provides spatial invariance to the neural network and, thereby, allows for the recognition of shapes in different parts of the image. The weights for the interaction between two layers in the network are shared by convolving the filter over the whole input. This leads to a sparse connection between the layers and a reduced amount of parameters compared to a fully-connected layer, mitigating the risk of overfitting, and reducing computation costs.

**Spatial pooling** The second common building block of CNNs are pooling layers, which are used to downsample the feature maps, by applying a pooling function to patches of the input in a sliding-window approach. The most prevalent pooling function is maximum pooling, which outputs the maximum value over a spatial region, as it has been found to offer the highest performance [197]. In contrast to convolutional layers, the pooling operation is a predefined function without learnable parameters. By reducing the dimensions of the feature map, the number of parameters in the following layers is drastically reduced, *e.g.* a 2x2-filter without overlap reduces the number of parameters by 75%. These properties reduce computational costs and mitigate overfitting, allowing for deeper neural networks.


**Basic convolutional neural networks** Convolutional neural networks are comprised of multiple sequential convolutional layers and pooling layers that progressively reduce the dimensions of the feature maps, which is commonly followed by several fully-connected layers, depending on the specific task. Figure 2.10 shows a sketch of a simple CNN for object classification. Each neuron in a convolutional layer has a so-called receptive field in the input image, which corresponds to the filter size in the first convolutional layer and successively increases in each following layer. Through the progressive processing of the input through non-linear transformations in the convolutional layers





**Figure 2.10:** Schematic of a simple convolutional neural network consisting of multiple consecutive and alternating convolutional and pooling layers with fully-connected layers appended for classification. Figure adopted from [144].


a feature hierarchy emerges, in which the semantic information increases in each successive feature map while its resolution decreases. Thereby, early layers learn to identify simple features like edges and later layers learn increasingly complicated and large features. As shown by Lecun et al. [125] CNNs exhibit to some degree translational, rotational, and scale invariance, and are robust to variation in illumination intensity, which makes these networks effective for image processing tasks.



**Figure 2.11:** Schematic of a residual block: the inputs  $X$  of a previous layer are added to the weighted inputs of the current layer  $\mathcal{F}(X)$  followed by a non-linear activation  $a(Z)$  calculated over the sum  $\mathcal{F}(X) + X$ . Figure adapted from [80].

**Residual CNNs** The performance of CNNs can be improved by more complex architectures and a higher amount of trainable parameters if a sufficient amount of data is available [217, 208]. However, for very deep networks vanishing or exploding gradients during backpropagation become an issue, leading to instability of the optimization and can eventually result in divergence. He et al. [80] introduced residual blocks that contain skip connections, as illustrated in Figure 2.11, which improve cross-layer connectivity and, thereby, alleviate the aforementioned gradient pathologies. Furthermore, residual connections facilitate the learning of an identity mapping between the layers, which allows information to propagate to later layers and encourages the reuse of features. Thereby, the neural networks becomes more adaptive towards different learning tasks and the network is less likely to overfit the training data.

**U-Net architecture** In applications that perform predictions on a local level, such as object detection [133], segmentation [187] or 3D-pose estimation [160] a high spatial resolution of the extracted feature maps is required. However, in the previously discussed CNNs, an increase in semantic value in the emerging feature hierarchy is associated with a reduction of the spatial resolution. Multiple approaches [187, 160, 133] have been developed to alleviate this problem by combining layers of different scales to enrich lower-level feature maps with the semantic value from higher-level feature maps. While these approaches differ in detail, the general concept is an extension of the contracting network architecture by a second expanding pathway. Particularly, in the *U-net*



**Figure 2.12:** Schematic of the *U-net* architecture constituted by a contracting pathway (left half) of consecutive convolution and pooling operations and an expanding pathway (right half) of convolution and up-sampling operations with lateral connections (gray arrows) for feature aggregation on multiple scales. Figure adopted from [187].

approach, additional convolutional layers are appended to the network, but instead of pooling layers an upsampling operation, such as nearest neighbor upsampling is integrated in-between the convolutional layers in an alternating manner, as can be seen in Figure 2.12. Furthermore, the contracting and expanding pathways are interconnected by lateral connections. Thereby, fine details from early feature maps with a high resolution are introduced to the later feature maps, allowing for a prediction on the basis of feature maps with both high semantic values and high resolution.

## 2.4.4 Physics-informed neural networks

Physics-informed neural networks (PINNs) [176] were introduced as universal function approximators for problems that can be described by non-linear partial differential equations (PDEs). The governing equations are encoded in a neural network by training on the residuals of the PDEs alongside the available data and, thereby, impose a regularization to the network optimization that constrains the space of possible solutions to the ones that adhere to the governing equations. Vanilla PINNs as introduced by Raissi et al. [175] consider a system of parameterized partial differential equations:

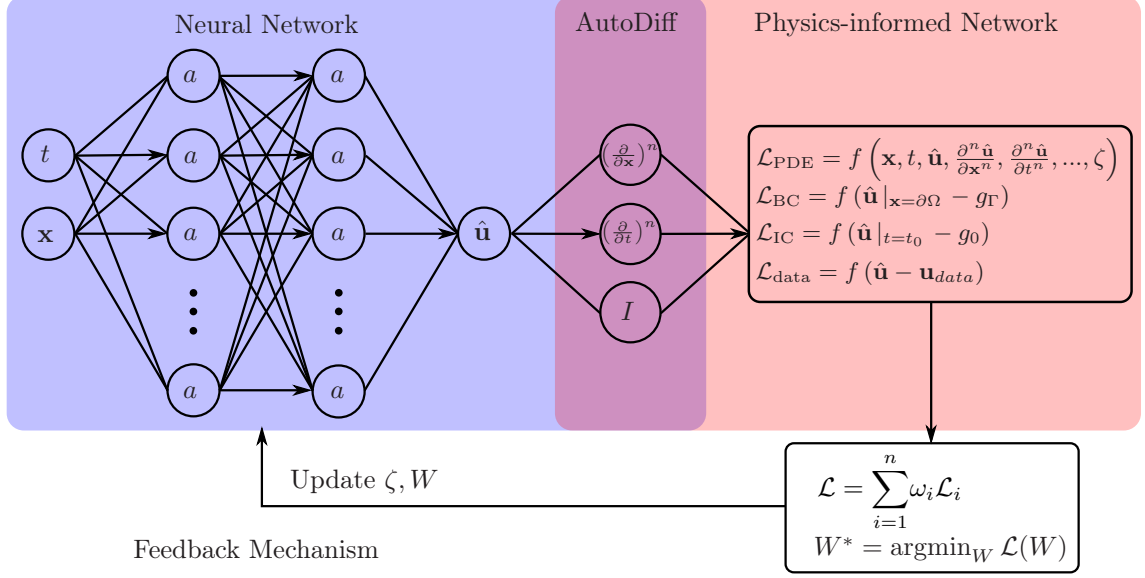
$$\mathcal{R}\left(\mathbf{x}, t, \mathbf{u}, \frac{\partial^n \mathbf{u}}{\partial \mathbf{x}^n}, \frac{\partial^n \mathbf{u}}{\partial t^n}, \dots, \zeta\right) = 0, \quad \mathbf{x} \in \Omega, \quad t \in [0, T], \quad (2.43)$$

$$\mathbf{u}(\mathbf{x}, t_0) = g_0(\mathbf{x}), \quad \mathbf{x} \in \Omega, \quad (2.44)$$

$$\mathbf{u}(\mathbf{x}, t) = g_\Gamma(t), \quad \mathbf{x} \in \partial\Omega, \quad t \in [0, T], \quad (2.45)$$

in which  $\mathbf{x} \in \mathbb{R}^d$  are the spatial coordinates,  $t$  the time coordinate and  $\mathcal{R}$  the residuals of the PDEs, which depend on the derivatives of the solution  $\mathbf{u}(\mathbf{x}, t)$  to the PDEs with respect to space  $\frac{\partial^n \mathbf{u}}{\partial \mathbf{x}^n}$  and time  $\frac{\partial^n \mathbf{u}}{\partial t^n}$ .  $\mathbf{u}(\mathbf{x}, t)$  is the solution of the PDEs with the initial condition  $g_0(\mathbf{x})$  on the spatial domain





**Figure 2.13:** Schematics of physics-informed neural networks (PINNs). A neural network is trained to predict the solution  $\mathbf{u}(\mathbf{x}, t)$  of partial differential equations (PDEs) based on the spatio-temporal input coordinates  $\mathbf{x}$  and  $t$ . The physics-informed network computes the derivatives of  $\mathbf{u}$  with respect to  $\mathbf{x}$  and  $t$  to calculate the residuals of the PDEs. The parameters  $\mathbf{W}$  of the neural network are optimized by the minimization of a composite loss calculated from the residuals of the PDE, as well as initial and boundary conditions or available measurement data.

$\Omega$  and the boundary conditions  $g_\Gamma(t)$  on the boundary of the domain  $\partial\Omega$ .  $\zeta$  represents the known or unknown parameters of the PDEs. In the context of fluid mechanics, the relevant PDEs could be the Navier-Stokes equations with the solution  $\mathbf{u} = (u, v, w, p)$  for the velocity field  $(u, v, w)^\top$  and pressure  $p$ , and parameters  $\zeta$  corresponding to the density or dynamic viscosity of the fluid. In the PINN approach, a neural network is trained to approximate the solution  $\hat{\mathbf{u}}(\mathbf{x}, t) \approx \mathbf{u}(\mathbf{x}, t)$  through the minimization of a composite loss function, which is constituted by the residuals of the PDEs, available data, as well as the initial and boundary conditions. As illustrated in Figure 2.13 PINNs consist of three main components, a neural network, a physics-informed network, and a feedback mechanism. During training, the neural network, parameterized by the weights  $\mathbf{W}$ , predicts the solution of the PDEs  $\hat{\mathbf{u}}(\mathbf{x}, t)$  from the spatio-temporal input coordinates  $\mathbf{x}$  and  $t$ . Subsequently, the physics-informed network computes the derivatives of  $\hat{\mathbf{u}}(\mathbf{x}, t)$  with respect to the input coordinates, from which the physics-informed loss terms are calculated. In the final step, a feedback mechanism updates the weights of the neural network through the minimization of the physics-informed composite loss function, defined as

$$\mathcal{L} = \omega_1 \mathcal{L}_{\text{PDE}} + \omega_2 \mathcal{L}_{\text{data}} + \omega_3 \mathcal{L}_{\text{BC}} + \omega_4 \mathcal{L}_{\text{IC}}, \quad (2.46)$$

with the weighting coefficients  $\omega_{1,\dots,4}$  for the different loss terms related to the residuals of the PDEs  $\mathcal{L}_{\text{PDE}}$ , available measurement data  $\mathcal{L}_{\text{data}}$ , as well as boundary and initial conditions,  $\mathcal{L}_{\text{BC}}$  and  $\mathcal{L}_{\text{IC}}$ , respectively. In vanilla PINNs [176] the neural network is realized by a multi-layer perceptron, however, further work explored the use of CNN and various other neural network architectures for PINNs [41]. The parameters  $\mathbf{W}$  of the neural network and the unknown parameters of the PDEs  $\zeta$  are jointly optimized by gradient-based optimization methods. Thereby, the PDEs are encoded within the network itself and the space of possible solutions is restricted to those that adhere to the underlying governing equations. The introduction of prior knowledge from the PDEs introduces an inductive bias that regularizes the optimization of the neural network. Thereby,

overfitting is mitigated, which facilitates the training of large networks capable of the approximation of complicated non-linear systems, while also allowing for training with limited or even no data.

The composition of the loss function is problem-dependent, as for inverse problems the initial and boundary conditions are typically unknown, while the solution of forward problems does not require any data. However, PINNs can be used to solve the over-determined system, in which initial and boundary conditions, as well as data are all integrated. The contributing loss terms are calculated over a set of sampling points across the spatio-temporal domain, typically with a varying number of samples for each loss term. The physics-informed losses are obtained by means of automatic differentiation (AD) [10], which calculates the explicit expression for the derivatives of  $\hat{\mathbf{u}}(\mathbf{x}, t)$  with respect to the input coordinates by the backward pass through the computational graph of the neural network (see Section 2.4.2). Note, that the procedure is similar to the optimization of the neural network during backpropagation, in which the derivatives of the loss function with respect to the network weights are calculated in order to obtain the weight update. Therefore, AD functions as the key mechanism to combine physics and data, *i.e.* the (physics-uninformed) neural network with the physics-informed network, which can be seen as two networks that share the weights  $\mathbf{W}$ . PINNs are generally considered a meshless method, as all derivatives are calculated by automatic differentiation at randomly sampled point coordinates, and thus no grid generation is required for the solution of the PDEs with PINNs. The calculation of the PDEs residuals by automatic differentiation requires  $n$ -times differentiable activation functions, depending on the highest order of derivatives in the PDEs. Therefore, infinitely differentiable activation functions, such as the hyperbolic tangent and, more recently, the sine activation function are commonly employed in PINNs.

The potential of PINNs for scientific computing in fluid mechanics lies in their flexibility and their capability as highly expressive function approximations. The unified framework of PINNs allows for the seamless integration of data from measurements with the underlying physics represented by governing equations in both forward and inverse problems. Specifically, it has been shown that a few measurement points or noisy data can lead to a significant increase in accuracy already [110]. One promising application of PINNs in fluid mechanics is the direct reconstruction of flow fields from flow visualization techniques [177]. Further examples, particularly for two-phase flows will be elaborated in Chapter 7.

### 3 Development of the reconstruction framework

In this chapter, first, a review of existing methods for the reconstruction of two-phase flows based on optical measurement techniques is conducted in Section 3.1, which is followed by a review of the current state-of-the-art regarding deep learning methods for volumetric reconstruction in Section 3.1. On the basis of this literature research the problem statement for the reconstruction framework, in particular, the solution space and requirements for the integrated measurement technique and reconstruction method are defined in Section 3.3. Consecutively, the reconstruction framework and its components are conceptualized and the selected deep learning reconstruction technique is outlined in Section 3.4. Parts of this chapter have previously been published in the articles *Three-dimensional encoding of a gas-liquid interface by means of color-coded glare points* by Dreisbach *et al.* [DBB<sup>+</sup>23], *Spatio-temporal reconstruction of droplet impingement dynamics by means of color-coded glare points and deep learning* by Dreisbach *et al.* [DKS24], and *Interface reconstruction of adhering droplets for distortion correction using glare points and deep learning* by Dreisbach *et al.* [DHK<sup>+</sup>25]. Where findings have already been published in one of the referenced works, these are clearly marked and identified as related quotations.

#### 3.1 Review of two-phase flow reconstruction

“For the reconstruction of the three-dimensional gas-liquid interface in two-phase flows various measurement techniques based on different optical phenomena have been proposed, including methods based on refraction [156, 172, 47, 48], fluorescence [97, 189], light scattering [119, 72, 17, 46, DBB<sup>+</sup>23] and structured light techniques [91, 263, 95].” from [DKS24, Section 1.1]

**Refraction-based methods** “Morris and Kutulakos [156] used the refractive properties of the gas-liquid interface to reconstruct the three-dimensional surface of a fluid film. The authors captured the deformation of a known reference pattern placed below the fluid surface in a stereo setup and subsequently determined the 3D shape by the comparison to a calibration image. Recently, Qian *et al.* [172] extended the dynamic-refraction stereo-based approach by further constraining the reconstruction to match the local 3D geometry, which allowed for a more precise estimation of the position and orientation of the surface. Dehaeck *et al.* [47] utilized the fringe pattern that is created by the refraction of a liquid droplet imaged from below in a Mach-Zehnder interferometry setup in order to reconstruct the three-dimensional volume by means of a 1D wavelet transform, later expanding their setup for the reconstruction of an asymmetric droplet through 2D fan wavelets [48].” from [DBB<sup>+</sup>23, Section 1]

**Structured light techniques** “Conversely, structured light techniques rely on the deformation of a light pattern projected onto the surface of a liquid by the diffuse reflection for a three-dimensional reconstruction. Zhang *et al.* [263] determined the thickness of a fluid film through a cross-correlation between the deformed image and a calibration image using a digital image

pattern. Hu et al. [95] applied the Fourier transform profilometry [see 218], a method in which a fringe pattern is modulated by the fluid surface, for the volumetric reconstruction of liquid films and the wind-driven droplet flow on an airfoil.” from [DBB<sup>+</sup>23, Section 1]

**Fluorescence-based methods** “Ihrke et al. [97] exploited the emission from fluorescent dye dissolved in the fluid imaged from multiple viewpoints for the volumetric reconstruction of a stream of water. More recently, Roth et al. [189] combined laser induced fluorescence (LIF) [115] and the projection of a fringe pattern to reconstruct the 3D surface of a pendent droplet with a single camera setup.” from [DBB<sup>+</sup>23, Section 1]

**Methods based on specular reflection** “Horbach and Dang [91] used the reflection of a structured light pattern on a specular surface for a 3D reconstruction through a region-growing approach that considers local curvature features. A variety of methods for the monocular reconstruction of fluid surfaces from specularity have been developed in the area of computer vision. Li et al. [131] utilized the shape from shading method [92] for the reconstruction of a wavy water surface from specular reflection, followed by a refinement of the reconstructed surface through physical constraints by a shallow water model. Similarly, [260] estimated the height of a fluid surface with waves through specular reflection and employed the Stokes wave model in order to improve the accuracy of the reconstruction.” from [DBB<sup>+</sup>23, Section 1]

**Methods based on glare points** “The glare points that result from the scattering of light on a liquid droplet can be used for a measurement of its fluid mechanical properties. König et al. [119] used the fringe pattern created by the interference of defocused glare points from reflected and refracted light (produced in wide-angle forward scatter) to measure the diameter of droplets. Later, Glover et al. [72] applied this method to determine the diameter, position, and velocity of droplets in a polydisperse spray. Conversely, Dehaeck et al. [46] used in-focus glare points to measure the size and velocity of bubbles in a fluid flow. The authors encode additional information through an additional reflective glare point produced by a light source at a different illumination angle. This glare point can be used for the detection of non-spherical bubbles and the determination of the relative refractive index between the gas and liquid phases. More recently, Brunel et al. [17] used interferometric particle imaging (IPI) with three perpendicular views in order to reconstruct the 3D shape and orientation of irregularly-shaped rough particles from their speckle patterns.” from [DBB<sup>+</sup>23, Section 1]

**Methods based on shadowgraphy** “The most common approaches build upon multi-view shadowgraphy experiments, which are, however, also complemented by single cameras/view efforts. Tomiyama et al. [226], for instance, determined the integral volume of a bubble from a single shadowgraph image. The authors assumed an oblate spheroidal shape of the bubble and determined its major and minor axis from the shadowgraph contour. Fujiwara et al. [67] reconstructed the 3D shape of deformed bubbles from two orthogonal shadowgraph views by fitting multiple cross-sections along the third orthogonal direction as ellipses whose main axes are estimated from the horizontal extent of the bubble in the two views. Honkanen [90] extended this slicing approach by first determining an oriented 3D bounding box for the bubble, in which then horizontal cross-sections are fitted by ellipses.” from [DKS24, Section 1.1]

“Higashine et al. [85] reconstructed the 3D gas-liquid interface shape of droplets deformed by gravitational or centrifugal forces through an analytical model based on the three-dimensional Laplace equation. While the analytically calculated droplet shape agrees well with experiments,

only the equilibrium state is considered and an accurate measurement of the advancing contact angle is required.” from [DHK<sup>+</sup>25, Section 1]

“Fu and Liu [66] employed the space carving technique [124] to combine the contours of a bubble at four different viewing angles into a virtual hull, which represents a maximum estimate of the bubble volume. Subsequently, this virtual hull is smoothed by spline fits on multiple cross-sections in a slicing approach, in order to represent the effects of surface tension and reach a more accurate reconstruction of the three-dimensional bubble shape. Masuk et al. [143] introduced additional virtual cameras to the space carving technique in order to consider surface tension. In their approach, the reconstructed hull is projected to the novel views of virtual cameras, in which locations of high curvature are iteratively smoothed, while the contour of the gas-liquid interface in all real views is respected. Ríos-López et al. [182] reconstructed the 3D shape of a deformed, non-axisymmetrical droplet sliding on a flat surface from two orthogonal views through a polynomial fit of the contours with the assumption of plane symmetry. More recently, Gong et al. [73] trained a neural network for the volumetric reconstruction of one side of a bubble in form of a depth map from grayscale information of a single shadowgraph image. The authors employed a pyramidal convolutional neural network [127] that was trained on rendered synthetic images of bubbles and their respective ground truth volumetric shapes.” from [DKS24, Section 1.1]

“The above introduced single view methods are in many cases only applicable at a limited degree of surface deformation or curvature, which as yet renders multi-view methods mandatory for any advanced volumetric reconstruction of an impinging droplet.” from [DBB<sup>+</sup>23, Section 1]

## 3.2 Review of deep learning techniques for volumetric reconstruction

“In recent years, deep learning methods for the volumetric reconstruction from images have evolved rapidly and became a promising prospect for the analysis of multi-phase flows. The principal idea is training a neural network for the representation of a deformable 3D geometry on a large paired dataset of input images and output 3D shapes or even just multi-view images. The prior knowledge learned by the neural network resolves ambiguities in the input and thereby allows for the volumetric reconstruction of the 3D geometry from a single image. Different 3D representations have been proposed in contemporary works, namely voxel-based [71, 37, 251, 179], point cloud [62, 132] and mesh-based [241], implicit representations [34, 169, 147, 191] and neural rendering techniques [161, 149].” from [DKS24, Section 1.2]

**Voxel-based representations** “Early work in this field leans on the success of convolutional neural networks (CNN) [127] as the 2D convolution operation can directly be extended to the 3D domain for the prediction of discrete volumetric representations, such as voxel grids. Girdhar et al. [71] trained a joint low dimensional representation for 2D-images and 3D voxel shapes through a 3D convolutional auto-encoder network [114], which allows for the reconstruction of 3D geometries from a single image. Choy et al. [37] employed a recurrent long short-term memory network (LSTM) [88] that consecutively processes multiple images and efficiently merges information of previously unseen parts of the 3D geometry with each novel viewpoint. Wu et al. [251] combined generative adversarial training [75] with 3D convolutional neural networks, which leads to a more realistic shape generation.” from [DKS24, Section 1.2] Furthermore, through the integration of variational autoencoders (VAE) [114] their 3D-VAE-GAN architecture allows for the volumetric reconstruction from a single image.

“While voxel representations can handle arbitrary topology and deliver accurate results, they are not suitable for the reconstruction of fine details, as the computational and memory requirements increase rapidly with the resolution. Furthermore, voxels represent 3D geometry inefficiently, as only those voxels representing the surface contain all the necessary information. Riegler et al. [179] proposed a voxel representation through adaptive hierarchical octrees [146], which have a fine resolution near the surface and a coarse resolution in the remainder of the domain, thus alleviating the drawbacks of voxel representations.” from [DKS24, Section 1.2]

**Mesh and point cloud-based representations** The representation of a 3D geometry as a mesh or point cloud allows for a more compact and scalable encoding of the surface, with low memory and computational requirements. Fan et al. [62] propose to learn a point cloud representation of the 3D geometry for the volumetric reconstruction from monocular images. In that regard, the authors introduce a novel conditional generative network, consisting of a CNN-based image encoder, followed by two branching paths for point-cloud prediction, a deconvolution network that preserves spatial continuity and a multi-layer perceptron (MLP) that predicts accurate fine details. Lin et al. [132] employ a 2D-convolutional encoder-decoder network to predict multi-channel images at different novel viewpoints as an intermediate representation, which encodes the coordinates of a point cloud. Subsequently, these intermediate representations are fused into a dense point cloud of the 3D surface through a transformation into a canonical space. Point-cloud representations are simple to implement and to assess with neural networks since they are unordered and no connectivity needs to be represented, however, they require considerable effort in post-processing for the retrieval of the 3D surface shape. Wang et al. [241] consider a mesh-based representation through a graph-based fully convolutional network, that reconstructs a 3D geometry by deformation of a template mesh in a coarse-to-fine manner. Their mesh-based representation allows for information flow between neighboring vertices during training, which helps to provide a regular smooth output, however, the topology is restricted by the mesh template.

**Implicit representations through continuous functions** “Recent advances have led to the development of implicit representations of continuous 3D shapes using neural networks to approximate level set functions. Chen and Zhang [34] and Mescheder et al. [147] trained an MLP for the implicit field representation of a 3D shape as an occupancy function, that takes the value of one if a point coordinate lies inside of the shape and zero otherwise. In this approach, a binary classification network is trained for the approximation of a decision boundary, which represents the surface. Park et al. [169] proposed the representation by a signed distance function (SDF) that takes values greater than zero on the outside and values smaller than zero on the inside, thus placing the surface at  $\text{SDF} = 0$ . An MLP evaluates this implicit function for randomly sampled 3D coordinates that are concatenated with the global image features extracted from an input image, thus making the method suitable for monocular volumetric reconstruction. Due to the continuous nature of the implicit representation, it can be evaluated at any arbitrary resolution and consequently, a high surface quality can be reached, while the memory requirement is comparably low. The marching cubes algorithm [136] is commonly employed to reconstruct a coherent 3D mesh from the evaluated point coordinates. Since methods based on implicit functions rely on global context for the prediction of 3D shapes, the local alignment with the input image is not guaranteed. Saito et al. [191] proposed the extraction of pixel-aligned local features by first processing the input image through a fully convolutional hourglass network [160] prior to the prediction of a level set function with an MLP. The combination of local features with a global 3D representation by an implicit function allows for the reconstruction of fine details, while the global shape is preserved.” from [DKS24, Section 1.2]



More recent advancements on the basis of the aforementioned pixel-aligned implicit function (PIFu) approach [191] have aimed at a further improvement of the accuracy and resolution of the prediction with the scope of a more detailed volumetric reconstruction. Noteworthy contributions in the context of monocular volumetric reconstruction are briefly introduced in the following. Saito et al. [192] extended the PIFu approach to a multi-scale architecture by the introduction of a second convolutional hourglass network that processes the input images at a higher resolution. The combination of a coarse and a fine resolution pathway during feature extraction enhances the ability of PIFu to utilize both global context and small-scale local information during the prediction, which improves the accuracy and the level of detail of the reconstruction. Similarly, He et al. [81] integrate the PIFu approach with a coarse voxel-representation that is learned by a second neural network in order to resolve ambiguities of the 2D projection and increase the quality of the reconstruction.

**Neural rendering techniques** Recently, differentiable rendering techniques have become increasingly relevant, as they produced impressive results in novel view synthesis and volumetric reconstruction and can be trained without 3D supervision. Niemeyer et al. [161] employ differentiable volumetric rendering and an implicit representation of the 3D geometry for monocular volume and texture reconstruction. Mildenhall et al. [149] introduced neural radiance fields, a continuous volumetric representation of a scene that is encoded inside of a simple MLP. From this representation, novel views from any arbitrary angle can be obtained by volumetric rendering.

### 3.3 Reconstruction framework

The scope of the present work, as outlined in Section 1.2, is the spatio-temporal reconstruction of the 3D gas-liquid interface dynamics of two-phase flows, in particular droplets, from monocular optical experiments. For that purpose, a reconstruction framework has to be conceptualized, which integrates data-driven techniques for spatio-temporal reconstruction with an optical measurement method for the observation of the droplet dynamics in the experiments. Furthermore, the integration of both components requires a data processing pipeline that conditions the experimental data prior to reconstruction and allows for the generation of appropriate training data.

**Problem statement** As discussed in Section 3.1, reconstruction methods based on monocular experiments are generally restricted to a low degree of deformation of the gas-liquid interface. This limitation is caused by several issues related to the projection of three-dimensional bodies into two-dimensional images. In most cases, this projection is ambiguous, as multiple 3D shapes can result in the same 2D projection, which makes the inverse problem of a volumetric reconstruction fundamentally ill-posed. Furthermore, the projection leads to a loss of information, which is particularly apparent in the self-occlusion of significantly deformed gas-liquid interfaces. Consequently, monocular experimental data often only offer an incomplete and sparse representation of the three-dimensional nature of gas-liquid interfaces, which limits the accuracy and area of application of monocular reconstruction techniques. To compensate for the sparsity of information and address the issue of ill-posedness of the reconstruction problem, often model assumptions, such as symmetries or simple geometric shapes of the interface, are employed. This introduction of inductive bias reduces the possible degrees of freedom that arise from the ambiguity of the projection and thereby allows for the reconstruction based on the limited information in the experimental recordings. Here, data-driven methods offer a promising alternative for the reconstruction of more complicated interface geometries, as prior knowledge can be obtained through machine

learning techniques directly from the database, without manual modeling. Thereby, a complex model of the interface dynamics can be acquired. The effectiveness of such data-driven methods, however, is highly dependent on the availability of sufficient amounts of high-quality data that represents the underlying physical principles governing the gas-liquid interface dynamics. In that regard, numerical simulation provides a rich source of information, as the problems of adhering droplets in shear flows and droplet impingement have been extensively investigated in numerical studies. Previous numerical investigations on the impingement of droplets on flat and structured surfaces [64, 250, 194, 228] and the dynamics of adhering droplets subjected to external shear flows [145, 122, 26] provide a large repository of numerical data that is readily available. Specifically, these numerical investigations provide highly-resolved three-dimensional information of the gas-liquid interface, as well as 3D velocity and pressure fields of the flow in both phases, as outlined in Section 2.2.4, which can be used as ground truth data for the optimization of data-driven models. Regarding the experimental observations, which are intended to serve as the basis for the reconstruction, a plethora of optical measurement methods are available that can be employed to produce large amounts of high-quality image data. Given that both experimental recordings for the input side, as well as 3D data for the output side, are easily accessible and further data can be acquired if necessary, supervised learning emerges as the optimal choice, as it provides access to a wide range of reconstruction techniques, which are often easier to optimize than unsupervised learning approaches. Furthermore, supervised learning allows the introduction of domain knowledge from both experiments and simulation, which in combination with a more constrained optimization process, offers the prospect of higher reconstruction accuracy. On the basis of this first elementary design decision, the components of the reconstruction framework are determined in the following, in particular the experimental imaging method and the data-driven reconstruction technique. Subsequently, a suitable data processing pipeline will be derived to integrate these two main components. The fundamental strategy that guides the conception of the reconstruction framework is the combination of simple methods, with the objective of facilitating a straightforward implementation. The underlying incentive for this approach is to create a functional version of the reconstruction framework as early as possible in order to demonstrate the feasibility of the approach at an early stage and to allow for an evaluation of the framework that informs the iterative development of the components that have been identified as weak points in the analysis. In this context, the optimization target of the method development is the accuracy of the gas-liquid interface reconstruction.

**Experimental method** For the observation of the droplet dynamics in the experiments, the shadowgraphy technique appears to be the most suitable option, as it combines a low degree of complexity with a high spatial accuracy and temporal resolution. As elaborated in Section 2.3.1, the shadowgraphy method accurately maps the gas-liquid interface to the images, which provides a substantive foundation for a high reconstruction accuracy. Furthermore, an adequate temporal resolution of the droplet dynamics can be reached in combination with high-speed imaging, which enables the spatio-temporal reconstruction based on snapshot images. The simplicity of the shadowgraphy setup facilitates an easy implementation and operation of the experiments and, furthermore, provides flexibility for the adaptation in different fluid mechanical investigations. Consequently, the shadowgraphy technique is the most commonly employed optical method for the observation and volumetric reconstruction of two-phase flows, which brings about a wealth of knowledge and experience documented in the current literature. Its wide range of applications covers both droplet impingement and adhering droplet experiments. From previous experimental investigations at the Institute of Fluid Mechanics, an experimental apparatus for the investigation of droplet impingement along with data and expertise on this experimental method are available. Additionally, Burgmann et al. [25] employed shadowgraphy for the investigation of adhering



droplets in shear flow, which likewise provides experimental hardware and expert knowledge for this thesis. On the basis of the aforementioned arguments, the shadowgraphy technique is chosen as the base of the reconstruction framework for data collection in the experiments.

A drawback of this method, however, is the limited content of information in the images, which is primarily located on the contour of the shadowgraph and, most importantly, represents only a 2D projection of the 3D gas-liquid interface. In that regard, techniques based on the specular reflection of light on the gas-liquid interface present a promising complementation for encoding additional information on the 3D interface shape within the images. As elaborated in Section 3.1, a simple setup with lateral light sources at different illumination angles can be used to produce glare points that result from the scattering of light at the gas-liquid interface. Therefore, the canonical shadowgraphy setup is extended by glare points from two lateral light sources that are anticipated to enrich the shadowgraph images with additional three-dimensional information on the gas-liquid interface. In order to facilitate the differentiability of the information, differently colored light sources are employed, specifically a blue light source for the backlight illumination, in combination with red and green light sources for the glare points. In this setup, the individual glare points from each light source, as well as the shadowgraph contour, can be captured in the respective channels of an RGB image. The scope of the proposed method is to provide a sufficient amount of 3D information on the gas-liquid interface to enable a volumetric reconstruction of the interface while retaining the simplicity of a monocular measurement setup. While the increase in complexity from the additional lateral light sources appears justified on a theoretical basis, the effectiveness of the method has to be proven in practice by its impact on the accuracy of the reconstruction. In order to investigate this question, both the canonical shadowgraphy setup and the extended glare-point shadowgraphy setup will be evaluated in a comparative experimental study, in which both methods are employed as data sources for the reconstruction. Furthermore, it should be noted that the combination of additional 3D information on the interface shape by the glare points and the utilization of data-driven modeling for the volumetric reconstruction of the interface can potentially result in redundancy. To which degree and under which circumstances the information in the glare points provides a valuable addition for the reconstruction, will be a subject matter of the following investigation.

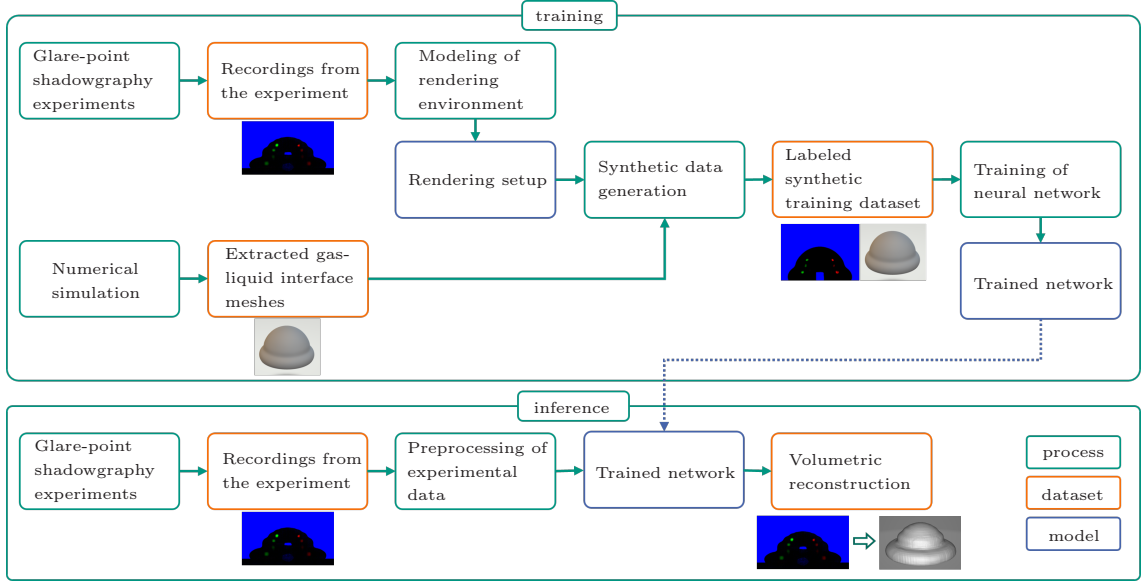
**Data-driven reconstruction technique** The core component of the reconstruction framework is a data-driven reconstruction technique that predicts the 3D shape of the gas-liquid interface from monocular images recorded in the experiments. On the basis of the literature research on deep learning techniques for volumetric reconstruction techniques presented in Section 3.2, methods based on the implicit representation of 3D shapes emerge as the most suitable option, as this approach combines a computationally efficient representation of the interface with a high reconstruction accuracy. In comparison to voxel-based techniques, implicit representations facilitate a higher reconstruction accuracy, while simultaneously demanding fewer computational and memory resources. In contrast to mesh- and point cloud-based methods, implicit representations are flexible towards the shape and topology of the 3D reconstruction target and, therefore, allow for the representation of topological changes of the interface, such as break-up or coalesce. Techniques based on neural rendering are primarily oriented towards unsupervised learning tasks but might be a further interesting candidate method in the future. Although there are extensions for supervised learning, the adaptation of such methods appears less straightforward in comparison to methods based on implicit representations. The PIFu approach [191] appears to be the optimal choice among the methods based on implicit representations, offering both a high reconstruction accuracy and the prospects of a straightforward implementation into the framework. As outlined in Section 3.2, PIFu incorporates developments of extensive research on monocular volumetric reconstruction by means

of deep learning, which provides an extensive theoretical background and results in a high technical maturity. As a result, PIFu features a comparably high spatial resolution of the predicted three-dimensional shapes that allows for the reconstruction of fine details, while also ensuring a high accuracy on the global scale. Concurrently, its purposefully developed neural network architecture is computationally efficient, while retaining a manageable degree of complexity. Furthermore, high-quality open-source code is available, which facilitates the implementation of this reconstruction technique and its integration into the reconstruction framework. Consequently, the choice of PIFu enables an early evaluation of the feasibility of the conceived reconstruction framework and, thereby, serves as the baseline for the further development of data-driven reconstruction techniques in this work.

As elaborated in Section 3.2, further research on the extension of the PIFu approach has led to improvements in the accuracy and resolution of this reconstruction technique. However, these approaches introduce additional complexity to the neural network architecture and training algorithms. Considering that the capacity of the neural network has to be adequate with respect to the complexity of the model that is to be trained, it appears more reasonable to initially implement a method with lower complexity and potentially lower capacity, providing the pathway to evaluate the performance of the framework and, subsequently, increase the complexity if required. Furthermore, neural networks with a higher capacity in general require larger datasets for the optimization of a higher number of parameters, which increases training time and raises the risk of overfitting. On the basis of these considerations, the baseline PIFu approach offers a well-rounded trade-off between performance and simplicity and is, therefore, chosen as the reconstruction technique for the initial concept of the reconstruction framework.

While supervised learning facilitates the incorporation of prior knowledge from both experimental and numerical investigations, the governing equations of the gas-liquid interface dynamics are only implicitly considered through the numerical training data. A more direct incorporation of the physical laws during reconstruction can be achieved by physics-informed neural networks (PINNs), which encode the governing equations, such as the Navier-Stokes equations in the neural network by training on a physics-informed loss, as introduced in Section 2.4.4. Thereby, PINNs regularize the optimization of the neural network and force the solution to adhere to the governing equations, which is expected to further enhance the reconstruction quality. However, this approach requires a significant modification of the neural network architecture and the training algorithms and likely will increase the difficulty of the optimization process, as PINNs require a fine tuning of the physics-informed and data loss terms. Consequently, the implementation of PINNs requires substantial development efforts to first determine the feasibility of the approach and, then, optimal training hyperparameters. Therefore, first a purely data-driven reconstruction based on the PIFu approach will be employed to prove the feasibility of the proposed reconstruction framework and based on the insights gained during the evaluation of this approach, the integration of PINNs into the reconstruction framework will be developed in a second step.

**Overview of the reconstruction framework** The training of the neural network on both data from experiments and numerical investigations in a supervised learning approach requires a labeled dataset, *i.e.* pairs of matching input images and respective 3D output shapes. For both considered fluid mechanical problems, specifically droplet impingement on flat and structured substrates and adhering droplets subjected to external shear flows, numerical datasets are available and experiments can be conducted under matching conditions to obtain a labeled dataset. “However, due to uncertainty in the experiments, and errors from modeling and approximations in the simulation, the droplet shapes are not identical. This results in a mismatch of the input images with their respective ground truth labels, which consequently introduces an error to the neural



**Figure 3.1:** Overview of the reconstruction framework, detailing the processing steps and associated datasets during training (top) and inference (bottom).

network. This kind of matching problem is resolved by training the neural network on fully synthetic image data that is obtained using the ground truth of the simulation, thus resulting in a perfect agreement between input and output in the dataset. The synthetic images are generated by means of a render-pipeline in *Blender* [39] with the *LuxCore* [19] package that allows for physically correct ray tracing. The virtual rendering setup replicates the optical configuration in the experiments in order to produce realistic image data. The PIFu neural network [191] is trained for the spatio-temporal reconstruction of the droplet dynamics [...] on this synthetic dataset.” from [DKS24, Section 2]


Subsequently, the trained model of the droplet dynamics is employed for the reconstruction of real observations in the experiments. Consequently, two distinct data processing pipelines are required during the training and inference of the neural network. An overview of the reconstruction framework that integrates the previously defined components for data generation, imaging, and reconstruction along with their respective data streams and processing steps during training and inference is illustrated in Figure 3.1.

The starting point for the training routine are experiments within the setup of the extended glare-point shadowgraphy technique in order to obtain a dataset of recordings that is representative of the optical phenomena involved in the image formation in the experiments. This dataset, in conjunction with theoretical considerations regarding image formation, is used to model the rendering setup, with the aim of generating synthetic images that accurately reproduce the visual appearance of the recordings from the experiment. The modeled rendering setup is then employed for the generation of synthetic images on the basis of 3D meshes of the gas-liquid interface extracted from the results of numerical simulations. The resulting labeled synthetic dataset is used for the optimization of the PIFu neural network in order to obtain a trained model of the droplet dynamics. This model of the droplet dynamics encoded in the trained neural network is employed for the spatio-temporal reconstruction of the droplet’s gas-liquid interface in novel experimental investigations. The starting point of the inference routine are experiments involving droplets impinging on solid substrates or adhering droplets subjected to external shear flow, which are imaged by the glare-point shadowgraphy technique. The image data recorded in the experiments has to be pre-processed prior to the actual reconstruction, to attain adequate conditioning as input for the PIFu neural

network. Finally, for each snapshot in the experiments, the volumetric shape of the gas-liquid interface is reconstructed by the trained network to recover the spatio-temporal dynamics of the droplet's interface.

Confronting the neural network with two distinct types of input data – synthetic images during training and experimental recordings during inference – necessitates a dedicated and careful pre-processing to ensure the consistency of the input data during training and inference, and thereby a consistent treatment by the neural network. The developed data processing techniques will be elaborated in further detail in the respective chapters regarding glare-point shadowgraphy and synthetic image data generation. In the following chapters, the components of the reconstruction framework will be discussed in further detail, in an order that considers the processing sequence and resulting interrelations in the reconstruction framework. First, the neural network architecture and optimization procedure of PIFu will be covered, followed by the discussion of the experimental setup in Chapter 4, and synthetic data acquisition techniques in Chapter 5. Afterwards, the methodology for the pre-processing of the recordings obtained in the experiments will be outlined in Section 6.1.3, followed by the validation and evaluation of the reconstruction framework in Sections 6.2 and 6.3, respectively. The development and integration of PINNs into the reconstruction framework will be discussed in Chapter 7.

### 3.4 Neural network architecture and optimization procedure



**Figure 3.2:** Sketch of the PIFu neural network architecture. The input images are first processed by a convolutional hourglass network, which extracts image features  $\mathcal{F}(x, y)$  that are subsequently sampled locally by a second network (MLP), which predicts the occupancy function at multiple points in 3D space. Figure adopted from [DKS24].

“For the volumetric reconstruction of the gas-liquid interface the state-of-the-art deep learning method based on the concept of the Pixel-aligned Implicit Function (PIFu) [191] is used as outlined in Figure 3.2. The core concept of PIFu is the implicit representation of the three-dimensional topology as a level-set function through a neural network. In this approach, an MLP predicts the three-dimensional occupancy field for various locations in the image plane by probing a defined set of distances in the out-of-plane coordinate for the prediction of whether the point coordinate lies within or outside of the predicted geometry. The input to the MLP are image features extracted by a convolutional neural network (CNN), in particular, the so-called hourglass network [160], which subsequently down- and upsamples the input image. The resulting feature maps are therefore pixel-aligned, *i.e.* each pixel position of the extracted feature maps has an accordingly aligned area in the input image. Furthermore, the relation between distant locations in the image can be expressed by the pixel-aligned feature maps, thus enabling the neural network for both global and local reasoning.” from [DKS24, Section 2.3]

**Network architecture** Specifically, the hourglass network consists of four stacked encoder-decoder pairs, the so-called hourglass modules, one of which is shown as a representative example of the network in Figure 3.2. In the encoder part of each hourglass module, features are extracted and a feature pyramid with increasing semantic information in each consecutive convolutional layer is built, while the resolution of the feature maps decreases. In the decoder part, the spatial resolution is increased again through upsampling operations, while feature maps with rich semantic context are fused with high-resolution feature maps via residual connections (see Section 2.4.3). Thereby, information on different scales is taken into account during feature extraction. In this manner, the hourglass network preserves the spatial structure of the input image while enriching the feature maps with the information from neighboring pixels. The repeated up- and downsampling through the multiple stacked hourglass modules in the network further enhances the propagation of the information in the feature maps and, thereby, introduces global context between potentially distant regions in the output. The flow of information, as well as the stability of the network optimization, are improved by residual connections between the hourglass modules, as indicated by the arrow in Figure 3.2. During training, intermediate predictions of the occupancy function are performed by the MLP based on the final feature maps in each hourglass module, which allows for the calculation of a loss term for each level of the hourglass network. Thereby, intermediate supervision is provided at each hourglass module, which further stabilizes the optimization of this deep neural network architecture [160]. In contrast to the prediction on the basis of global features [34, 147, 169], the hourglass network extracts high-resolution features that are spatially aligned with the input image, which facilitates the accurate reconstruction of local details. In order to allow for global reasoning, the receptive field of the output neurons in the hourglass network has to cover the whole input image in order to enable the propagation of information from each pixel in the input image into each location of the last feature map [192]. As introduced in Section 2.4.3, the receptive field increases with the number of consecutive layers. The overall size of the network, however, is restricted by memory limitations, which creates a trade-off between the resolution of local features and the accuracy of the global shape estimation. Therefore, the size of the input images is limited to 512 x 512px and the size of the final feature map in each hourglass network is limited to 128 x 128px in 256 channels. It should be noted that extensions of the PIFu approach in more recent works [192, 81] facilitate the reconstruction of smaller details through a higher resolution of the extracted features and the combination of multiple pathways at different resolutions.

On the basis of the extracted features an MLP learns the implicit representation of the 3D occupancy field in the form of a level-set function, as illustrated in Figure 3.2. The 3D reconstruction on the basis of pixel-aligned features  $\mathcal{F}(x, y)$  in the image plane, requires additional parametric input for the depth coordinate  $z$ . As all 3D positions in the projection along a ray in the depth direction have the same image features, the PIFu network focuses on the varying input depth  $z$  to infer the occupancy along the ray. Consequently, for a prediction at the spatial point  $(x, y, z)$ , the pixel-aligned features  $\mathcal{F}(x, y)$  are sampled at the location  $(x, y)$ , while the coordinate  $z$  is directly inputted to the MLP. The MLP consists of four fully-connected layers with 1024, 512, 256, and 128 neurons and skip connections at each layer to propagate the image features  $\mathcal{F}(x, y)$  and depth information  $z$ . The input layer of the MLP has a size of 257 to match the size of the feature vector (256 channels) with the addition of the depth coordinate, while the output has a size of one. In the hidden layers of the MLP, the leaky ReLU activation function is used, while the output neuron features the sigmoid activation function in order to constrain the output between zero and one, ensuring a meaningful prediction of the occupancy. In the hourglass network, the ReLU activation function is employed. Further implementation details can be found in the published work by [191] and the code repository of this thesis.

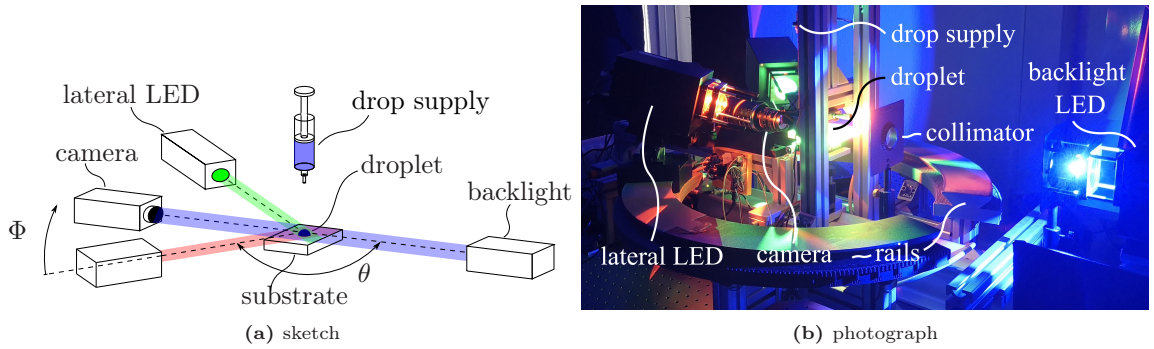
**Network training** “The coupled neural networks were jointly trained on the labeled synthetic dataset by supervised learning with the RMSProp optimization algorithm [225], which is an extension of stochastic gradient descent with momentum (SGDM) [171]. In particular, a forward pass through the network consists of the computation of feature maps from the input images by the hourglass network, followed by  $n$  passes through the MLP to predict the occupancy at  $n$  spatial locations  $\mathbf{x} = (x, y, z)$  in continuous 3D space. As illustrated in Figure 3.2, the MLP receives the pixel-aligned local features at the position  $(x, y)$  in the image plane and an additional parametric input for the depth coordinate  $z$ , which are used to predict the occupancy  $\hat{\phi}(\mathbf{x}_i)$  at sampling point  $i$ , for  $i = 0 \dots n$ . In order to allow for continuous sampling, the local features are obtained by bilinear interpolation from the extracted feature map. During training, a purposefully developed sampling scheme introduced by Saito et al. [191], is employed to sample spatial coordinates  $(x, y, z)$  close to the interface, for which the label  $\phi(\mathbf{x}_i)$  of the occupancy function is determined by whether they lie within or outside of the ground truth mesh. This sampling scheme has been found to significantly enhance the learning of 3D shapes, as the optimization of the network is effectively focused on the accurate prediction of the surface location [191]. In the backward pass, the mean squared error loss for all  $n$  sampling points  $\mathcal{L} = \frac{1}{n} \sum_{i=1}^n (\phi(\mathbf{x}_i) - \hat{\phi}(\mathbf{x}_i))^2$  is calculated and backpropagated first through the MLP and subsequently through the hourglass network for one iteration of weight updates. During inference, the coordinates are sampled uniformly from an equidistant grid, as the location of the gas-liquid interface is unknown. The training hyperparameters were originally chosen according to the original publication [191]. However, it was found that a reduction of the training iterations led to better results for the considered dataset. Therefore, the networks were trained for eight epochs [...], with a batch size of 12 and an initial learning rate of 0.001 that is reduced by a factor of ten after epochs four and six. As a technical note, the training duration [for a representative dataset size of 37,300 training samples] amounted to 58 h on a single Nvidia RTX A5000 graphics processing unit. The network requires 20 s on average per time step for the volumetric reconstruction at an output resolution of  $512^3$  grid nodes. Due to the implicit representation of the surface, the reconstruction can be performed at an arbitrary resolution, and consequently, a speed-up is possible by lower output resolutions [191]. Furthermore, processing the predicted grid nodes by the marching cubes algorithms [136] is required in order to obtain a mesh of the gas-liquid interface and PIFu uses the octree structure [146] for a more efficient and thus faster inference. The inference time is required proportionally for the following processes on average, *i.e.* 13.2 s for the prediction by the neural network, 1.6 s for marching cubes, and 5.2 s for data handling.” from [DKS24, Section 2.3]



## 4 Glare-point shadowgraphy

In this chapter, a novel optical measurement technique that extends the shadowgraphy method by color-coded glare points is introduced by the example of droplet impingement on flat solid substrates. While the canonical shadowgraphy captures the contour of the gas-liquid interface with high spatial accuracy, the resulting images only provide sparse information limited to a two-dimensional projection of the three-dimensional interface. The glare points created by the reflection of light on the gas-liquid interface from additional lateral light sources are anticipated to encode further information on the three-dimensional droplet shape in the images, which lay the foundation for the volumetric reconstruction of the dynamically deformed interface during impingement. Within this scope, the theoretical aspects of encoding three-dimensional information on the gas-liquid interface of non-axisymmetrical droplet shapes by means of glare points are examined. The results discussed in this chapter have previously been published in the article *Three-dimensional encoding of a gas-liquid interface by means of color-coded glare points* by Dreisbach *et al.* [DBB<sup>+</sup>23].

### 4.1 Experimental setup



**Figure 4.1:** Sketch (a) and photograph (b) of the experimental setup of the applied glare-point shadowgraphy method. Figure adopted from [DBB<sup>+</sup>23].

“The present approach extends the shadowgraphy method by supplementing two lateral light sources to the standard shadowgraphy setup as shown in Figure 4.1. These additional light sources are mounted on a custom-built rail system as can be seen in Figure 4.1b in order to allow for the adjustment of the azimuth angle  $\theta$ , while the elevation angle  $\Phi$  can be adjusted by a hinge. This mounting system facilitates the independent and continuously variable adjustment of both angles and, thereby, allows for the evaluation of the influence of the direction of illumination on the glare-point location and intensity on the gas-liquid interface. The angular adjustment uncertainty is estimated to be lower than  $\pm 2^\circ$ . In order to ensure a reproducible droplet volume over multiple experiments and to avoid oscillations of the droplet from the introduction of momentum, the water droplet is produced by an automatic drop-application system. This system consists of a syringe with a cannula diameter of  $d_s = 0.1$  mm and a stepper motor with  $l_m = 10$  mm linear displacement

and a step angle of  $\alpha_m = 7.5^\circ$ , which is used as a linear actuator for the syringe. The resulting droplet diameter for distilled water is  $2a \approx 2\text{ mm}$  at room temperature. [...]

The two lateral light sources, as well as the backlight, are high-power *ILA\_5150 LPSv3* LEDs with narrow-banded spectra and maxima in the visible spectrum at  $\sim 455\text{ nm}$  ("blue"),  $\sim 521\text{ nm}$  ("green") and  $\sim 632\text{ nm}$  ("red") for the desired blue, green and red light sources, respectively. The light of the LEDs is captured either at 3,000 frames per second (fps) by a *Photron Fastcam Mini WX* RGB Camera or at 7,500 fps by a *Photron Nova R2*, both equipped with a *Schneider-Kreuznach Apo-Componon 4.0/60* enlarging lens, where the three-colors  $i = (\text{"red"}, \text{"green"}, \text{"blue"})$  of illumination are expected to be mainly captured with the corresponding channels of the RGB camera chip.

The backlight produces a classical shadowgraphy image on the corresponding image channel, where only non-deflected light reaches the camera chip. In order to capture the shape of the droplet as accurately as possible a parallel light beam is required, leading to a better resolution of small features and optimizing contrast [201]. This is achieved by placing an optical collimator consisting of a pinhole aperture with a diameter of  $d_a = 4\text{ mm}$  and a *Spindler & Hoyer* biconvex collimator lens with 300 mm focal length between the backlight and the droplet. The blue LED is used for the backlight since the camera chip possesses the lowest relative sensitivity in this range of wavelengths and the highest response in the chip is expected from direct illumination of the backlight LED.

The lateral light sources are focused by *Thorlabs* plano-convex lenses with a broadband anti-reflective coating (reflectivity  $< 0.5\%$ ) and illuminate the droplet from an azimuthal angle  $\theta > 90^\circ$  in respect to the blue LED. Consequently, only reflected light reaches the camera, resulting in the creation of  $p_{gp} = 0$  glare points [233]. The illuminating light beams possess a cross-section that is larger than the droplet and are approximately parallel, which implies that the beams can be represented as plane waves. Since the light is partially transmitted at the first interface between air and water, it is reflected again on the following interfaces in the path of the light ray, *e.g.* while leaving the droplet or at internal phase boundaries of air inclusion within the droplet, resulting in higher-order glare points ( $p_{gp} \geq 2$ ).


Due to the smooth surface of the phase boundary, pure interface reflection is considered, which is captured by the camera as differently colored glare points on each side of the droplet and on any present air bubbles, thereby encoding additional volumetric information on the droplet in the remaining red and green image channels. Since the information is separated between the channels of the RGB image, three individual grayscale images with accordingly different illumination conditions can be extracted from the RGB image, each providing a unique representation of the droplet that can be used for the volumetric reconstruction of the droplet.” from [DBB<sup>+</sup>23, Section 2.2]

## 4.2 Image processing

“The spectral power distributions of the light sources do not perfectly match with the spectral sensitivity of the corresponding camera channels  $k = (\text{R}, \text{G}, \text{B})$  as illustrated in Figure 4.2. Therefore, each light source causes a perturbation of the images in the other channels, which consequently needs to be considered as follows.

The response of the  $k^{\text{th}}$  color channel at pixel  $x$  on the camera chip  $I_k(x)$  can be calculated by spectral integration from the spectral power distribution (SPD) of the incoming light  $L(\lambda)$ , the





**Figure 4.2:** Spectral power distribution of the LED light sources  $L_i$  (color-coded solid lines); spectral sensitivity distribution of the camera's color channels  $S_k$  (color-coded dashed lines); transmittance of the objective lens  $T_o$  (dashed black line) and the collimator lenses of the light sources  $T_l$  (solid black line). Figure adapted from [DBB<sup>+</sup>23].

spectral transmittance of the color filter  $T_k(\lambda)$  for the  $k^{\text{th}}$  color channel and the spectral responsivity  $s(\lambda)$  of the camera [202] as

$$I_k(x) = \int_{\lambda} L(\lambda) T_k(\lambda) s(\lambda) d\lambda. \quad (4.1)$$

Equation (4.1) implies the assumption of a linear camera-chip behavior, *i.e.* the measured signal at a pixel's location  $x$  is proportional to the intensity of the incident light at this position. This assumption generally holds true, however, with little error for CMOS cameras [239]. The spectral transmittance  $T_k(\lambda)$  and the spectral responsivity of the camera  $s(\lambda)$  are often expressed in a combined form as the spectral sensitivity of the camera  $S_k(\lambda) = T_k(\lambda)s(\lambda)$  for the  $k^{\text{th}}$  color channel. The term of the camera's spectral sensitivity  $S_k(\lambda)$  also includes the wavelength-dependent transmittance of the objective lens  $T_o(\lambda)$ , while the transmittance of the focusing lenses of the light sources can be assumed to be constant in the relevant wavelength-band, as seen in Figure 4.2. Since spectral integration is a linear transformation, an integration over  $L(\lambda)$ , which is the combined superposed spectral power distribution of all light sources  $i$  via  $L(\lambda) = \sum_{i=1}^n \int_{\lambda} L_i(\lambda)$ , results in the same camera response as the sum of an integration over all light sources  $L_i(\lambda)$  [117]. Therefore, Equation (4.1) can be rewritten as

$$I_k(x) = \sum_{i=1}^n \int_{\lambda} L_i(\lambda) S_k(\lambda) d\lambda. \quad (4.2)$$

Considering that light transport from the two lateral light sources to the camera chip can occur either directly or by scattering on the droplet, thus creating glare points, the light transport coefficient  $l(\lambda, x)$  at pixel position  $x$  on the camera sensor is introduced as

$$I_k(x) = \sum_{i=1}^n \int_{\lambda} l_i(\lambda, x) L_i(\lambda) S_k(\lambda) d\lambda. \quad (4.3)$$

In order to consider the contributions of direct illumination, reflection at the droplet interface, and transmission through the droplet, the light transport coefficient is expressed as

$$l(\lambda, x) = d_1(x) + r_1(\lambda, x) + t_1(\lambda, x), \quad (4.4)$$

where  $d_1(x)$  represents the direct light transport,  $r_1(\lambda, x)$  the spectral reflectance and  $t_1(\lambda, x)$  the transmission. Generally, spectral reflectance and transmission are dependent on direction and

wavelength of the light incident on the surface and the geometrical shape of the illuminated object, whereas the participation from direct light only depends on the illumination angle. Furthermore, the spectral reflectance can be divided into interface (specular) reflection  $r_s(\lambda, x)$  and body (diffuse) reflection  $r_d(\lambda, x)$ , according to the dichromatic reflection model [203], *i.e.*

$$r_l(\lambda, x) = r_s(\lambda, x) + r_d(\lambda, x). \quad (4.5)$$

For a water droplet, the diffuse reflection  $r_d(\lambda, x)$  is neglected due to the smoothness of the surface, so that reflection can be assumed to be only specular reflection  $r_s(\lambda, x)$  in the following.

The light transport by transmission  $t_l(\lambda, x)$  is determined by the two-fold refraction at the droplet interface  $R_f(x, \lambda)$  and the spectral transmittance  $T_D(x, \lambda)$  of the droplet liquid. Following the Beer-Lambert law [216], the absorption of distilled water in the visible spectrum is negligibly low for the considered distance traveled within the droplet [77], even considering longer distances for higher-order glare points. Therefore, the spectral transmittance can be considered to be  $T_D \approx 1$  in the range of the considered wavelengths ( $\lambda \approx 400 - 700$  nm) and its influence on the light transport through transmission can be neglected.

Interface reflection and refraction are in general dependent on the refractive index  $n_r$  following the Fresnel equations [83]. Since  $n_r$  is furthermore dependent on the wavelength of the incident light, *i.e.*  $n_r = n_r(\lambda)$ , the intensity of both reflected and transmitted light on the interface likewise depend on the wavelength. Note that the light of the LEDs is unpolarised. Therefore, in the following, the effects of polarisation on the reflection are not considered.

The neutral interface reflection (NIR) model is a good approximation for a water droplet. It states that the spectral power distributions (SPD) of the reflected and incident light are approximately equal [128]. This assumption is valid for materials with insignificant variation of refractive index in the visible spectrum, which holds true for water. With the assumption of dominating specular reflectance and the NIR model, the spectral reflectance  $r_l$  becomes independent from the wavelength, *i.e.*  $r_l(x) = r_s(x)$ . Furthermore, for  $n_r \approx \text{const.}$  the refraction becomes independent of the wavelength  $r_f(x)$ . Consequently, the light transport only depends on the geometry of the problem  $l(x)$  and Equation (4.3) can be accordingly rewritten as

$$I_k(x) = \sum_{i=1}^n l_i(x) \int_{\lambda} L_i(\lambda) S_k(\lambda) d\lambda. \quad (4.6)$$

Since both the SPD of the light sources and the spectral sensitivity of the camera are known, an integration across all relevant wavelengths for the camera sensor results in a transfer coefficient

$$t_{k,i} = \int_{\lambda} L_i(\lambda) S_k(\lambda) d\lambda \quad (4.7)$$

to quantify the transfer of light from each light source  $L_i$  to each channel of the camera chip  $I_k$ .

For the presented measurement setup, consisting of three LEDs with different maximal wavelengths in the red, green, and blue spectral range of the visible light ( $L_r, L_g, L_b$ ), and a camera chip with a red, a green, and a blue channel ( $I_R, I_G, I_B$ ), the resulting transfer function is represented by

$$\begin{pmatrix} I_R \\ I_G \\ I_B \end{pmatrix} = \begin{pmatrix} t_{Rr} & t_{Rg} & t_{Rb} \\ t_{Gr} & t_{Gg} & t_{Gb} \\ t_{Br} & t_{Bg} & t_{Bb} \end{pmatrix} * \begin{pmatrix} l_r \\ l_g \\ l_b \end{pmatrix} = \mathcal{T} * \mathbf{1}. \quad (4.8)$$

Finally, the inverse of the transfer matrix  $\mathcal{T}$  provides a correction matrix  $\mathcal{C} = \mathcal{T}^{-1}$  for a determination of the image  $\mathbf{l}$  created as a result of the light transport from only a singular light source

$$\mathbf{l} = \mathcal{T}^{-1} * \mathbf{I} = \mathcal{C} * \mathbf{I} \quad (4.9)$$

from the response of the three image channels  $\mathbf{I}$ . Therefore, the analytical spectral correction (ASC) function (4.9) allows for a separation of the images created by each only one of the light sources, resulting in three independent images of the droplet.

The transfer coefficients in Equation (4.8) and thus the correction matrix can also be obtained in-situ from the experimental apparatus by means of calibration images, in the following called in-situ color correction (ISC). ISC is particularly useful if the spectral characteristics of the optical components are unknown and, consequently, ASC cannot be applied. For this purpose, three sets of calibration images have to be recorded in the presented setup through an RGB camera, where consecutively only a single light source  $L_i$  is active per recording. Afterward, the channels  $I_k$  of the RGB images are split from each other to achieve nine grayscale images that contain the required information to calculate the transfer coefficients  $t_{k,i}$  of Equation (4.8). An accurate estimate for the transfer coefficients can be obtained by calculating the maximum channel intensities and averaging over a large set of calibration images, thus reducing the influence of random noise. While ISC proved to be effective, additional measurement errors have to be considered in comparison to the above-introduced analytical method.” from [DBB<sup>+</sup>23, Section 2.3]


## 4.3 Validation

In the following section, the proposed measurement technique is validated, starting with an evaluation of the effectiveness of the color-correction function for the removal of cross-talk in the recordings. Subsequently, the usability and accuracy of the Lorenz-Mie theory (LMT) for the prediction of glare points of various orders at different scattering angles in the experiments are examined. Lastly, the influence of the elevation angle of the lateral light sources on the visibility of glare points on the dynamic gas-liquid interface during droplet impingement is investigated.


### 4.3.1 Spectral reconstruction of image data

“The images of a water droplet impinging onto a flat substrate recorded for the present work with the above-introduced RGB shadowgraphy setup and their processing are discussed in the following section. The required information for the ASC function (4.9) is calculated from the SPD of the LED lights and the spectral sensitivity of the camera with the values provided by the manufacturers; see Figure 4.2.

Figure 4.3 shows an RGB frame of the droplet shortly after impact and the separated response of the uncorrected channels of the raw image. The comparison of the three separate image channels to the RGB image demonstrates that the additional information from reflective images of lateral light sources of different wavelength bands can be separated into the channels of an RGB image, which however contain some spurious crosstalk between the images. Saliiently, the background of the green channel is considerably illuminated by the blue light source and the reflective images from the lateral light are apparent on the other channels, but with a lower intensity. The annotations in Figures 4.3b, 4.3c and 4.3d highlight the most noticeable perturbations in the respective color of the causative light source.



**Figure 4.3:** Raw color-imaging of an impinging water droplet on a substrate with SiOx-coating; (a) RGB snap shot; separated (b) red, (c) green, and (d) blue raw channels of (a). The color-coded annotations highlight the perturbations from the respective other channels. Figure adopted from [DBB<sup>+</sup>23].




**Figure 4.4:** Processed color-imaging of an impinging water droplet on a substrate with SiOx-coating; (a) RGB snap shot; separated (b) red, (c) green, and (d) blue corrected channels of (a). Figure adopted from [DBB<sup>+</sup>23].


To reduce the cross-talk effect, the above-discussed ASC function (4.9) has been applied to the RGB image of Figure 4.3. The results are shown in Figure 4.4 below the raw images to allow an immediate comparison of all spurious patches. Obviously, most disturbances in Figure 4.3 have been successfully removed or at least reduced significantly. This is most saliently emphasized by the removal of the background light in the green channel and the transmitted light around the vertical center line resulting from the blue light source. The glare points from the red light source on the right side of the droplet in the green image are reduced as well. In the blue and red channels, the glare points from the green light source on the left side of the droplet are likewise removed successfully. The full series of RGB frames shortly before and after the drop impact, corrected with the ISC function, is provided in Appendix A.2 to further contrast the raw images (Figure A.3) to the corrected images (Figure A.4).

A quantitative accuracy evaluation of the image color correction has been conducted on images from the RGB imaging with only one activated light source. In this setup, a complete color correction is expected to result in zero intensities for all but the considered image channels and only the channel corresponding to the active light source is expected to contain an image. The quality of the color correction is measured by the mean image intensity in the omitted image channels, *i.e.* the remaining disturbance after correction. Figure 4.5 shows the temporal evolution of mean image intensities for all channels in the first 50 images of the drop impact before (dashed lines) and after (solid lines) color correction.

As expected, each light source causes the highest response in the corresponding image channel and additionally a smaller response in each of the two other channels. The intensity of the disturbance in the omitted channels is significantly reduced by the color correction, as becomes obvious from the reduction of average intensity for the respective channel to near-zero values for most frames. This evaluation, therefore, confirms the accuracy of the applied analytical color-correction approach. While the correction of the blue light is nearly perfect, the disturbance in the other channels cannot be fully eliminated for particular frames recorded with only the lateral green or red light source. The corresponding frames are characterized by high overexposure, such that clipping effects lead to an effective overestimation of the other channels. One such event is shown in Figure 4.6, which occurred in the 38<sup>th</sup> frame of an RGB shadowgraphy experiment with only the green lateral



**Figure 4.5:** Mean image intensities of the unprocessed color channels (dashed lines) in comparison to the ASC-corrected color images (solid lines) for the first 50 frames of the drop-impact measurements with a single lateral red (a), green (b) light source, or blue backlight (c). Figure adopted from [DBB<sup>+</sup>23].



**Figure 4.6:** Processed color-imaging with overexposure (frame 38), produced with only the green lateral light source activated; RGB snap shot (a); separated red (b), green (c), and blue (d) corrected channels of (a). Figure adopted from [DBB<sup>+</sup>23].


light source activated (*cp.* Figure 4.5b, where the considered frame is marked with a black arrow).” from [DBB<sup>+</sup>23, Section 3.1]

### 4.3.2 Glare point prediction accuracy


“To evaluate the accuracy of the recorded glare-point locations and intensities, the occurrence of the glare points in the experimental images is compared to the theoretically expected results as can be calculated from the Lorenz-Mie theory for scattering on a spherical particle with Equation (2.36). In the following, the glare-point position is defined by the location of the maximum glare-point intensity. First, the position of the glare points on the projected gas-liquid interface of the droplet, as well as their relative intensity is evaluated for the simple case of a spherical droplet illuminated at  $\Phi = 0^\circ$  elevation angle. The azimuth angle is varied from  $\theta = 95^\circ$  to  $145^\circ$  in increments of  $10^\circ$ , which corresponds to the angular range for which back-scattering occurs and that was simultaneously possible with the presented experimental setup. A corresponding set of recordings is shown in Figure 4.7, where only the red channel appears as grayscale images. To guide the unaided eye, the uncorrected image channel with *e.g.* spurious backlight is shown, even though the corrected image channels were processed and are discussed below.

Figure 4.8 shows the column-wise averaged intensity of the red image channel after applying the previously introduced ISC method, therefore revealing the distribution intensity of the glare points over the projected gas-liquid interface of the droplet. The column-wise intensity was calculated for five rows of the color-corrected image around the maximum horizontal expanse of the droplet and subsequently averaged over five frames. Note that the third (center) frame of this temporal average corresponds to Figure 4.7. The reconstructed glare point signal allows for a comparison of the position of the glare points in the experiment to the position expected from the calculations by van de Hulst and Wang [233] following Equation (2.36), shown in Figure 4.9.


The good agreement of the position of the glare-point peaks in the experimental images and the theoretically expected intensities allows for an interpretation and subsequent differentiation of the



**Figure 4.7:** Red image channel of the frame with a spherical water droplet before impact for different azimuth angles  $\theta$ . The identical order of angles is found in Figures 4.8 and 4.9. The uncorrected channel was chosen for the better visibility of the droplet contour. (Note that the  $p_{\text{gp}} = 1$  glare point from the blue backlight is visible in the center of the droplet due to cross-talk of the camera and has the same intensity as the background.) Figure adopted from [DBB<sup>+</sup>23].




**Figure 4.8:** Glare-point intensity measured from the ISC corrected red channel from a recording of a spherical droplet in the experiments. The column-wise averaged glare point intensity is plotted on the logarithmic ordinate against the dimensionless variable  $w$  on the abscissa, that describes the position on the scattering plane relative to the projection of the droplet, as introduced in Section 2.3.2. Figure adapted from [DBB<sup>+</sup>23].



**Figure 4.9:** Mie glare-point images calculated by means of Equation (2.36). The logarithmic glare-point intensity is plotted on the ordinate against the dimensionless variable  $w$  on the abscissa. Figure adapted from [DBB<sup>+</sup>23].

contributing orders in the observed glare-point signal from the experiment. In Figure 4.8a only the zeroth-order glare point ( $p_{\text{gp}} = 0$ ) is clearly visible at around  $w = 0.67$ , which is in good agreement with the theoretically expected position. Towards higher scattering angles the zeroth-order glare point moves towards the center of the droplet and reaches  $w = 0.29$  at  $145^\circ$ , as can be seen in Figure 4.8f. The exact glare point positions in geometrical optics were calculated by Equation (2.34) and subsequently Equation (2.35) for further comparison. The results showed a good agreement with the experimentally determined positions.



**Figure 4.10:** Comparison of the glare-point images at scattering angles of (a)  $95.0^\circ$  and the ideal angle for  $p_{\text{gp}} = 0$  glare points  $96.5^\circ$ , (b)  $125.0^\circ$  and the  $p_{\text{gp}} = 3$  rainbow angle  $129.5^\circ$  and (c)  $135.0^\circ$  and the  $p_{\text{gp}} = 2$  rainbow angle  $137.7^\circ$  calculated by means of Equation (2.36). The logarithmic glare-point intensity is plotted on the ordinate against the dimensionless variable  $w$  on the abscissa. Figure adapted from [DBB<sup>+</sup>23].




In Figure 4.8b additionally a third-order glare-point peak ( $p_{\text{gp}} = 3$ ) becomes visible at  $w = 0.73$ , despite the order-of-magnitude difference in intensity to the main peak. The comparison of the predicted intensity with the measured signal confirms the presence of a third-order peak to the right of the strong zeroth-order peak (*cp.* Figure 4.9b). The intensity of the third-order glare point increases towards higher scattering angles as it approaches its rainbow angle at  $129.5^\circ$ . In Figure 4.8c the third-order glare point is clearly visible at  $w = 0.8$ . The theory predicts a second third-order glare point at the  $w = 1$  edge of the droplet that can be confirmed by a clearly visible intensity peak in the experiment. Both third-order glare points are predicted to merge into a single, bright glare point at their rainbow angle, as can be seen in Figure 4.10. However in Figure 4.8d they can still be clearly differentiated at a scattering angle of  $125^\circ$  as evident from the experimental image and the corresponding intensity curve.

A second-order glare point ( $p_{\text{gp}} = 2$ ) first emerges in Figure 4.8e as a large and bright glare point around  $w = -0.92$ . Since the scattering angle lies close to the second-order rainbow angle of  $137.7^\circ$ , the intensity of the glare points is high. In Figure 4.8f two second-order glare points are visible at  $w \approx -0.61$  and close to the edge. As predicted by theory, the second-order glare point, therefore, behaves in the opposite way as the third-order glare point and splits into two peaks for scattering angles larger than the rainbow angle. The comparison of the measured and computed glare points close to the  $p_{\text{gp}} = 2$  rainbow angle (see Figure 4.10) reveals, that in fact, the scattering angle in the experiments has been closer to  $137.7^\circ$  than  $135.0^\circ$ , as the expected relative intensity peaks produce a better match. This elucidates that the intensity of the glare points close to their respective rainbow angles reacts sensibly to small angular changes. However, the position of the glare points does not change significantly, therefore allowing for an accurate prediction of glare point occurrence by Equation (2.36).


The first-order glare point ( $p_{\text{gp}} = 1$ ) was not visible in the range of evaluated azimuthal angles  $\theta$ , as it appears only at scattering angles smaller than  $82.6^\circ$  according to geometrical optics and following Equation (2.34). In a similar manner, the intensity of the fourth and fifth-order glare points is predicted to be too low to result in visible glare points, which is confirmed by the experimental results. Higher-order rainbows accumulate at the  $w = \pm 1$  edges, resulting in intensity peaks from potentially multiple higher-order glare points, complicating the interpretation of the edge peaks. Furthermore, the higher-order glare points become increasingly more sensible to deviations in the scattering angle, since the greater number of internal reflections amplify any deviation. Therefore, the interpretation of higher-order glare points becomes less reliable.

According to van de Hulst and Wang [233], additionally surface waves might contribute to the edge peaks as well. However sixth ( $p_{\text{gp}} = 6$ ) and seventh-order glare points ( $p_{\text{gp}} = 7$ ) appear plausible following the descriptions of van de Hulst and Wang [233]. The intensity peak at the  $w = -1$  edge in Figures 4.8d and 4.8e matches the theoretical position of the sixth-order glare point, while in Figures 4.8e and 4.8f the peaks at  $w = 0.96$ , respectively at  $w = 0.98$  match the expected seventh-order glare point. Additionally, the diffraction on the droplet is visible around the contour of the droplet in all images, contributing to all the glare-point signal in Figure 4.8 and in particular at the edges.

It should be noted that according to Walker [238] the rainbow angle calculated in wave optics differs by  $\approx 0.15^\circ$  for  $p_{\text{gp}} = 2$  and  $\approx 0.21^\circ$  for  $p_{\text{gp}} = 3$  in comparison to geometrical optics. The resulting deviation in the  $p_{\text{gp}} = 3$  glare-point position was calculated to be  $\delta_w < 0.0005$ , which is considered negligible. The angular difference for the  $p_{\text{gp}} = 2$  rainbow angle was even lower, which further confirms that the theoretical calculation provides a valuable and meaningful comparison to the experimentally observed glare points.” from [DBB<sup>+</sup>23, Section 3.2]



**Figure 4.11:** Mean image intensities of the color-corrected red and green image channels for the first 150 frames of the drop impact; (a)  $\Phi = 45^\circ$  elevation angle and (b)  $\Phi = 0^\circ$  elevation angle of the lateral light sources. Figure adapted from [DBB<sup>+</sup>23].



**Figure 4.12:** Frame with the lowest glare point intensity (frame 50) at an elevation angle  $\Phi = 45^\circ$  (a) in comparison to a frame for a similar physical time at  $\Phi = 0^\circ$  (b). The insert in (a) is a contrast-enhanced and brightened repeat of the droplet tip to emphasize the occurrence of glare points. Figure adopted from [DBB<sup>+</sup>23].

### 4.3.3 Influence of the elevation angle

“To further study the influence of the elevation angle of the lateral light sources on the visibility of glare points on the deforming droplet during impact, the lateral light sources were successively raised and simultaneously tilted in order to illuminate gas-liquid interface from different angles, as sketched in Figure 4.1. The elevation angle was varied in a range of  $0^\circ \leq \Phi \leq 45^\circ$  with increments of  $15^\circ$  and the additional angle of  $\Phi = 50^\circ$ , while the azimuth angle was varied from  $\theta = 95^\circ$  to  $145^\circ$  with  $10^\circ$  increment. A drop impact experiment with the same kinematic boundary conditions was recorded for each possible combination of elevation and azimuth angles. The images from the experiments were post-processed with the previously introduced ISC in order to obtain the undisturbed glare point images. Figure 4.11 shows the development of the glare points for experiments performed at  $95^\circ$  azimuth angle in comparison between  $\Phi = 0^\circ$  and  $\Phi = 45^\circ$  elevation angles by the mean intensity of the color-corrected red and green image channels.

The comparison of Figures 4.11a and 4.11b provides evidence that the glare-point intensity fluctuates significantly more at the lower elevation angle of  $\Phi = 0^\circ$  with some frames reaching close to zero intensity, thus indicating that no glare points were visible in these frames. Such low-intensity frames were already apparent in previous experiments at  $\Phi = 0^\circ$  elevation angle in which only the lateral light sources were active, as can be seen in Figure 4.5. At higher elevation angles, in contrast, significantly more frames contained glare points and the development of the intensity between the frames was more consistent as well. This effect increased towards  $45^\circ$ , after which no further improvement was observed. Figure 4.12 shows the least luminous frame in the experiments at  $\Phi = 45^\circ$  elevation angle (Figure 4.12a) and in comparison a frame at  $\Phi = 0^\circ$  (Figure 4.12b) for a similarly shaped droplet.



As evident from the images at  $\Phi = 45^\circ$  elevation angle even the least luminous frame contains clearly visible glare points although faint, as opposed to  $\Phi = 0^\circ$  where no glare points are visible. This implies that glare points must be visible in all other frames, which was confirmed by visual inspection. At the chosen azimuth angle of  $\theta = 95^\circ$ , as expected from the results of Section 4.3.2, only zeroth-order glare points appear independent from the elevation angle. Analogous results were obtained for both the intensity and consistency of the glare points at all other azimuthal angles tested in the experiments, including those that resulted in higher-order glare points. For elevation angles larger than  $\Phi = 0^\circ$ , however, no higher-order glare points were observed after a certain deformation of the droplet due to the interaction of the light with the substrate after drop impingement.” from [DBB<sup>+</sup>23, Section 3.3]

## 4.4 Discussion

“The measurement for glare-point images on the deforming gas-liquid interface with consistent scatter intensity throughout the recorded frames of the drop-impact dynamics is important for a seamless surface reconstruction. As such, the observations and insights of the previous sections will be conflated in order to identify the optimal viewing angles for the lateral light sources below. Further evaluation of 3D-encoding options and limitations follows in the subsequent sections.” from [DBB<sup>+</sup>23, Section 4]

### 4.4.1 Optimal position of the lateral lights

“Figures 4.3, 4.4 and 4.7 indicate that the zeroth-order glare points appear only at locations on the surface of the droplet that align with the complementary angle  $\Phi_c = 90^\circ - \Phi$  to the direction of the incoming light as viewed from the image plane. In general, the phase boundary of the droplet takes any angle that ranges between the contact angle at the contact-line position and  $0^\circ$  degrees at the top of the droplet. Higher angles are possible due to the deformation of the droplet, as can be seen for example in Figure 4.12, but do not necessarily occur. If the contact angle does not exceed the complementary angle to the direction of the incident light, accordingly, frames with no visible glare points might be produced. Consequently, those frames lack the desired additional three-dimensional information of the gas-liquid interface.

The completely horizontal setup with  $\Phi = 0^\circ$  elevation angle reveals various frames without glare points for the stages of the drop impact while the contact angle drops below  $90^\circ$ , which becomes obvious *e.g.* in Figure 4.12. To guarantee the consistent occurrence of glare points, the complementary angle has to remain below the expected range of dynamically varying contact angles in the experiment. Therefore, an elevation angle of  $\Phi = 90^\circ - \Theta_{\min}$  is required, in which  $\Theta_{\min}$  is the minimum receding contact angle observed in the experiments (*cp.* also Figure 4.12). Thus, the optimal elevation angle lies within the range between  $0^\circ$  and the minimum receding contact angle  $\Theta_{\min}$  and immediately relies on the wettability of the substrate. Conversely, an increase in the elevation angle leads to a decrease in the distance between the zeroth-order glare points, which in turn results in a larger relative measurement error of the distance between the glare points. The theoretical optimum for the elevation angle, therefore, is the lowest angle at which glare points are still visible on all frames during impact. An equilibrium contact angle of  $\Theta_{\text{eq}} = 80^\circ$  and a receding contact angle of  $\Theta_{\text{rec}} = 66^\circ$  characterizes the considered water-droplet impact experiment onto a SiOx-substrate. For this configuration, an elevation angle of  $45^\circ$  was found to produce the most continuous glare point signal.


At elevation angles larger than  $0^\circ$ , higher-order glare points mostly disappear as the light that is refracted while entering the droplet interacts with the substrate surface instead of reflecting on the second gas-liquid interface of the droplet. Furthermore, the position of the higher-order glare points and the light reflected from the substrate surface quickly become chaotic when the shape of the droplet deviates significantly from the spherical shape. Both of these effects render an interpretation of higher-order glare points at elevation angles larger than  $0^\circ$  difficult, as the glare point positions become ambiguous. Combined with the aforementioned requirement for elevation angles  $\Phi > 90^\circ - \Theta_{\text{rec}}$  for a continuous glare point signal – and in particular for hydrophilic substrates – this leaves the zeroth-order glare points as the best option for encoding the droplet’s shape. The restriction to the zeroth-order glare point yields the additional benefit of a straightforward description of the underlying physics and interpretation of the results.

The optimal azimuth angle consequently is the angle at which only the zeroth-order glare point is visible. According to GO a scattering angle of  $\theta = 96.5^\circ$  leads to the strongest intensity of the zeroth-order glare point relative to all other glare points at a wavelength of  $\lambda = 632\text{ nm}$ . At this angle the glare points of orders  $3n, n \in \mathbb{N}_0$  coalesce, with the rays forming an equilateral triangle in the droplet, as can be calculated by Equation (2.34). The  $p_{\text{gp}} = 1$  and  $p_{\text{gp}} = 2$  glare points do not appear at  $\theta = 96.5^\circ$ , as the  $p_{\text{gp}} = 1$  rays are limited to  $\theta < 82.6^\circ$ , which is the angle of the edge ray. Likewise, the  $p_{\text{gp}} = 2$  rays are limited to  $\theta > 137.7^\circ$ , which is the rainbow angle.

Since higher-order glare points quickly become less intense outside of a limited range close to their rainbow angle, these orders can be neglected as well. Therefore,  $\theta = 96.5^\circ$  can be regarded as a scattering angle that only yields zeroth-order glare points. The theoretical derivation above matches well with the experimental results at  $\theta = 95^\circ$  azimuth angle (see Figure 4.7), which was sufficiently close to the optimal azimuth angle to only produce visible zeroth-order glare points. Furthermore, the comparison of the theoretically expected glare points for  $\theta = 95.0^\circ$  and  $96.5^\circ$  (see Figure 4.10) reveals that the difference in the intensity distribution is negligibly small. This high tolerance for angular misalignment in turn facilitates a robust, straight-forward calibration of the experimental setup.

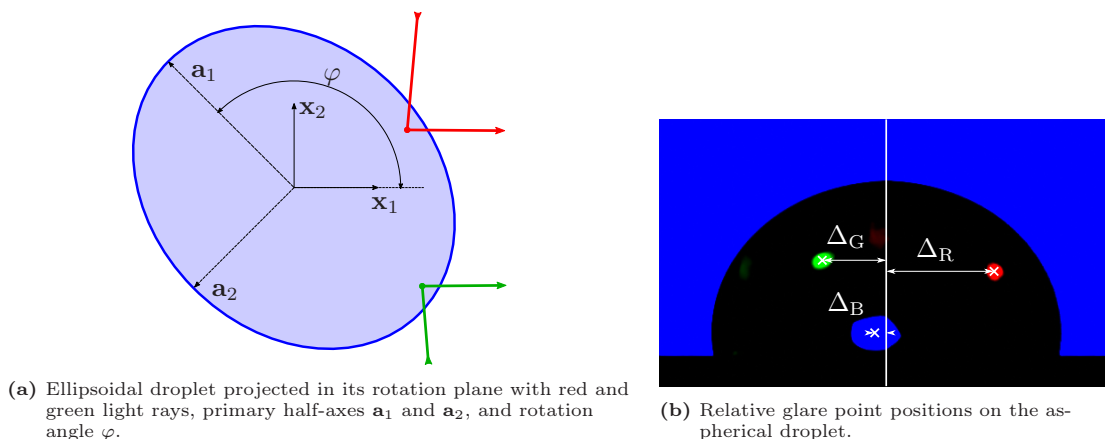
The mild influence of the wavelength-dependent refractive index  $n_r(\lambda)$  of water becomes obvious from the above discussion for red light ( $\lambda_R = 632\text{ nm}$ ) in comparison to the green light ( $\lambda_G = 521\text{ nm}$ ) of the second lateral light source. Accordingly, the green glare points are determined to occur at an optimal angle of  $\theta_G = 96.3^\circ$  by Equation 2.34. This color-dependent  $0.2^\circ$  difference, however, can be considered negligible in comparison to the angular alignment accuracy of the measurement setup.” from [DBB<sup>+</sup>23, Section 4.1]

#### 4.4.2 Three-dimensional encoding of the gas-liquid interface



**Figure 4.13:** Renderings of an ellipsoidal cap with an aspect ratio of  $A_r = 1.25$ . The droplet is rotated in clockwise direction from (a) to (e). The images additionally contain the shadowgraph contour (white line) and the glare point positions (white crosses) from the  $\varphi = 0$  orientation for comparison. Figure adopted from [DBB<sup>+</sup>23].

“In the following section the informative value of the glare points regarding the three-dimensional shape of the droplet will be analyzed. For this purpose, a render pipeline for the generation of



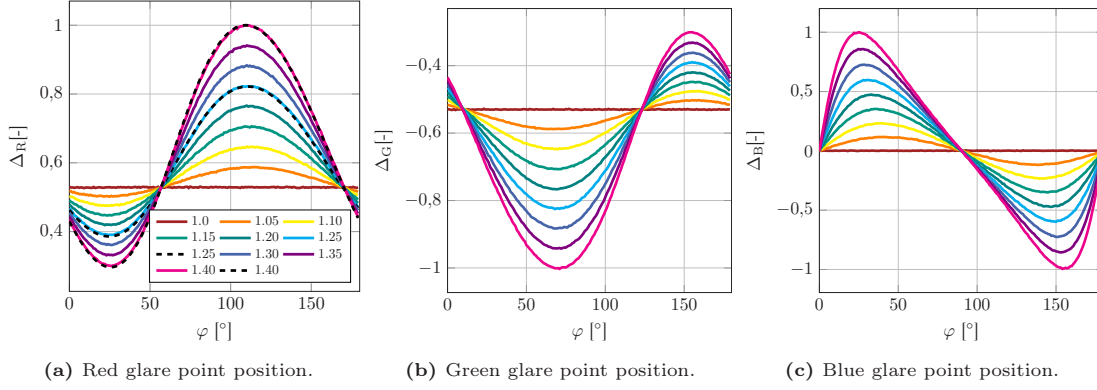
**Figure 4.14:** Geometrical description of the aspherical droplet (a) and definition of the glare point positions  $\Delta_R$ ,  $\Delta_G$  and  $\Delta_B$  relative to the symmetry axis of the shadowgraph contour (b). Figure adopted from [DBB<sup>+</sup>23].

synthetic images was set up in the 3D rendering software *Blender* [39] in combination with the ray-tracing library *LuxCore* [19], which allows to create realistic images of arbitrary droplet geometries that closely match the experimental images<sup>1</sup>. In order to evaluate the evolution of the glare-point images for non-axisymmetric droplets, ellipsoidal cap geometries with varying aspect ratios were produced by stretching the geometry of a gas-liquid interface extracted from a numerical simulation along one axis in the horizontal plane by factors of  $A_r = 1.0$  to  $A_r = 1.4$  with increments of 0.05. Finally, the render setup, configured with an azimuth angle of  $\theta = 95^\circ$  and an elevation angle of  $\Phi = 45^\circ$  was used to render glare-point images from the droplet geometries at different rotation angles. The renderings created by the aforementioned render pipeline can be seen in Figure 4.13 for the aspect ratio  $A_r = 1.25$ .

Figure 4.14a shows the geometrical description of the ellipsoidal droplet that has the aspect ratio  $A_r = \|\mathbf{a}_1\|/\|\mathbf{a}_2\|$  and is rotated by  $\varphi$  relative to the space-fixed frame of reference. Figures 4.13a and 4.13e show the geometry rotated by  $\varphi = 0^\circ$  and  $180^\circ$ , so that the major axis of the ellipsoid is pointing towards the observer. In order to emphasize the differences between the rendered glare point images at different  $\varphi$  all images in Figure 4.13 contain the contour of the  $\varphi = 0^\circ$  droplet and the positions of the glare points marked by white crosses. The relative position of the glare points in dependence of  $\varphi$  and  $A_r$  was measured relative to the center line of the projection, as indicated in Figure 4.14b. The measured positions are plotted in Figures 4.15a, 4.15b and 4.15c for the red and green zeroth-order glare points, and the blue first-order glare point, respectively. The aspect ratio  $A_r$  appears as a family parameter in the diagrams and angle  $\varphi$  is the functional variable.

As can be seen in Figure 4.13 the position of all three glare points moves relative to the shadowgraph contour in dependence of the rotational angle  $\varphi$  and the aspect ratio  $A_r$ . The positions of the red and green glare points appear as phase-shifted and mirrored functions since both glare points are created through interface reflection ( $p_{gp} = 0$ ). Their location, therefore, depends on the point on the droplet's gas-liquid interface that matches the required angle for interface reflection  $\theta_S = (180^\circ - \theta)/2$ , which is half the angle included between the camera and the respective light source. The relative position of the blue glare point shows a different dynamic, as it arises from the transmission through the droplet with correspondingly two refractions on the gas-liquid interface ( $p_{gp} = 1$ ).

<sup>1</sup> The setup of the rendering pipeline is described in further detail in Chapter 5.



**Figure 4.15:** Positions  $\Delta_i$  of (a) the red and (b) green  $p_{\text{gp}} = 0$  glare points and (c) the blue  $p_{\text{gp}} = 1$  glare point relative to the shadowgraph contour of the droplet for increasing aspect ratios  $A_r$  and normalized with the maximum absolute value. The dotted lines in (a) represent the analytical solution for comparison. Figure adapted from [DBB<sup>+</sup>23].

With the assumption that the shape of the droplet can be described as an ellipsoidal cap with two planes of symmetry along the major and minor axis, two parameters – the rotational angle  $\varphi$  and the aspect ratio  $A_r$  – are sufficient to uniquely determine the three-dimensional shape and position of the droplet. As can be seen in Figure 4.15, the relative position of the three (color-specific) glare points are three mutually independent families of functions, since their amplitudes, as well as the positions of their maxima, depend on the aspect ratio. For a given measurement setup that determines the position of the light sources and the volume of the droplet,  $\varphi$  and  $A_r$  can thus be unambiguously determined from the system of the two  $p_{\text{gp}} = 0$  glare-point positions. The relative glare-point positions were determined analytically (see dashed lines in Figure 4.15) and show a good agreement with the results from the rendering and, therefore, provide further evidence on the applicability of the considered approach. The derivation of the underlying necessary Equations (A.8) and (A.9) can be found in the Appendix A.1. The two  $p_{\text{gp}} = 0$  glare points obviously encode the location of the droplet’s phase boundary in the object plane and the local orientation at two additional points outside the shadowgraph plane. For the presented setup the local orientation of the gas-liquid interface at the position of the lateral glare points is  $\theta_S = 42.5^\circ$  in the azimuth angle and  $\Phi = 45^\circ$  elevation angle.” from [DBB<sup>+</sup>23, Section 4.2]

### 4.4.3 Limitations of the method

“The positive results of the color correction, as seen in Figure 4.5 suggest that the proposed model for the light transport and the underlying assumptions are valid. After the color correction, the blue image channel shows only a pure shadowgraphy image without any specular reflection of the lateral light sources. This result demonstrates that the presented method is able to produce standard shadowgraphy images in a high quality with the addition of further illumination angles of the droplet’s interface that encode additional three-dimensional information of the droplet’s gas-liquid interface in the glare points as shown in Section 4.4.2.

The two evaluated methods (ASC and ISC) allow for an accurate reconstruction of the glare point and shadowgraph images by the correction of the mutual disturbance in the image channels, that resulted from polychromatic light and spectral sensitivity of the camera sensor. The analytical reconstruction (ASC) can reach a potentially higher reconstruction quality, if the spectral characteristics of all optical components can be determined accurately a-priori. If such information is missing or incomplete, then the in-situ method (ISC) can be used to obtain the correction matrix, which introduces additional measurement imperfections such as *e.g.* image noise.

It should be noted that the proposed color-correction models neglect chromatic aberration and dispersion. Due to the dependence of the refractive index on the wavelength of the incident light the (polychromatic) transmitted light would be deflected at a slightly different angle. Therefore all glare points of order  $p_{\text{gp}} = 1$  and higher are affected by dispersion. Due to a higher amount of interactions with the gas-liquid interface, the dispersion becomes increasingly severe for higher-order glare points. As seen in Figure 4.9 first, second, and third-order glare points, each with an increasing magnitude of dispersion can be observed in the experiments. Higher-order glare points might appear in the images as well, however, since these lie close to the edges of the projected droplet, their spatial extent is very small compared to the lower-order glare points so that the effects of dispersion can be neglected. A significant level of dispersion could lead to a faulty color correction since different areas of the glare points on the images would be composed of different wavelengths. This effect would consequently distort the pixel-wise reconstruction.

For the optimal angular position of the lateral lights determined in Section 4.4.1 only  $p_{\text{gp}} = 0$  glare points do appear, which are produced by pure interface reflection and thus remain unaffected by dispersion. In contrast, the  $p_{\text{gp}} = 1$  glare point of the backlight is affected by dispersion, since it is refracted by the gas-liquid while entering and again while leaving the droplet. However, the influence of dispersion can be assumed to be negligibly small due to a small change in the index of refraction for the narrow-banded spectrum of the lateral LED light sources. The maximum angular dispersion in the first-order glare point, *i.e.* the deviation in the angles of the refracted light, can be estimated to be lower than  $0.25^\circ$  for the given spectra of the light sources using Snell's law [83].

Chromatic aberration causes differences in defocus and magnification between the color channels of the image, which can introduce errors to the color correction. Therefore, in the presented experiments a *Schneider-Kreuznach Apo-Componon 4.0/60* enlarging lens with apochromatic correction has been used to mitigate chromatic aberration. Even though beyond the scope of the present work, further investigation into the effect of chromatic aberration on the quality of the color correction is envisioned to further improve the accuracy of the method.

The intensity of the light sources was found to have no significant influence on the results of the color correction. The intensity of the color-corrected images scaled linearly with the light source intensity, but otherwise the channels had undergone the same color correction for different input intensities, thus validating the assumption of the linearity of the camera sensor. While the assumption of a linear response of the chip to the incident light holds true for intensities within the dynamic range, it neglects clipping in the regions of glare points and other highlighted image regions. Since one or potentially multiple color channels can be clipped, changes in hue occur [165] and a proper calculation of the color correction is inhibited. The resulting loss of information for intensities above the dynamic range skews the correction function such that the mutual disturbance of the color channels cannot be fully removed in the overexposed frames, which has been discussed along selected frames in Figure 4.5. Note that such frames can still be improved by the correction function. However, reduced light intensity mitigates this problem at the cost of likewise less pronounced glare points, especially on frames with larger areas of reflection. Therefore, the light source intensity needs to be fine tuned for conserving a maximum of information in dim and bright images.

The influence of the chosen elevation angle  $\Phi$  on the resulting continuous registration and homogeneous intensity of recorded glare points has been already discussed in Section 4.4.1 for convex gas-liquid interfaces. Non-convex droplet shapes, which occur during the dynamic deformation of the droplet during impact, comprise multiple glare points at each position on the droplet surface that has the orientation  $\Phi_c$  and  $\theta_s$ . These glare points split and merge depending on the deformation of the gas-liquid interface as indicated *e.g.* in Figure 4.3, where multiple glare points are

visible at each capillary wave (see also Figure A.4 for the full time series). The appearance of  $p_{\text{gp}} = 0$  glare points on non-convex deformed droplets facilitates the encoding of additional three-dimensional information on a larger amount of points on the surface. However, a reconstruction of these droplet shapes will be more complicated in comparison to the spheroidal cap shape, for which an unambiguous encoding of the volumetric shape was confirmed in Section 4.4.2.

The overall good agreement of the glare-point positions in the experiment in dependence of the azimuthal angle  $\theta$  to the theoretically expected positions (see Figures 4.7) for all tested azimuthal angles, confirms the assumption of a circular cross-section for the undeformed droplet. However, small differences in the position of the glare points were observed that might originate either from a deviation from the spherical shape and/or from alignment errors in the experimental setup. On the one hand, small changes in the scattering angle due to alignment errors in the azimuth angle of the lateral light sources can result in significant changes in both glare-point intensity and position, which particularly cumulates for higher-order glare points due to a larger number of interactions with the gas-liquid interface. In particular, the intensity increases sharply if the azimuth angle approaches the rainbow angle of the respective order of scattering, as can be seen in Figure 4.10. The reduction of the measurement system to  $p_{\text{gp}} = 0$  and  $p_{\text{gp}} = 1$  glare points, therefore, significantly limits this error. On the other hand, alignment errors in the elevation angle of the lateral light sources result in a shift of the glare-point position. For low elevation angles, the resulting error becomes negligible for small angular deviations, as the horizontal movement of the glare points is small due to its trigonometrical relation to the elevation angle. For large elevation angles, however, this error source has to be considered.

A third error source are deformations of the droplet from the spherical shape, resulting in an elliptical cross-section in the scattering plane for the lateral light sources. As shown in Section 4.4.2 the aspect ratio of an ellipsoidal droplet has a significant influence on the position of the glare points such that the deformation of the droplet in any axis has to be considered. Irregular deformations of the droplet will inevitably introduce deviations in the position of the glare point, which can lead to reconstruction errors. The present experiments of a droplet impinging in the normal direction onto a flat substrate, however, revealed no significant irregular deformations. Instead, a deformation in the vertical direction from the detachment of the droplet at the needle was observed that caused a vertical oscillation and resulted in a spheroidal shape of the droplet. Note that the aspherical shape of the droplet in this case can be determined analogously to Section 4.4.2 from the position of the glare points.” from [DBB<sup>+</sup>23, Section 4.3]

## 4.5 Concluding remarks regarding glare point shadowgraphy

“The present work has elaborated on how an extension of the shadowgraphy measurement method by lateral light sources with distinct spectra can be used to gather additional information about the three-dimensional shape of a gas-liquid interface for the considered example of a liquid droplet impacting onto a solid surface. From the spectral power distribution of the light sources and the spectral sensitivity of the camera, a color correction function is determined by spectral integration, which allows for a clear separation of the two reflection images from the lateral light sources and the shadowgraphy image from the backlight. Application of the proposed color-correction function successfully demonstrated that the narrow-banded polychromatic light from three differently colored LEDs can be used to gather three distinct representations of the gas-liquid interface of an impinging droplet.

The presented method of capturing the reflection images from the lateral light sources in each one channel and the shadowgraphy image in the third channel of an RGB image allows for the use



of a simple single-camera setup. This facilitates an easier calibration and reduces the complexity of a volumetric reconstruction from the images in comparison to a multi-camera setup, since no spatial or temporal matching of the produced images is necessary, which otherwise would require an exact measurement of the relative positions between the light sources and the camera in respect to the droplet. Note, however, that three representations captured with the single-camera [glare-point shadowgraphy] setup certainly contain less information on the droplet interface than three individual pictures from cameras at different angles.

It was shown that the measurement setup is still able to produce a clear shadowgraphy image that is not impaired by the additional light sources. The two additional channels can be consequently considered as a pure extension to a canonical setup. The additional perspectives from the glare points on the gas-liquid interface by the lateral light sources lay the ground for a volumetric reconstruction of the phase boundary. Since the geometrical properties of the experimental rig – in particular the relative positions of the camera and the light sources with respect to the droplet – are known, features of the droplet surface can be reconstructed by the images of the glare points using the theory of geometrical optics, [as shown by 82, 91]. It was accordingly demonstrated that the additional information encoded in the glare points can be used to determine the aspect ratio and rotation angle of an ellipsoidal droplet. Future studies might extend this approach to more complicated shapes that appear during drop impact, including non-convex shapes with multiple glare points. Additional glare points from further viewing angles could then be employed to determine the larger number of parameters necessary to describe such complex geometries. [...]

A strong and continuous glare point signal in the images over the complete range of droplet shapes during impact is required for a reconstruction of the whole droplet dynamics. The present work revealed that the optimal elevation angle for a continuous signal depends on the wettability of the substrate and is, therefore, case dependent, whereas an azimuth angle of  $96.5^\circ$  and  $96.3^\circ$  for the red and green lateral light sources, respectively, proved to be optimal to produce unambiguous glare point images. At the resulting scattering angle, only zeroth-order glare points appear, which prevents ambiguities from the reflection on the substrate from higher-order glare points that were previously refracted by the gas-liquid interface. In addition, the position of zeroth-order glare points is less affected by alignment errors of the lateral light sources in the scattering angle, compared to higher-order glare points that see a larger number of interactions with the gas-liquid interface and therefore amplify the error at every interaction with the gas-liquid interface. However, higher-order glare points have the potential for a more precise measurement of the surface orientation, therefore becoming suitable for experiments in which the measurement of small deviations in the droplet's surface is of interest.

For axisymmetric droplets, different elevation angles for the two lateral light sources could be exploited to produce a higher number of unique illumination orientations on the droplet, which could be used in order to detect sites of air entrainment. Additional glare points from further illumination angles could be envisioned to encode the position and orientation of the gas-liquid interface at a larger amount of points to gather additional information. However, caution needs to be taken to ensure that the resulting glare points can be clearly distinguished from each other in order to avoid ambiguities. [...] While the measurement setup was presented along the example of an impinging droplet, it can be expanded to other applications involving gas-liquid interfaces like the two-phase flow in fuel cells, for instance.” from [DBB<sup>+</sup>23, Section 5]





## 5 Synthetic training data generation

In this chapter, a data processing pipeline for the generation of synthetic training data by means of physically-based ray tracing on the basis of the results from numerical simulations is developed. The suitability of the rendered images as training data for the neural network in the reconstruction framework is evaluated through the examination of their physical correctness and visual similarity to experimental images. The chapter concludes with a brief overview of the numerical simulations and the respective datasets derived from them. Part of the results has previously been published in the articles *Spatio-temporal reconstruction of droplet impingement dynamics by means of color-coded glare points and deep learning* by Dreisbach *et al.* [DKS24] and *Interface reconstruction of adhering droplets for distortion correction using glare points and deep learning* by Dreisbach *et al.* [DHK<sup>+</sup>25]. Where findings have already been published in the referenced work, these are clearly marked and identified as related quotations.

### 5.1 Problem statement


The aim of the synthetic training data generation is to obtain labeled datasets that are suitable for training the neural network for the reconstruction of recordings obtained in the previously introduced glare-point shadowgraphy experiments. As introduced in Section 3.3, labeled data, *i.e.* one-to-one paired input images and output three-dimensional (3D) ground truth shapes of the droplets gas-liquid interface are required due to the chosen supervised learning approach. The three-dimensional ground truth is obtained by means of numerical simulation. However, due to the aforementioned matching problem between simulation and experiment (see Section 3.3), the use of synthetic image data generated on the basis of the numerical ground truth droplet shapes is necessary. For that purpose, a setup for synthetic image generation has to be developed that accurately represents the relevant optical phenomena involved in the image formation in experiments with the glare-point shadowgraphy technique, in order to reproduce the recordings from the experiments. The appearance of the glare points, in particular their location and size has to be reproduced in a physically correct manner to ensure that the training data faithfully represents the encoded three-dimensional information in the glare points. Furthermore, the defocus of the glare points and the shadowgraph contour, as well as the intensity of the different light sources has to be modeled accurately. “It is well known that neural networks trained on synthetic data exhibit a reduced accuracy once employed on real data. This performance gap is dependent on the difference between the feature distributions of the real and synthetic datasets [40, 207].” from [DKS24, Section 2.2] Therefore, it is crucial that the synthetic training images accurately reproduce the visual appearance of the recordings from the experiments, to ensure that the convolutional neural network (CNN) can learn the extraction of similar visual features from the synthetic data. Furthermore, the image generation setup has to be efficient, in order to allow for the production of large amounts of image data within a reasonable timeframe. The data from two different numerical investigations is employed, specifically, the phase-field simulation of droplet impingement on flat and structured substrates and the Volume-of-fluid (VoF) simulation of adhering droplets in shear flows. The 3D direct numerical simulation of impinging droplets on flat and structured surfaces conducted by Fink *et al.* [64] in the framework of the phase-field method exhibits a high spatial

and temporal resolution. Conversely, the 3D VoF simulation of adhering droplets conducted by Dreisbach *et al.* [DHK<sup>+</sup>25] has a significantly lower resolution in comparison. The calculations of both numerical simulations were conducted with the open-source computational fluid dynamics (CFD) software OpenFOAM<sup>®</sup>, which provides the results of the simulation in a data format that allows for the extraction of the three-dimensional gas-liquid interface shape. However, the different resolution of the numerical simulations necessitates a dedicated processing of the extracted three-dimensional gas-liquid interface shapes.

## 5.2 Synthetic image generation

As discussed in Section 2.3.2 and further demonstrated by the experimental results in Section 4.3.2 the theory of geometrical optics (GO) offers an accurate approximation for the image formation in the glare-point shadowgraphy experiments, specifically for the location and size of the glare points and the shadowgraph contour. The formulation of the synthetic image generation problem within the framework of geometrical optics enables a simplified approach compared to wave optics. Furthermore, ray-tracing algorithms on the basis of geometrical optics with computationally efficient codes are available facilitating fast image synthesis. Therefore, the setup for the generation of synthetic images is implemented in the rendering environment *Blender* [39] with the *LuxCore* [19] package that enables physically correct ray tracing. Specifically, a three-dimensional virtual scene that accurately represents the setup of the glare-point shadowgraphy experiments is modeled in the open-source computer graphics software *Blender*. Subsequently, the rendering engine *LuxCore* is employed to perform the necessary calculations for image formation in order to convert the 3D scene into a 2D image, *i.e.* rendering. *LuxCore* is based on the approach of physically based rendering (PBR) [170], a ray-tracing algorithm aimed at the physically accurate calculation of light transport. Within the framework of PBR, the specular reflection and refraction of light at the gas-liquid interface are accurately calculated from the Fresnel equations [83]. Furthermore, *LuxCore* entails physically-based models for the diffuse reflection and transmission of light at rough surfaces, which allow for an accurate representation of the interaction of light with the various substrates in the experiments. For a comprehensive review of physically based rendering the reader is referred to the monograph by Pharr *et al.* [170].

**Rendering setup** The rendering setup depicted in Figure 5.1 was implemented in *Blender* with the aim of accurately modeling the major optical components relevant to the image formation in the experiments. The lateral light sources and the backlight are modeled as unpolarised LED



**Figure 5.1:** Optical setup modeled in the rendering environment *Blender*. Insert: Magnified view of the immediate vicinity of the droplet, indicating the position and orientation of the lateral light sources.

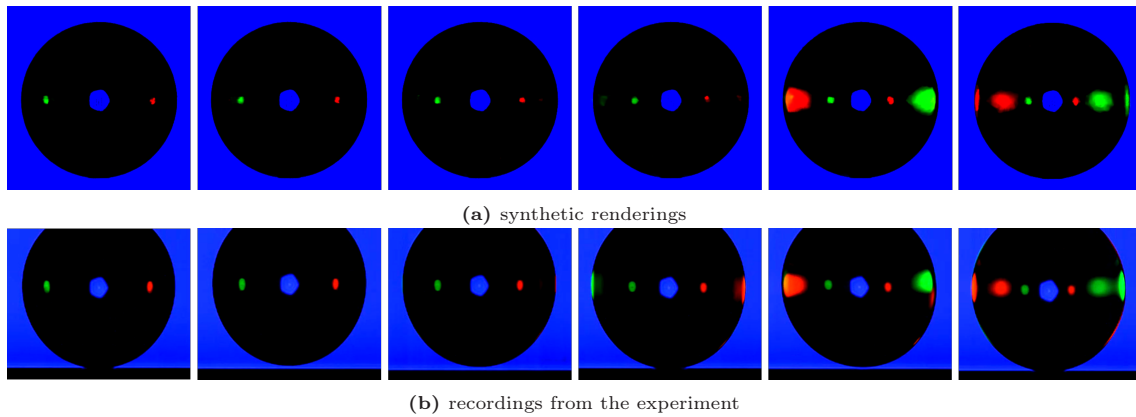
lights, each with a single wavelength and a square-shaped emitting area. The divergence angle of the lateral light sources was set to  $4^\circ$  in order to reproduce the spatial extent of the glare points observed in the experiments. As can be seen in Figure 5.1, the lateral light sources are positioned close to the droplet to achieve maximum yield of the diverging light. Furthermore, the light sources (including the backlight) are larger than the maximum expected droplet size in order to ensure an illumination of the whole droplet during its dynamic deformation, similar to the experiments. It should be noted, that the spatial extent of the glare points depends on the divergence angle of the light and the size of the receiving aperture, but not the size of the illuminating light beam, as discussed in Section 2.3.2. The backlight is modeled as a source of parallel light, which is an accurate representation of the collimated light beam in the experiments. The lateral light sources are oriented relative to the droplet so that the scattering and elevation angles match the experimental setup. The particular angles of the setup demonstrated in Figure 5.1 are  $\theta = 95^\circ$  scattering angle and  $\Phi = 45^\circ$  elevation angle, which were identified as an optimal setup in Section 4.4.1. The position and orientation of the lateral light sources can be changed with little effort, which allows for rendering at various scattering and elevation angles, required for different variations of the experimental setup. “The intensities of the lateral light sources, as well as the backlight, were tuned empirically to be in accordance with the images observed in the experiments.” from [DKS24, Section 2.2] The droplet is modeled as a transparent volumetric body with a constant index of reflection of  $n_r = 1.324$  and, consequently, dispersion is neglected. “In fact, dispersion only needs to be considered for the  $p_{gp} = 1$  glare point that is produced by the backlight, since the lateral glare points only undergo direct interface reflection ( $p_{gp} = 0$ ). Since the LED lights have a narrow-banded spectrum the effect of dispersion is small and can therefore be neglected [...]” from [DKS24, Section 2.2] The gas-liquid interface of the droplet is modeled as a perfectly smooth surface to ensure that only specular reflection and refraction are calculated by the rendering engine. “As the higher-order glare points ( $p_{gp} \geq 1$ ) are intentionally suppressed by the experimental setup, only the  $p_{gp} = 0$  glare points resulting from the lateral light sources and the  $p_{gp} = 1$  glare point from the backlight are modeled. Furthermore, reflection on the substrate is not considered.” from [DKS24, Section 2.2] The multi-lens system of the objective lens is represented by one aspheric plano-convex lens with a numerical aperture of  $N_A = 0.42$  and 60 mm focal length, resulting in a similarly low optical aberration. This approach significantly simplifies the calculation of light transport during rendering and, thus, enables reasonably low rendering times in contrast to the simulation of a multi-lens system. In order to ensure high spatial accuracy of rendered images, the object meshes representing the droplet and the objective lens are required to have a sufficiently high resolution. Therefore, the mesh of the objective lens is resolved by 525,620 vertices, while the mesh of the droplet is resolved by approximately 2,500,000 vertices, depending on the shape of the gas-liquid interface. The aperture of the objective lens is modeled as an orifice in the shape of a regular pentagon with a circumdiameter of  $D_A = 10.6$  mm, which is placed within the objective lens, 5.165 mm behind its principle image plane. Finally, the image of the illuminated droplet is projected by the objective lens on a screen, from which the images are rendered.

**Data processing** “A *Python*-script was set up to import the droplet geometries into this setup and optionally apply transformations through rotation and scaling before rendering a synthetic image of the droplet. The required geometry of the gas-liquid interface was extracted from the results of direct numerical simulations (DNS) of droplet impingement conducted by Fink et al. [64] in the framework of the phase-field method. Hereby, the liquid-gas interface is retrieved by extracting isosurfaces of the phase indicator at  $C = 0$  with  $C = -1$  and  $C = 1$  corresponding to gas and water, respectively.” from [DKS24, Section 2.2] Analogously, the droplet geometries for adhering droplets subjected to external shear flows are obtained from the VoF simulations conducted by Dreisbach *et al.* [DHK<sup>+</sup>25] by extracting isosurfaces of the volume fraction at  $\alpha = 0.5$ . “In order

to accurately simulate the glare point behavior, smooth surface meshes with a sufficiently high resolution are required.” from [DHK<sup>+</sup>25, Section 2.2] The resolution of the droplet geometries obtained from the phase-field DNS is refined by two consecutive iterations of mesh subdivision with linear interpolation [57]. The coarser meshes of the adhering droplet geometries are refined by “three consecutive iterations of mesh subdivision using linear interpolation [57], followed by a smoothing of the mesh with the Taubin filter [221] for 50 iterations” from [DHK<sup>+</sup>25, section 2.2], which approximately conserves the volume of the droplets. Subsequently, the adhering droplet meshes are again refined by two iterations of mesh subdivision, in order to attain a sufficiently high resolution. For each time step in the respective simulations, multiple synthetic images are rendered in the previously introduced setup, while the droplet geometry is rotated around the vertical axis in increments of  $10^\circ$  for a total of  $360^\circ$ , in order to obtain a training dataset that enables the reconstruction of the droplet dynamics from arbitrary observation angles. Moreover, the rendering at different observation angles increases the size and variation of the training dataset with respect to the visual appearance of the input images, which is aimed to improve the optimization of neural networks. The average rendering time per image amounts to 18.7 seconds for two parallel processes.


### 5.3 Validation

The suitability of the rendered synthetic images as training data for the convolutional neural network is evaluated with respect to the stated goals in Section 5.1. First, the physical correctness of the modeled rendering setup is validated by the analysis of the predictive accuracy for the relative location and spatial extent of the glare points and the shadowgraph contour. Subsequently, the quality of the generated images is evaluated based on the degree of visual similarity to the images recorded in the experiments.




**Figure 5.2:** Comparison of synthetic images and glare-point shadowgraphy recordings of spherical droplets at various scattering angles  $\theta$ , ranging from  $\theta = 95^\circ$  (left) to  $\theta = 145^\circ$  (right) in increments of  $10^\circ$ .

**Physical correctness** Synthetic images were rendered for a spherical droplet geometry with a variation of the azimuth angle  $\theta$  of the lateral illumination sources from  $\theta = 95^\circ$  to  $\theta = 145^\circ$  in increments of  $10^\circ$  at a constant elevation angle  $\Phi = 0$ , analogously to the experimental investigation conducted in Section 4.3.2. For validation purposes, higher-order glare points ( $p_{gp} = 2$  and  $p_{gp} = 3$ ) are calculated by the rendering engine in the current investigation, in contrast to the later data generation. The rendered images are shown in Figure 5.2a in comparison to images recorded in the experiments in Figure 5.2b. It should be noted that the recordings presented in Figure 5.2b are the same as those employed in the previous investigation, for which only the red image channel is shown



**Figure 5.3:** Comparison of the glare-point intensity between the synthetically rendered images (colored lines) and images recorded in the experiments (black lines), measured from the red color channel of the RGB image at different scattering angles  $\theta$ , ranging from  $\theta = 95^\circ$  (a) to  $\theta = 145^\circ$  (f). The logarithmic glare-point intensity is plotted on the ordinate against the dimensionless variable  $w$  on the abscissa.




**Figure 5.4:** Theoretical intensity of the glare points according to Lorenz-Mie theory, calculated through Equation (2.36) at different scattering angles  $\theta$ , ranging from  $\theta = 95^\circ$  (a) to  $\theta = 145^\circ$  (f). Figure adapted from [DBB<sup>+</sup>23] and repeated from Figure 4.9 to allow for direct comparison.

in Figure 4.7. The comparison of synthetically rendered images with their counterparts from the experiment reveals an overall good agreement. The relative location and spatial extent of the glare points and the shadowgraph contour are consistently reproduced in the renderings at the different scattering angles. In particular, higher-order glare points are accurately calculated, as evidenced by the red second-order glare points on the left hemisphere of the droplets at  $\theta = 135^\circ$  and  $\theta = 145^\circ$  (see Figures 5.2a right and second from right) and, respectively, the green second-order glare points on the right hemisphere.

In order to evaluate the accuracy of the glare point simulation by the rendering setup in more detail, the horizontal intensity distribution of the glare points was calculated and compared to the experimental results, as well as the glare-point intensity predicted by Lorenz-Mie theory (LMT) [148]. For that purpose, the column-wise averaged intensity in the red image channel was calculated for the five central rows around the maximum horizontal expanse of the droplet in the synthetically rendered images shown in Figure 5.2a. The intensity of the red glare points in the recordings from the experiment was determined analogously. The glare-point intensity of the synthetic images rendered at different scattering angles is plotted in comparison to the intensity distribution of the images from the experiments in Figure 5.3 and the intensity distribution according to LMT calculated by Equation (2.36) is plotted in Figure 5.4.

The comparison of the glare-point intensities between the synthetic rendering and the experiments in Figure 5.3 confirms the previous findings, demonstrating a good agreement for the position and size of the glare points. In particular, the  $p_{\text{gp}} = 0$  glare point was accurately reproduced at all scattering angles. The accurate reproduction of the spatial extent of the glare points validates the modeling for the divergence of the lateral light sources and the aperture of the objective lens in the rendering setup. The comparison of the intensity distributions for the rendering and the experiment at  $\theta = 115^\circ$  reveals that the lower intensity third-order glare points can be accurately reproduced in the renderings. Furthermore, the comparison with the glare-point intensity distributions predicted by LMT in Figure 5.4, which serves as a more accurate reference for the location of the glare points, reveals a precise agreement of the simulated and theoretically predicted glare points. As higher-order glare points are increasingly susceptible to any modeling errors due to a higher number of interactions with the gas-liquid interface, the accurate simulation of the position of the third-order

glare point demonstrates a high accuracy of the optical model. The position of the third-order glare point for  $\theta = 125^\circ$  at  $w = 0.8$  precisely matched the theory, however, its intensity was underestimated. The glare point predicted for  $\theta = 125^\circ$  close to the right edge of the droplet ( $w = 1$ ) and visible in the experiment, was not reproduced by the rendering model. However, the  $p_{\text{gp}} = 2$  glare point close to the left edge of the droplet at  $\theta = 145^\circ$  and the  $p_{\text{gp}} = 3$  glare point at  $\theta = 115^\circ$  are accurately represented, which indicates that the glare point at the edge might be of order ( $p_{\text{gp}} \geq 3$ ) and, therefore, was not simulated by the rendering setup. The existence of high-order glare points ( $p_{\text{gp}} \geq 3$ ) at the edges of the droplet was suggested by van de Hulst and Wang [233] and is, furthermore, indicated by the intensity peaks in Figure 5.4 at  $w = -1$  and  $w = 1$ .



**Figure 5.5:** Comparison of synthetic images (left) and recordings from the experiment (right) at similar time steps.

**Representation of real images** Synthetic images were rendered from geometries of the gas-liquid interface extracted from the aforementioned phase-field simulation of droplet impingement on a flat surface [64]. The numerical simulation considered the impingement of a  $d_0 = 2.1$  mm water droplet on a flat hydrophobic Polydimethylsiloxane (PDMS) substrate at an impact velocity of  $u_0 = 0.62$  m/s. The position and orientation of the lateral light sources were configured for a scattering angle of  $\theta = 95^\circ$  and an elevation angle of  $\Phi = 45^\circ$ , to produce glare points on all frames during the dynamic deformation of the droplet during impingement.

“The resulting synthetic [glare-point shadowgraphy] images are shown in Figure 5.5 (left) in comparison to real images from the experiment for a similar physical time after droplet impingement. The experiments were conducted with the setup introduced in Subsection 4.1 and featured the impingement of a  $d_0 = 2.08$  mm water droplet at an impact velocity of  $u_0 = 0.7$  m/s on a flat hydrophilic silicon oxide (SiOx) substrate. It should be noted that while the impact parameters are not identical to the simulation, they are similar enough to allow for phenomenological comparison of the obtained images. As can be seen in Figure 5.5, the position and appearance of the glare points, as well as the focus of the shadowgraph contour and the glare points was reproduced accurately in the synthetic images. The illumination of the lateral green and red light sources resulted only in  $p_{\text{gp}} = 0$  glare points in the rendering, which agrees well with the experiments and theory [...] [see Section 4.4.1]. Furthermore, the blue  $p_{\text{gp}} = 1$  glare point from background illumination is reproduced in the rendering.

Note that the complex deformed shape of the gas-liquid interface in the early spreading phase [181] produces multiple glare points, as each capillary wave with an appropriate surface angle casts a glare point (Figure 5.5 upper row). These glare points split and merge over time due to the temporal evolution of the interface deformation, which leads to a highly non-linear behavior of



the color-coded glare points. Consequently, the employment of deep learning techniques for the reconstruction of the deformed gas-liquid interface on the basis of these glare points is suggested.

The shadowgraph contour and the glare points originate from different depth coordinates and the focal plane of the objective lens is located closer to the shadowgraph contour. Consequently, the glare points show a varying degree of defocus in the experiments, depending on the width of the droplet, as can be seen in Figure 5.5. Furthermore, their defocus is dependent on the focal length and aperture of the camera equipment and, therefore, specific to the experimental setup. As demonstrated by the rendering results, the defocus of the glare is reproduced well, which results from an accurate representation of the objective lens and aperture in the render setup. The shape of the  $p_{\text{gp}} = 1$  glare point is influenced by the shape of the aperture, which is captured well in the rendering, as can be seen in Figure 5.5. However, the contour of the  $p_{\text{gp}} = 1$  glare points appears to be too crisp in the rendering compared to real experiments, which indicates that some of the optical phenomena that were not modeled are relevant here. This might include dispersion, absorption, or light scattering by pollution in the water. It should be noted that the render engine uses numerical approximations for the calculation of reflection and refraction, which potentially influence the outcome of the  $p_{\text{gp}} = 1$  glare point. Furthermore, since the experimental setup is mirrored in the render setup, the projection of the droplet to the camera is closely approximated, so that perspective errors can be neglected. The overall good agreement between synthetic and experimental images demonstrates the capability of the rendering approach for realistic synthetic data generation and thus validates the assumptions made in the setup of the rendering environment.” from [DKS24, Section 3.1]

## 5.4 Training datasets

The labeled datasets used for the optimization of the neural network were obtained from the numerical simulations of droplet impingement and adhering droplets subjected to external shear flow by means of the data processing methods introduced in Section 5.2. The three-dimensional direct numerical simulation conducted by Fink et al. [64] within the framework of the phase-field method “considered the perpendicular impingement of  $d_0 = 2.1$  mm water droplets at an impact velocity of  $u_0 = 0.62$  m/s on flat and structured hydrophobic Polydimethylsiloxane (PDMS) substrates. The droplet impingement on flat substrates resulted in a deposition with axisymmetric droplet deformation, while impingement on structured surfaces led to a complete rebound of the droplet, followed by an anisotropic wetting of the substrate and non-axisymmetric droplet deformation. The surface structure consisted of regular square grooves that have a width, height, and spacing of  $60\text{ }\mu\text{m}$  and matched the experiments [64].” from [DKS24, Section 2.2] The average numerical time step was  $\Delta t \approx 0.1\text{ms}$ . The coupled Cahn-Hilliard (eq. 2.29) and single-field Navier-Stokes equations (eq. 2.22) were solved on a uniform grid with  $10\text{ }\mu\text{m}$  cell size by the solver *phaseFieldFoam* in OpenFOAM®. Thereby, the initial droplet diameter was resolved by 210 mesh cells. “The resulting dataset is comprised of 53,244 synthetic images associated with 1,479 ground truth droplet shapes, of which 1,015 represent the non-axisymmetrical and 465 the axisymmetrical case.” from [DKS24, Section 2.2]

The 3D VoF simulation of an adhering droplet subjected to external shear flow conducted by Dreisbach *et al.* [DHK<sup>+</sup>25] “features a  $20\text{ }\mu\text{l}$  water droplet deposited on a PMMA substrate in a fully developed turbulent channel flow at  $Re_{\text{ch}} = 8,500$  [122, 26]. The average numerical time step is  $\Delta t \approx 1\text{ms}$  and a total of  $1.1\text{ s}$  are simulated. The droplet is discretized by approximately 20 cells in the vertical direction and 17 cells in the streamwise and spanwise direction.” from [DHK<sup>+</sup>25, Section 2.2] The single-field Navier-Stokes equations (eq. 2.22) and the transport equation for

the volume fraction (cp. eq. 2.21) containing an additional compression term, which ensures the conservation and boundedness of the volume fraction in the interface region, are solved by *hysteresisInterFoam* in OpenFOAM®. The simulation includes a boundary condition for the contact angle that accounts for hysteresis effects and, thereby, models the pinning of the droplet at the solid wall. The resulting dataset consists of 1,134 three-dimensional ground truth droplet geometries and respective 40,824 rendered images. From these three simulation cases, additional composite datasets are derived.

In the following, the nomenclature for these different datasets is introduced. The dataset  $\mathcal{D}$  featuring **D**roplet impingement on **F**lat substrates is referred to as  $\mathcal{D}_{\text{DF}}$  and, accordingly, the dataset for **D**roplet impingement on **S**tructured substrates is denoted as  $\mathcal{D}_{\text{DS}}$ . The combination of both datasets is referred to as  $\mathcal{D}_{\text{DFS}}$ . The dataset of adhering **D**roplets subjected to the external **S**hear flows is indicated as  $\mathcal{D}_{\text{DSH}}$  and the combination of all datasets is referred to as  $\mathcal{D}_{\text{DFSSh}}$ . All datasets were split by a ratio of 0.7/0.1/0.2 into separate training, validation, and testing subsets. A comprehensive overview of the dataset can be found in Table 5.1.

**Table 5.1:** Overview of the different training datasets used throughout this work.

dataset	no. 3D geometries	no. renderings
$\mathcal{D}_{\text{DF}}$	465	16,740
$\mathcal{D}_{\text{DS}}$	1,015	36,540
$\mathcal{D}_{\text{DFS}}$	1,479	53,244
$\mathcal{D}_{\text{DSH}}$	1,134	40,824
$\mathcal{D}_{\text{DFSSh}}$	2,613	94,068

## 5.5 Concluding remarks regarding synthetic training data generation

The high accuracy of the simulated glare-point locations qualifies the synthetically rendered images as suitable training data for the neural network, offering a physically accurate basis for learning the relation of the glare-point dynamics to the dynamics of the gas-liquid interface. The high visual similarity of the synthetic images to their real counterparts provides image data with appropriate features for training a convolutional neural network to extract feature from images recorded in the experiments. The quality of the image features will determine the generalization capability of the neural networks from synthetic image data to real image data. Consequently, a detailed evaluation of the neural networks during the application for the reconstruction of real data is required, as it will also yield an assessment of the quality of the synthetic training data. Further improvement of the synthetic data generation could be reached by domain adaption [40], which can be used to modify the already rendered synthetic data to move its feature distribution closer to that of the real images. Such an approach was successfully employed for the refinement of synthetic training data by Dreisbach *et al.* [DLP<sup>+</sup>22], leading to an enhanced performance of neural networks for object detection. In future investigations, higher-order glare points produced at different scattering could be used to encode further 3D information at additional points of the surface to further constrain the reconstruction. The first results for a modified setup at a scattering angle of 125° indicate that higher-order glare points can be rendered successfully for the dynamic deformation of the gas-liquid interface during impingement and that a functional neural network can be trained on that basis. The rendering setup can be adapted for the rendering of gaseous bubbles in liquids with little modification, enabling the synthetic data generation for such reconstruction efforts.




## 6 Spatio-temporal reconstruction of gas-liquid interfaces in droplet flows

In this chapter, the capabilities and limitations of the reconstruction framework are evaluated. Part of the results has previously been published in the articles *Spatio-temporal reconstruction of droplet impingement dynamics by means of color-coded glare points and deep learning* by Dreisbach *et al.* [DKS24] and *Interface reconstruction of adhering droplets for distortion correction using glare points and deep learning* by Dreisbach *et al.* [DHK<sup>+</sup>25]. Where findings have already been published in one of the referenced works, this is clearly marked and identified as related quotations.

### 6.1 Methodology

In this section, first, the fundamental procedure for the evaluation of the developed reconstruction framework is outlined, followed by a description of the necessary adaptations of the experimental setup for the investigation of different droplet phenomena. Subsequently, the pre-processing methods for the conditioning of the recorded image data prior to the reconstruction are described. Lastly, the evaluation metrics employed for the quantitative assessment of the reconstruction results are introduced.

#### 6.1.1 Evaluation procedure



**Figure 6.1:** Symmetries of the droplets gas-liquid interface in different experiments. Axisymmetrical deformation of the droplet during impingement on flat substrates (left), deformation within two planes of reflective symmetries during droplet impingement on structured substrates (middle), and deformation in one plane of reflective symmetry for adhering droplets, as well as droplet impingement on patterned substrates with gradients in wettability.

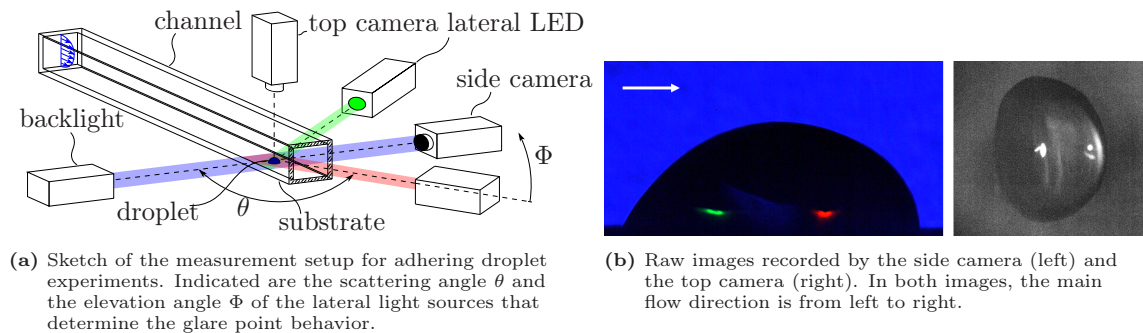
The fundamental approach of the following investigation is the application of the reconstruction framework to synthetic and experimental datasets that exhibit a gradually increasing complexity of the involved droplet phenomena, with the aim to progressively determine the capabilities and

limitations of the data-driven reconstruction method. First, the reconstruction framework is evaluated on different synthetic datasets in order to determine the capability of the neural network for learning the spatio-temporal dynamics of the various two-phase flow phenomena involving droplets. Moreover, the availability of a three-dimensional ground truth for the synthetic datasets allows for the quantitative evaluation of the reconstruction accuracy. The validation entails an investigation of the feature extraction in the convolutional neural network to uncover the relevance of glare points for the prediction of the three-dimensional droplet shape by the trained model of droplet dynamics. Afterward, the reconstruction framework is employed on experimental data in order to uncover the capability of the neural network, trained on synthetic data, to generalize to the real-world task. Initially, the reconstruction of the gas-liquid interface dynamics during droplet impingement on flat substrates is investigated. In this case, The deformation of the droplet is approximately axisymmetrical, as illustrated in Figure 6.1, aside from any perturbation in the experiments due to *e.g.* a minimally tilted surface or heterogeneity of the substrate due to inhomogeneous roughness or residues on the surface. Consequently, droplet impingement on flat substrates constitutes the most basic reconstruction task, for which in theory only the information on the shadowgraphy contour would be sufficient for the reconstruction, as the rotation of the contour yields the reconstructed three-dimensional shape of the interface. Therefore, this first step aims to validate the proposed approach of employing a neural network trained on synthetic data for the reconstruction of data obtained by experiments. The impact of the additional information on the three-dimensional interface shape encoded in the glare points is further elucidated through a comparative experimental study between versions of the neural network trained on synthetic data with and without glare points. The reconstruction framework is scrutinized on reconstruction tasks with gradually increasing difficulty, characterized by a higher degree of three-dimensional droplet deformation and a reduced symmetry of the interface in experiments involving droplet impingement on structured surfaces and adhering droplets in external shear flows. Specifically, the dynamic wetting of structured surfaces with regular grooves during droplet impingement results in the deformation of the gas-liquid interface within two planes of reflective symmetry, as illustrated in Figure 6.1. Conversely, the interface deformation of adhering droplets due to drag forces caused by the external flow occurs within one plane of symmetry. The degree of droplet deformation is further influenced by a variation of the experimental conditions, specifically, through the variation of the impact velocity and different substrates during droplet impingement, as well as different velocities of external flow for adhering droplet experiments. The Pixel-aligned implicit function (PIFu) neural network [191] is trained on the appropriate datasets for the respective reconstruction tasks and employed for the reconstruction of experiments with a varying deviation of the experimental conditions from the specific parameters of the training data in order to investigate the generalization capability of the trained models of droplet dynamics. In particular, the model trained for the reconstruction of adhering droplet dynamics is employed for the reconstruction of experiments that exhibit a significantly larger deformation of the gas-liquid interface in comparison to the training data, due to higher velocities of external flow. In the final step of the evaluation, the reconstruction framework is employed for experiments involving the impingement of droplets on patterned substrates with gradients in wettability, resulting in droplet dynamics that are significantly different from any of the training datasets. These experiments feature a higher degree of droplet deformation in comparison to the previous experiments and occur within only one plane of reflective symmetry. The generalization capability of the neural network to the resulting unknown interface shapes and the associated effects of different training datasets are evaluated in a comparative study.

### 6.1.2 Experimental setup

All experiments conducted in the present investigation employ the previously introduced glare-point shadowgraphy technique that is detailed in Chapter 4. The description of the experimental methods in this section is limited to details not entailed in the aforementioned chapter regarding the modification of the experimental apparatus for the different experiments. Specifically, the experiments involving the impingement of droplets on flat and structured solid substrates are conducted within the basic experimental setup introduced in Section 4.1. The experiments involving adhering droplets and droplet impingement on gradiented substrates feature an additional orthogonal camera view for validation purposes, which necessitates an extension of the experimental apparatus. Furthermore, the integration of a wind tunnel into the test rig, which is necessary for the experiments with adhering droplets subjected to external shear flows, is detailed in the following.

**Droplet impingement** “In order to evaluate the reconstruction method on surfaces with varying wetting properties, three solid substrates are considered: a flat silicon oxide (SiOx) surface and structured Polydimethylsiloxane (PDMS) and polylactide (PLA) substrates. The surface of the PDMS substrate forms regular square grooves with a height, width, and spacing of  $60\text{ }\mu\text{m}$ . The 3D-printed PLA substrate produced by fused deposition modeling (FDM) has circular arc-shaped ridges with a peak-to-peak spacing of  $154\text{ }\mu\text{m}$ . Images showing the structure of both substrates can be found in Figure C.3 in the Appendix alongside Table C.1, which lists the wettability properties of both substrates. The contact angles are determined through the detection of the droplet contour in the ten lowest rows of image pixels above the substrate, to which a linear fit is applied in order to obtain the angles relative to the substrate. In order to obtain the sub-pixel accurate contour of the droplet, first the horizontal intensity gradients are calculated by means of the Sobel operator, which are then fitted locally with a second-order polynomial. The uncertainty of the measured contact angles is estimated at  $\sigma_{\Theta} = 3.8^\circ$ , which is within the typical range for structured and/or imperfect smooth surfaces [167, 49, 64]. The impinging droplet is imaged by a *Photron Nova R2* equipped with a *Schneider-Kreuznach Apo-Componon 4.0/60* enlarging lens at a frame rate of  $f_r = 7,500$  frames per second (fps) and  $1,280\text{ px} \times 512\text{ px}$  resolution.” from [DKS24, Section 2.1] A sequence of images recorded for a water droplet impinging on the flat SiOx-surface is shown in Figure A.3 in the Appendix.




**Figure 6.2:** Experimental setup for the investigation of adhering droplets subjected to external shear flows (a) and resulting raw recordings (b). Figures adopted from [DHK<sup>+</sup>25].

**Adhering droplets in external shear flow** “The optical setup is adapted and integrated with a flow channel that allows for the investigation of adhering droplets in external shear flows at different fluid mechanical conditions [...] as shown in Figure 6.2a [...]. The shadowgraph images of the adhering droplet with glare points are recorded by a 36-bit color dynamic range *Photron*

*Nova R2* equipped with a *Nikon AF Micro Nikkor 2.8/105* macro lens at  $f_r = 7,500$  frames per second (fps) and  $1,280 \text{ px} \times 512 \text{ px}$  resolution. Moreover, a second orthogonal viewpoint from the top is recorded to obtain reference images for the evaluation of the out-of-plane accuracy of the reconstruction. For that purpose, a 12-bit monochrome dynamic range *Photron Fastcam SA4* equipped with a *Zeiss Milvus 2/100M* macro lens is used, which records at  $f_r = 3,000$  fps and  $1,024 \text{ px} \times 1,024 \text{ px}$ .

A flow channel with a  $22 \text{ mm} \times 22 \text{ mm}$  cross-section and  $1010 \text{ mm}$  length, made of acrylic glass (PMMA), is used in order to provide full optical access. The water droplets are placed at the downstream end of the channel, where a fully developed laminar or turbulent flow can be ensured. The volume flow rate of the air flow is measured by a *Testo 6451* compressed air meter and adjusted by a pressure valve. In order to investigate the influence of different degrees of droplet motion, the bulk velocity of the air flow is varied in three discrete steps at  $u_B = 5.85 \text{ m/s}$ ,  $u_B = 7.58 \text{ m/s}$ , and  $u_B = 8.32 \text{ m/s}$ , while the droplet volume is varied between  $5$  and  $22 \text{ }\mu\text{l}$ . The corresponding Reynolds numbers, with respect to the bulk velocity and the hydraulic diameter of the channel, lie in the range of  $Re_{ch} = 8,500 - 12,100$ .


The substrate on which the droplets are placed consists of PMMA, which is moderately hydrophilic with an advancing contact angle of  $85^\circ$  and a receding contact angle of  $44^\circ$ , measured by the tilting method [145]. Note that dynamic contact angles lower than  $90^\circ$  ensure that the top view on the contact line is never self-obstructed by the droplet's gas-liquid interface and thus a continuous tracking of the wetted area is possible. However, the low dynamic contact angles necessitate an adjustment of the lateral light sources to an elevation angle of  $60^\circ$  in order to ensure the occurrence of glare points on the gas-liquid interface in all frames." from [DHK<sup>+</sup>25, Section 2.1]




**Figure 6.3:** Sketch of the measurement setup for experiments involving droplet impingement on patterned substrates with gradients in wettability. Indicated are the scattering angle  $\theta$  and the elevation angle  $\Phi$ .

**Droplet impingement on gradiented substrates** In order to facilitate the accurate observation of droplet dynamics involving moving contact lines, the experimental setup is adapted for the simultaneous imaging of orthogonal bottom and side views in the shadowgraphy technique with one RGB camera. The modified setup is illustrated in Figure 6.3. Experiments with transparent substrates allow for the imaging of the droplet dynamics from below in a backlighting setup, which provides a higher contrast in the images in comparison to the previous reflected light setup and, thus, a more accurate mapping of the droplet contour. A red LED is used to illuminate the droplet from the top to generate a second bottom-view shadowgraph, while the previously red lateral illumination source is replaced by an additional green LED. This second shadowgraphy view is redirected towards the camera by a system of mirrors. The light paths from the bottom view and

the side view are aligned through a dichroic mirror that reflects high wavelengths of red light and transmits lower wavelengths of blue and green light, thereby allowing for the combination of all light paths. Both lateral illumination sources are realized by green LED lights that produce glare points on the gas-liquid interface that can be transmitted through the dichroic mirror. The lateral illumination is set to a scattering angle of  $96.3^\circ$  to produce only  $p_{gp} = 0$  glare points. Consequently, the blue image channel records the shadowgraph contour of the droplet in the side view, unchanged from the basic setup, the green channel records the glare points resulting from both lateral light sources, while the red image channel is reserved for the validation image in the bottom view. The simultaneous imaging with one camera facilitates an exact temporal correlation of the bottom and side views, which simplifies the evaluation of the image data involving fast-moving dynamic contact lines during droplet impingement.



**Figure 6.4:** Side view of the gas-liquid interface at  $t = 1.1$  ms, 8.0 ms, and 19.1 ms after droplet impingement on the *SMG2*-substrate.




**Figure 6.5:** Bottom view of the gas-liquid interface at  $t = 1.1$  ms, 8.0 ms, and 19.1 ms after droplet impingement on the *SMG2*-substrate.

The modified setup is employed for experiments involving the impingement of water droplets on patterned surfaces with gradients in wettability. The patterned substrates are produced by casting PDMS in negative molds obtained from arrested breath figures in polymer glues through the process described by Dent [50]. The result of this manufacturing process are patterned surface structures constituted by densely packed hemispheres with around  $d_p = 10 \mu\text{m}$  diameter that transition into a flat surface, thereby producing a smooth and continuous gradient in the spacing, as well as size of the structures. Different substrates with significant variation in the length of the transition region from the patterned structures to the flat surface are produced. The estimated gradient length ranges from  $l_g \approx 100 \mu\text{m}$  to a few millimeters. The substrates with spatially modulated gradients are denoted *SMG0* to *SMG4* in the following. Example images of these substrates obtained by confocal microscopy are shown in Appendix Figure C.4. Further characterization of the substrates is beyond the scope of this work, however, the reader is referred to the thesis of Dent [50] for further details.

The substrates are aligned in the experimental apparatus, so that the droplets impinge in the middle of the transition region, with the surface gradient oriented parallel to the side view of

the imaging setup. A sequence of images recorded in the side and bottom views during the impingement of a  $d_0 = 2.4\text{ mm}$  water droplet at  $u_0 = 0.84\text{ m/s}$  on the *SMG2* patterned substrate is shown in Figures 6.4 and 6.5, respectively. In the bottom view the transition from the patterned structures to the flat region, indicated by low and high background illuminance, respectively, becomes obvious. The surface patterning enhances the hydrophobicity of the PDMS substrate, as evident by the difference of the contact angles visible in Figure 6.4 for the structured and flat regions of the substrate. Consequently, a gradient in wettability emerges over the length of the transition region, whose magnitude is inversely related to the length of the transition region. Due to the gradient in wettability, an additional horizontal force acts on the impinging droplets, which in turn induces a motion directed toward the flat region of the substrate. A similar behavior during droplet impingement on chemically heterogeneous substrates was previously observed by Mock et al. [153]. The horizontal displacement  $\Delta_s$  of the droplet from the point of impact to the resting position at equilibrium is measured from the detected contours in the side view shadowgraph and is detailed in Appendix Table C.2. Thereby, a maximum displacement of  $\Delta_s = 1.8\text{ mm}$  is found for droplet impingement on the *SMG2* substrate. This horizontal motion of the droplet is superposed with the vertical motion of the falling droplet, which leads to a complicated three-dimensional deformation of the gas-liquid interface. As can be seen in Figure 6.4, the deformation of the droplet differs significantly from the previous experiments involving droplet impingement on flat or grooved substrates.

### 6.1.3 Data processing and network training



**Figure 6.6:** Image of a droplet that is deposited on a structured substrate (left), binary mask (middle), and masked image (right). Figure adopted from [DKS24].

The PIFu neural network [191] is trained for the physically correct reconstruction of the gas-liquid interface on the labeled datasets comprised of synthetically rendered images, produced by physically-based rendering as described in Section 5.2, and ground truth shapes of the three-dimensional gas-liquid interface.

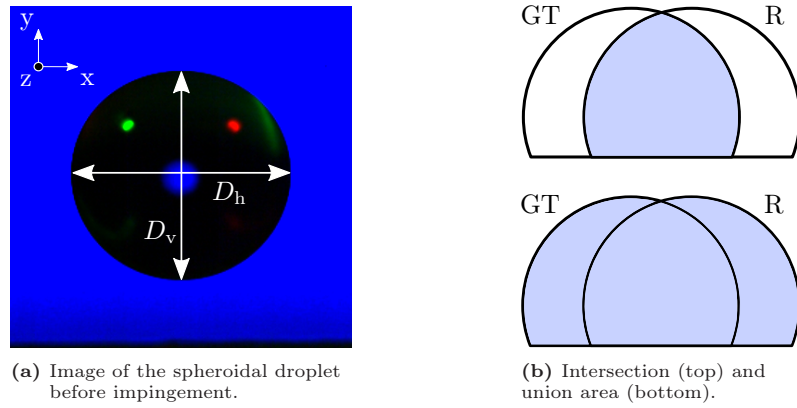
“The training images are pre-processed by the superimposition of a binary mask that covers the substrate ground in the images. This introduction of prior knowledge simplifies the reconstruction task, as the differentiation between the shadowgraph contour of the droplet and the liquid-solid contact area does not need to be learned by the neural network. In order to allow for the reconstruction of different droplet volumes, the input images and the corresponding three-dimensional ground truth shapes are scaled randomly, which is a common data augmentation technique to introduce scale invariance to neural networks [206]. The PIFu neural network is trained by the RMSProp optimization algorithm [225] for eight epochs with hyperparameters according to [...] [Section 3.4], in particular a batch size of 12 and learning rate of 0.001, which is reduced by a factor of ten at the beginning of epochs six and eight.

The images obtained in the experiments are pre-processed prior to the reconstruction in order to further enhance the similarity to the synthetic images, which is crucial for an optimal generalization



of the neural network between both image data domains [40, 207]. The mutual perturbation in the color channels of the RGB image that resulted from cross-talk in the camera sensor and the polychromatic light of the LEDs is compensated by the [ISC] correction method introduced [in Section 4.2] [...]. Thereby, the colored glare points and the shadowgraph contour are separated into the respective image channels for each light source color. In the second step, the substrate ground is masked analogously to the training data.” from [DHK<sup>+</sup>25, Section 2.3] “In order to allow for the distinction between the contours of the droplet and substrate, first the outline of the substrate is determined from an initial image, in which the droplet is not yet in frame. Afterward, the substrate contour is removed from all later masks, and thus the information of the solid-liquid interface is imposed on each frame, as evident from Figure 6.6 (middle). Subsequently, the input images are overlayed with their respective mask, as can be seen in Figure 6.6 (right), and inputted to the neural network in order to obtain the frame-wise volumetric reconstruction of the droplet during impingement.” from [DKS24, Section 2.4]

#### 6.1.4 Evaluation metrics



**Figure 6.7:** (a) Image of a spheroidal droplet before impact with markers indicating the horizontal  $D_h$  and vertical semi-axis  $D_v$  of the ellipse fitted in the image plane; (b) intersection and union areas of the reconstruction  $R$  and ground truth  $GT$  in two-dimensional space. Figure adopted from [DKS24].

“The performance of the neural network is evaluated considering the reconstructed three-dimensional shapes and the availability of ground truth data. The following metrics are used for the evaluation:

- The three-dimensional intersection over union

$$3D\text{-IOU} = \frac{R \cap GT}{R \cup GT} \quad (6.1)$$

is calculated as the fraction of the intersection volume between the reconstruction  $R$  and ground truth  $GT$  and the union volume of  $R$  and  $GT$ . The 3D-IOU is an extension of 2D-IOU [61] (see Figure 6.7b) to three dimensions and, therefore, provides a measure for the spatial volumetric accuracy of the reconstruction in 3D space.

- The bias error of the reconstructed volume  $\delta_V$  is calculated by the absolute deviation of the mean volume

$$\bar{V} = \frac{1}{n} \sum_{i=1}^n V_{R,i} \quad (6.2)$$

of the reconstructed shapes  $V_R$  from the ground truth volume  $V_{GT}$ , and given relative to the ground truth volume [11]

$$\delta_V = \left| \frac{V_{GT} - \bar{V}}{V_{GT}} \right|. \quad (6.3)$$

- The measured uncertainty of the reconstructed volume  $\sigma_V$  is calculated by the standard deviation of the deviation between the reconstructed volume and the ground truth volume, and given relative to the ground truth volume [11]

$$\sigma_V = \frac{1}{V_{GT}} \sqrt{\frac{1}{n-1} \sum_{i=1}^n (V_{R,i} - \bar{V})^2}. \quad (6.4)$$

The reconstruction of the synthetic validation dataset is evaluated by means of the 3D-IOU. The results can be interpreted as the baseline performance of the neural network, that can be reached for a perfect agreement of synthetic training data with real experimental data.

Since no 3D ground truth data is available for experiments, the reconstruction of experimental images is evaluated by means of the volumetric bias error  $\delta_V$  and uncertainty  $\sigma_V$  of the reconstruction relative to the integral ground truth volume  $V_{GT}$ . The ground truth volume in the [droplet impingement] experiments can be estimated from the shadowgraph contour of the droplet before its impact on the surface. For this, the shape of the droplet is assumed to be spheroidal with axisymmetry around the vertical axis, and consequently, the volume can be calculated as  $V_{GT} = \frac{\pi}{6} D_s^3$  with the equivalent spherical diameter  $D_s = \sqrt[3]{D_h^2 D_v}$ .  $D_h$  and  $D_v$  are the horizontal and the vertical semi-axis of the falling droplet, as indicated in Figure 6.7a that are determined by an ellipse fit to the shadowgraph contour with the method of Taubin [220]. The ground truth volume is averaged over all frames  $n_f$  of the undeformed droplet before impact (usually  $n_f = 11$  frames), in order to reach a high accuracy of the estimation for the ground truth volume. The uncertainty of the ground truth volume, measured by the standard deviation over all consecutive frames amounts to  $\sigma_{V,GT} = 0.06\%$  of  $V_{GT}$  and is therefore negligible. The impact velocity  $u_0$  was determined by tracking vertical motion  $\Delta s$  of the center point of the fitted ellipses over  $n_f$  frames in relation to the time difference  $\Delta t = n_f / f_r$  between the initial and final frame, which produces in the linear fit  $u_0 = \Delta s / \Delta t = f_r \Delta s / n_f$ . The uncertainty of the velocity measurement is estimated to be below  $\sigma_u = 0.006$  m/s. Furthermore, the reconstruction is qualitatively evaluated by the comparison of the contours of the reconstructed droplet shapes and the shadowgraph images.” from [DKS24, Section 2.5] The estimation of the ground truth volume for the experiments involving adhering droplets subjected to external shear flows was found unfeasible due to the introduction of significant errors from the oscillation of the droplet contour. Consequently, the measured uncertainty in those experiments is normalized by the arithmetic mean  $\bar{V}$  of the reconstructed volumes instead of the estimated ground truth volumes.

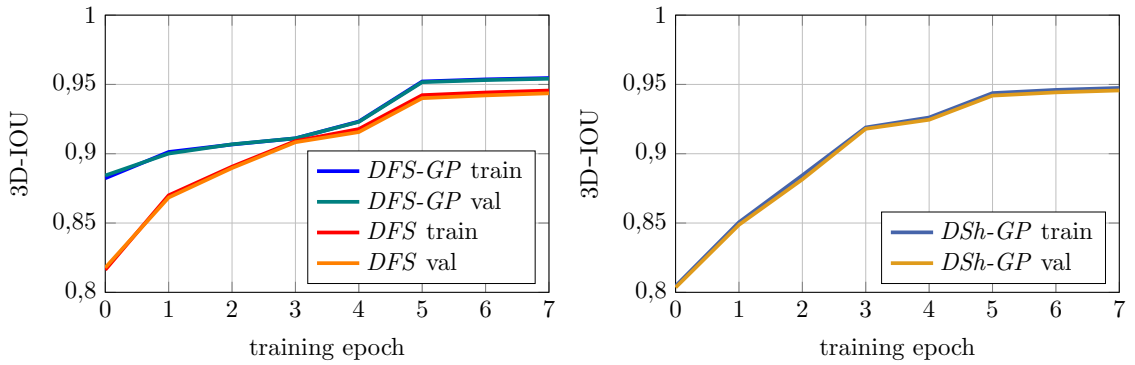
## 6.2 Validation

As elaborated in Section 6.1.1, the feasibility of the proposed framework for the data-driven reconstruction of gas-liquid interfaces in droplet flows is first evaluated on the grounds of synthetic images. Subsequently, the significance of glare points for the reconstruction process is investigated by an in-depth analysis of the feature extraction in the convolutional neural network.



### 6.2.1 Training and validation on synthetic image data

“In the following, the trained state of the PIFu neural network for the reconstruction [of the impinging droplets gas-liquid interface from glare-point shadowgraphy] [...] images is referred to as *DFS-GP* (Droplet impingement on Flat and Structured substrates with Glare Points). Additionally, a second synthetic image dataset without glare points, and therefore only featuring the shadowgraph contour, was generated from the same ground truth data. A second benchmark version of the neural network was trained on the dataset without glare points and is accordingly referred to as *DFS* in the following. Both datasets were split in a ratio of 0.7/0.1/0.2 into a training dataset and separate validation and testing datasets that are not used directly for the training of the network and can therefore be used to evaluate the reconstruction accuracy on unknown data. During training the reconstruction performance was evaluated by the 3D-IOU (Equation 6.1) on the training and validation datasets.



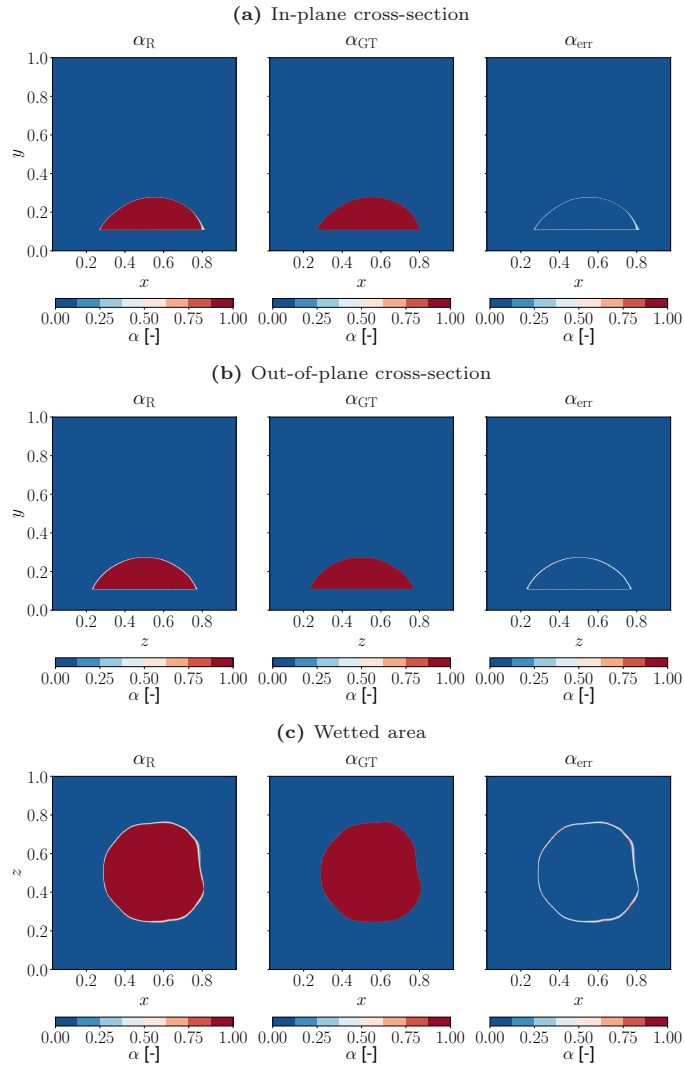
(a) 3D-IOU during training of the *DFS-GP* and *DFS* networks. Figure adopted from [DKS24]. (b) 3D-IOU during training of the *DSh-GP* network. Figure adapted from [DHK<sup>+</sup>25].

**Figure 6.8:** Average 3D-IOU during training on the  $\mathcal{D}_{\text{DFS}}$  (left) and  $\mathcal{D}_{\text{DSh}}$  (right) synthetic datasets, calculated by the arithmetic mean over 1,000 samples of the training and validation dataset at every epoch.

As can be seen in Figure 6.8 the reconstruction accuracy improves over time on both the training and validation data set, which indicates a successful learning of the neural network. In particular, the strictly monotonically increasing validation performance indicates that the network did not overfit the training data, but rather that a generalization capability to unknown data was maintained. However, the training and validation performance is very similar, which likely results from both data stemming from the same distribution with a fine temporal resolution, so that the validation samples were very similar to the training samples and thus easy to predict for the neural network. Precisely, *DFS-GP* reaches a performance of 3D-IOU=0.954 on the validation dataset at the end of the training and 3D-IOU=0.955 on the training dataset, while *DFS* reaches 3D-IOU=0.944 and 3D-IOU=0.946, respectively [see Figure Figure 6.8a]. Therefore, the reconstruction accuracy appears to be improved with glare points. In particular, considering that both models reach close to a perfect reconstruction result (3D-IOU=1) the difference is significant with a 1.05% lower error for *DFS-GP*. Consequently, the results on synthetic data suggest that glare points facilitate the learning of the three-dimensional droplet shape.” from [DKS24, Section 3.2]

Analogously, the 3D-IOU was tracked during the training of the *DSh-GP* neural network on the  $\mathcal{D}_{\text{DSh}}$  dataset for the reconstruction of the adhering droplets’ gas-liquid interface dynamics. “As can be seen in Figure 6.8b the neural network converges to a high volumetric accuracy on both the training and the validation dataset towards the end of the training (3D-IOU<sub>train</sub> = 0.947 and 3D-IOU<sub>val</sub> = 0.946), which is close to a perfect agreement with the ground truth (3D-IOU<sub>ideal</sub> = 1).

The comparative analysis of all validation samples reveals that the accurate volumetric shape was reconstructed consistently throughout the dynamic deformation and oscillation of the droplet, as indicated by a low standard deviation of  $\sigma_{3D-I\text{OU}} = 0.027$ . The median of the volumetric accuracy on the validation dataset at the final training epoch was notably higher than the arithmetic mean ( $3D-I\text{OU}_{\text{median}} = 0.954$  and  $3D-I\text{OU}_{\text{val}} = 0.946$ ), due to outliers in the first five time steps of the validation dataset, which had a minimal accuracy of  $3D-I\text{OU}_{\text{min}} = 0.793$ . The shape of these outliers deviated significantly from the rest of the data distribution, thus explaining the lower reconstruction accuracy. The overall high spatial accuracy of the prediction for the validation dataset, however, implies a good generalization to unseen samples.” from [DHK<sup>+</sup>25, Section 3.1] To conclude, the results indicate that the PIFu neural network was able to successfully learn a spatio-temporal representation of the gas-liquid interface dynamics for both impinging and adhering droplets from the synthetic training data.



**Figure 6.9:** Cross-sections of the volume fraction  $\alpha$ ; prediction by the neural network (left), ground truth (middle), and deviation of the prediction from the ground truth (right). In subfigures (a) and (c) the main flow direction is from left to right and in subfigure (b) the main flow direction is aligned to the image plane. Figure adopted from [DHK<sup>+</sup>25].

“Figure 6.9 shows the prediction of the phase distribution for a sample from the  $[\mathcal{D}_{\text{DSH}}]$  validation dataset represented by the volume fraction  $\alpha$  in three orthogonal cross-sections of the reconstructed

volume. The reconstruction results for a representative sample of the  $[\mathcal{D}_{\text{DSh}}]$  validation set are shown in comparison to the ground truth shape and the respective difference of both. [...]

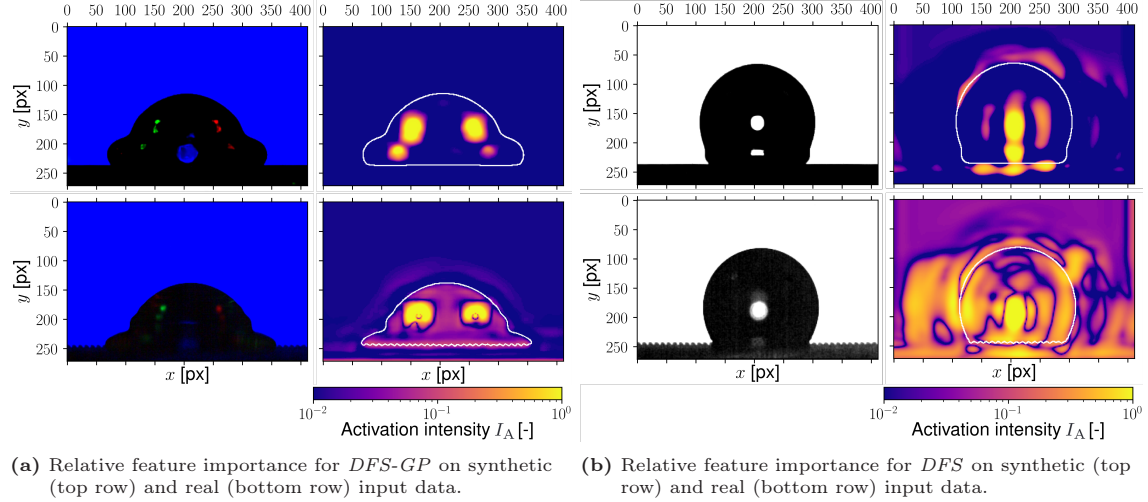
As can be seen, there is a good agreement with the ground truth in the global shape of the contour, as well as the local curvature of the reconstructed droplet shapes. However, the gas-liquid interface of the reconstruction is more diffuse than the ground truth, which is a result of the continuous implicit representation of the droplet by the neural network. In particular, the out-of-plane component of the reconstruction is more diffuse than the in-plane component, as can be seen by the comparison of Figures 6.9a and 6.9b. Furthermore, the out-of-plane component of the droplet was found to be less accurately reconstructed, which was more obvious for high error samples (see Figure C.2 in the Appendix). The more diffuse gas-liquid interface, as well as the higher deviation in the shape of the contour, indicate a higher uncertainty of the network’s prediction in the out-of-plane direction. This result was expected, as fewer image features, *i.e.* only the glare points, are available for the reconstruction of the out-of-plane component in comparison to the in-plane component, which has a strong basis for the reconstruction by the shadowgraph contour. Therefore, the network has to rely more on the learned knowledge from the training data for the out-of-plane reconstruction.

As can be seen in Figure 6.9 the streamwise cross-section of the droplet is nearly planar-symmetrical, while the in-plane cross-section deviates significantly from planar symmetry due to the deformation of the droplet by the external flow. Consequently, the in-plane geometry is more difficult to reconstruct, which is alleviated by the surplus of image features available for the reconstruction. As demonstrated by the reconstruction results, both the planar symmetry of the streamwise cross-section and the deformed contour in the lateral view of the droplet are reconstructed accurately, which makes the proposed approach well-suited for the task of adhering droplet reconstruction.” from [DHK<sup>+</sup>25, Section 3.1]

## 6.2.2 Feature visualization

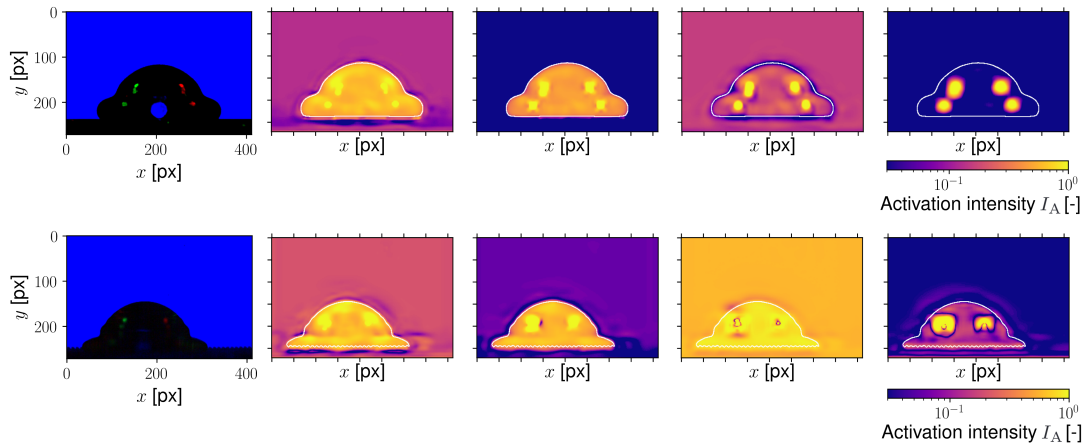
“In order to gain further insight into the relevance of glare points for the prediction of the three-dimensional gas-liquid interface, intermediate results from feature extraction are analyzed in the following. As the features extracted by the hourglass network directly serve as the input for the occupancy prediction, they can be used to determine the meaningful regions in the input image for the neural network. A popular method for the visualization of the feature relevance is based on the successive occlusion of different regions in the input image, which allows measuring their importance by the drop in network performance. Following the score-based class activation map (Score-CAM) approach introduced by Wang et al. [240], the feature maps obtained by the final convolutional layer of the hourglass network are used to occlude the input images. The weighted sum of the feature maps, which are scored according to their influence on the prediction, gives the activation map that highlights the most determining regions for prediction in the input.

Figure 6.10a compares the activation maps of the hourglass network in *DFS-GP*, obtained by processing samples of synthetic images and images recorded in the experiments, while Figure 6.10b contrasts the activation of *DFS*. The normalized intensity of the activation maps is plotted on a logarithmic scale in order to highlight the different orders of magnitude in the feature importance of the glare points, the shadowgraph, and the background. The comparison reveals that both models predominantly focus on the glare points during feature extraction, as indicated by a high activation. Furthermore, both models have a similar response to the synthetic and real image data, which indicates that the synthetic data generation already reaches a sufficiently high quality in terms of representing the real data distribution. In particular, *DFS-GP* concentrates on the lateral



**Figure 6.10:** Score-CAM feature importance for *DFS-GP* and *DFS* on synthetic and real data. The raw images are shown on the left side of each subfigure and the normalized intensity of the activation map  $I_A$  is plotted on a logarithmic scale and overlaid with the contour line of the shadowgraph, indicated by a white solid line on the right side. Figure adopted from [DKS24].

glare points and almost completely disregards the available spatial information from the central glare point, while *DFS* concentrates on the central glare point. Both models consider the contour and bulk of the droplet to a lesser degree, with *DFS* relying more on the contour information than *DFS-GP*. Also, both models put more emphasis on the contour for real image data – particularly on frames that are different from the training data – due to dim or diffused glare points or unknown droplet shapes. Furthermore, the reliance on additional features and the focus on the contour might be related to an increased uncertainty of the neural network due to noise and aberration in the recordings from the experiment. This indicates that additional information is extracted from the images for the volumetric reconstruction of high-uncertainty samples. The feature visualization reveals that glare points effectively enhance reconstruction by encoding 3D information of the gas-liquid interface, and the PIFu network can be trained to utilize this information. Furthermore, the network might be highly adaptable towards available information and is able to exploit multiple different features for prediction.” from [DKS24, Section 3.3]



**Figure 6.11:** Relative feature importance for *DFS-GP* in the intermittent levels ( $l_{HG,0}$  to  $l_{HG,2}$ ) and the final level ( $l_{HG,3}$ ) of the hourglass network on synthetic (top row) and real (bottom row) input data.

In order to further investigate the processing of the information in the image data during feature extraction, the activation maps at intermittent steps are compared. As discussed in Section 3.4, the

hourglass network consists of four stacked hourglass modules ( $l_{\text{HG},0}$ ,  $l_{\text{HG},1}$ ,  $l_{\text{HG},2}$ ,  $l_{\text{HG},3}$ ) that each produce a feature map, which is used to predict the three-dimensional gas-liquid interface during training for intermittent supervision of the optimization. The intermittent supervision causes each hourglass level to learn a feature map that can be used for prediction, with each level becoming increasingly accurate in comparison to the previous levels. During inference, only the prediction on the basis of the final feature map ( $l_{\text{HG},3}$ ) is used. Figure 6.11 shows the activation maps for the intermittent hourglass levels in *DFS-GP* obtained for synthetic and real input data. As can be seen, the neural network relies on the combined information of the shadowgraph contour and glare points in the early levels ( $l_{\text{HG},0}$  and  $l_{\text{HG},1}$ ), while the later levels ( $l_{\text{HG},2}$  and  $l_{\text{HG},3}$ ) increasingly rely on the glare points. This effect can be observed in the processing of both synthetic and real data, as well as in both versions of the network. Furthermore, in the earlier levels, a larger number of glare points are considered, including less intense glare points, which are neglected in the later levels. The limitation on fewer features in later levels indicates the removal of redundant information, as the information in the input is increasingly propagated throughout the feature maps, due to the growing receptive field in later levels. This progressive refinement of the features throughout all levels of the hourglass network reveals that this deep neural network architecture is effective for feature extraction on the glare-point shadowgraphy images. This allows the PIFu network to leverage the encoded information of the three-dimensional gas-liquid interface in the glare points for the reconstruction of the droplet dynamics.

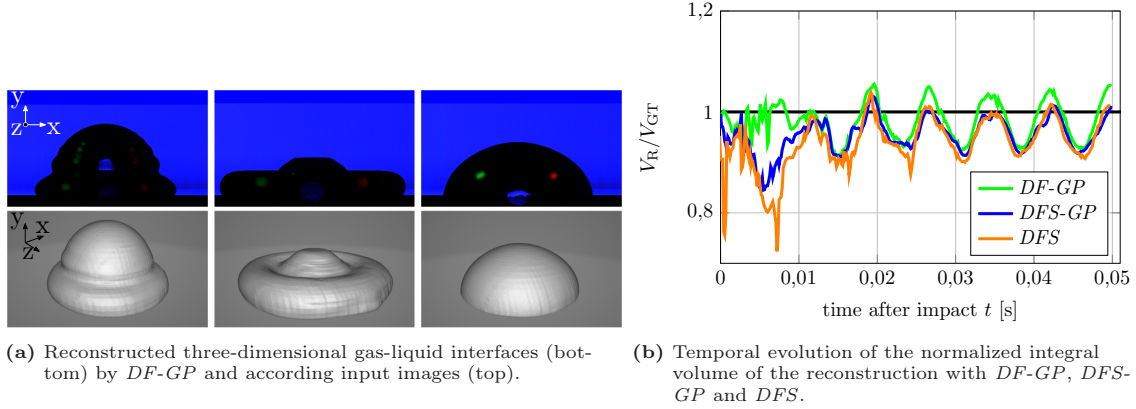
## 6.3 Results and Discussion

In the following section, the results for the spatio-temporal reconstruction of gas-liquid interfaces on the basis of images recorded in experiments using the glare-point shadowgraphy technique are evaluated. As discussed in Section 6.1.1, the capabilities and limitations of the data-driven reconstruction framework are elucidated by experiments involving droplet impingement on various solid substrates, as well as adhering droplets subjected to external shear flows.

### 6.3.1 Axisymmetrical droplet impingement

“The experiments were conducted with the test rig described in Subsection 6.1.2 and the recorded images were pre-processed according to Subsection 6.1.3. Furthermore, prior to the reconstruction, all images were cropped uniformly in order to reach a high resolution in the reconstruction, while conserving the scale between time steps. The experiments covered the impingement of water droplets with an equivalent diameter of  $d_0 = 2.08$  mm onto a flat hydrophilic SiOx surface with an impact velocity of  $u_0 = 0.7$  m/s. The outcome of the impact was a deposition of the droplet with axisymmetrical deformation. Both states of the neural network trained on simulation data with anisotropic wetting [ $\mathcal{D}_{\text{DFS}}$ ], *DFS-GP* and *DFS*, were employed. Additionally, a benchmark model *DF-GP* was trained solely on the part of the training dataset [ $\mathcal{D}_{\text{DF}}$ ] with glare points that comprised isotropic wetting of the flat substrate and therefore had strictly axisymmetrical deformation.

In Figure 6.12a three exemplary three-dimensional gas-liquid interface geometries reconstructed by the *DF-GP* network are illustrated with their respective input images. As can be seen, the network was able to infer physically reasonable shapes from the input images that align well with the contour of the input image. Furthermore, the reconstruction results are axisymmetric on the global scale, thus revealing that the rotational symmetry that is characteristic for a droplet impingement on flat substrates was learned well by the network. The results of the *DFS-GP* and



**Figure 6.12:** Reconstruction results for droplet impingement on the flat SiOx substrate. Figures adopted from [DKS24].

$DF-GP$  networks deviate slightly more from axisymmetry (not shown here) than those of  $DF-GP$ , which indicates an influence of the training data on the learning of symmetries.

These results demonstrate the successful application of the neural network trained on synthetic data to the real-world task and, consequently, validate the general approach of using synthetic data for the optimization of the neural network. Furthermore, this confirms the earlier conjecture that the synthetically generated training data sufficiently represents the real data distribution. Further improvement of the reconstruction accuracy could be reached by domain adaption [40], which is aimed to modify the already rendered synthetic data to move its feature distribution closer to that of the real images.

Figure 6.12b shows the temporal evolution of the integral volume  $V_R$  of the reconstruction relative to the ground truth volume measured from the images in the experiment  $V_{GT}$ , as described in Subsection 6.1.4. It can be observed that all versions of the PIFu network were able to reconstruct the volume of the gas-liquid interface in good agreement with the experiment.


**Table 6.1:** Uncertainty  $\sigma_V$  and bias error  $\delta_V$  of the reconstructed integral volume for  $DFS-GP$  and  $DFS$  in percent of ground truth volume  $V_{GT}$ . Table adopted from [DKS24].

network	$DF-GP$	$DFS-GP$	$DFS$
$\sigma_V$ [%]	3.7	3.5	5.1
$\delta_V$ [%]	1.8	4.7	6.2

In particular, the  $DF-GP$  network reached a close agreement indicated by a low bias error of  $\delta_V = 1.8\%$ . In comparison, the versions trained on datasets including non-axisymmetrical wetting produced higher bias errors in the reconstruction, as can be seen in Table 6.1. The uncertainty of the reconstructed volume is similar across all versions of the network. Overall the networks trained on image data with glare points reached a lower bias error and uncertainty in the reconstruction in comparison to the model trained without glare points, even though the droplet dynamics were axisymmetrical and thus sufficiently represented by the shadowgraph contour. Furthermore, it was found that the error of the reconstructed volume is closely related to the oscillation of the droplet after impact, which can directly be observed in the oscillation of the reconstructed volume.

In Figure 6.13 the reconstruction result for an image frame with self-occlusion of the gas-liquid interface is shown. The droplet contour in the shadowgraph image appears to have a sharp corner between the higher central region and the lamella surrounding it, however in reality the gas-liquid interface is smooth and there is a trough between the outer rim and the central peak [185] that






**Figure 6.13:** Reconstruction of a frame with self-occlusion of the gas-liquid interface (right) and respective input image (left) overlaid with the contour of the reconstruction. Figure adopted from [DKS24].

is hidden in the 2D-projection. The white line overlaid with the input image on the left indicates the gas-liquid interface of the 3D reconstruction for the cross-section in the image plane in the middle of the droplet. As can be seen, the smooth contour of the gas-liquid interface, as well as the large curvature that should occur are successfully estimated by the neural network in a physically correct manner for the regions that are obscured in the input image. The successful reconstruction of occluded regions within the three-dimensional shape reveals the strength of the deep learning approach to interpolate large unknown regions while respecting the underlying physics contained in the numerical training data.” from [DKS24, Section 3.4]

### 6.3.2 Non-axisymmetrical droplet impingement




**Figure 6.14:** Contours of the shadowgraph images for deposited droplets on the structured PLA and PDMS substrates at parallel (||) and transversal (⊥) viewing orientation. Figure adapted from [DKS24].

“In the following subsection the volumetric reconstruction on the basis of experiments of water droplet impingement on two different structured substrates, in particular hydrophilic polylactide (PLA) and hydrophobic Polydimethylsiloxane (PDMS), is presented. The anisotropic wetting of structured surfaces leads to non-axisymmetrical droplet deformation during impingement, with a larger spreading ratio and lower contact angles in the direction parallel to the grooves in comparison to the transversal direction. The resulting static shape of the droplet after impact on the PLA and PDMS substrates in parallel and transversal direction is illustrated in Figure 6.14. Due to the non-axisymmetrical deformation of the gas-liquid interface, the information from only a shadowgraph projection is not sufficient for a volumetric reconstruction. Consequently, the additional three-dimensional information of the gas-liquid interface encoded in the glare points has to be exploited by the neural network to achieve an accurate reconstruction. The results of the reconstruction for *DFS-GP* and *DFS* are compared in order to evaluate the effectiveness of glare points for the reconstruction of asymmetric droplet deformation. Different observation angles, *i.e.* the angle between the orientation of structures and the camera axis, are considered and the results of the volumetric reconstruction are compared in order to determine the influence of the observation angle



on the reconstruction accuracy. In the following for both substrates an observation angle of  $\alpha = 0^\circ$  denotes parallel alignment of the camera with substrate features (*e.g.* grooves), while  $\alpha = 90^\circ$  denotes a perpendicular alignment.

In the first series of experiments the droplet impingement on a hydrophilic 3D-printed substrate, that was produced from polylactide (PLA) by Fused Deposition Modeling (FDM), is reconstructed. A black filament was chosen in order to minimize the reflection of the lateral illumination on the substrate that would result in unwanted additional glare points, which, in turn, could disturb the reconstruction. The pattern of the stacked layers resulting from the 3D-printing process is horizontally aligned with the substrate surface, forming a wave pattern with a characteristic length of  $154\text{ }\mu\text{m}$ . The experiments featured the impingement of water droplets with an equivalent diameter of  $d_0 = 2.27\text{ mm}$  at an impact velocity of  $u_0 = 0.45\text{ m/s}$ , that was recorded at orientation angles of  $0^\circ$ ,  $45^\circ$  and  $90^\circ$ . In the experiments an equilibrium contact angle of  $\Theta_{\text{eq,p}} = 76^\circ$  in the parallel direction and  $\Theta_{\text{eq,t}} = 63^\circ$  in the transversal direction was measured. The dynamic contact angles for both directions are found in Table C.1 in the Appendix.

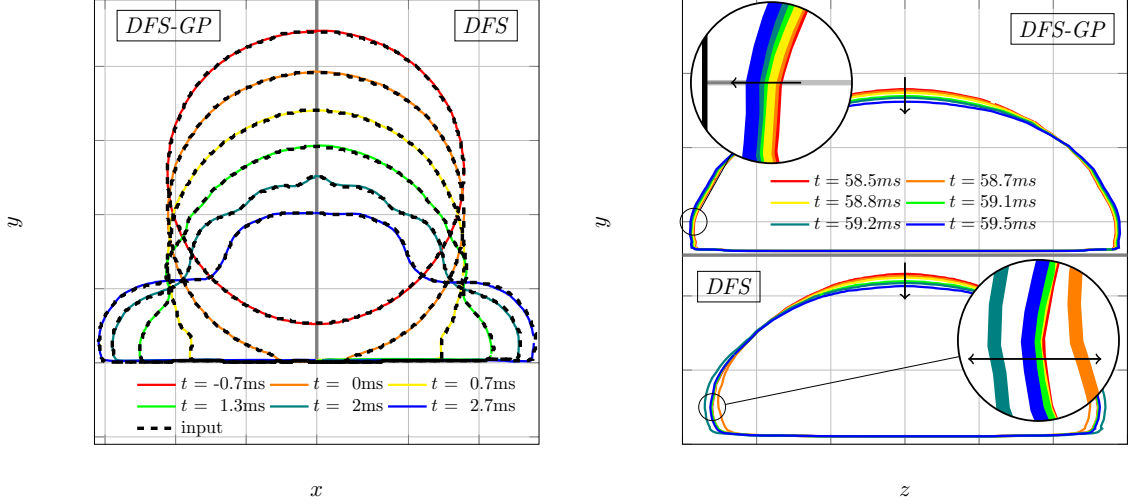


**Figure 6.15:** Temporal evolution of the normalized integral volume of the reconstruction for droplet impingement on the structured PLA substrate at  $\alpha = 45^\circ$  observation angle with *DFS-GP* and *DFS*. Figure adopted from [DKS24].

The results of the volumetric reconstruction from experimental images of droplet impingement on the PLA substrate indicate that both states of the network trained with and without glare points successfully reconstructed the non-axisymmetrical droplet deformation. *DFS-GP* reached an uncertainty of 5.7% and a bias error of 4.9% averaged over all evaluated observation angles, compared to an uncertainty of 6.0% and bias error of 8.0% of the reconstruction by *DFS*. Consequently, the training with glare points resulted in lower errors for the reconstruction of non-axisymmetrical droplet dynamics. More detailed results can be found in Table C.3 in the Appendix.

The temporal evolution of the integral volume for the reconstructed droplet from the  $45^\circ$  observation angle is illustrated in Figure 6.15. As can be seen *DFS-GP* overestimates the volume of the droplet, while *DFS* does underestimate it. Furthermore, the oscillation of the droplet is visible in the error of the reconstruction volume as a low-frequency oscillation, as was already observed for the reconstruction of asymmetric droplet deformation in Subsection 6.3.1. Additionally, a high-frequency fluctuation can be found in the reconstructed volume for the neural network trained without glare points (*DFS*), which indicates further random errors in the reconstruction.

In order to determine the cause for the high-frequency oscillation the contour of the 3D reconstruction is projected in two orthogonal views, in particular, the viewing angle matching the input image (in-plane) and a view rotated by  $90^\circ$  around the height axis, representing the reconstruction in the depth coordinate (out-of-plane). Figure 6.16a shows the temporal evolution of the in-plane



(a) In-plane contour of the reconstructed droplet shapes over time (colored) in comparison to the contour of the input shadowgraph (dashed black) for the reconstruction of droplet impingement on the structured PLA substrate at  $\alpha = 90^\circ$  by *DFS-GP* (left) and *DFS* (right).

(b) Out-of-plane contour of the reconstructed droplet shapes for consecutive frames for the reconstruction of droplet impingement on the structured PLA substrate at  $\alpha = 45^\circ$  by *DFS-GP* (top) and *DFS* (bottom).

**Figure 6.16:** In-plane and out-of-plane contours of the reconstructed droplet shapes. Figures adapted from [DKS24].


contour for the droplet impingement on the PLA substrate at an observation angle of  $45^\circ$  reconstructed by *DFS-GP* (left) and *DFS* (right) in comparison to the contour of the respective input shadowgraph. As can be seen, the in-plane reconstruction reaches a perfect agreement with the input data. This observation holds true for all other frames of the input sequence and for both models with very little exception. Consequently, the influence of the in-plane reconstruction on the error in the volume can be ruled out.

Figure 6.16b shows the temporal evolution of the out-of-plane contour from the volumetric reconstruction at an observation angle of  $45^\circ$  around a time step with particular high volumetric error for both *DFS-GP* and *DFS* (see Figure 6.15). The oscillation of the droplet is already significantly dampened around this time instance, as indicated by the minimal movement of the apex of the droplet. As can be seen, the resulting deformation of the droplet in the out-of-plane coordinate with *DFS-GP* conserves the volume of the droplet. However, *DFS* reconstructs two of the consecutive time frames with a significantly deviating extent in the out-of-plane coordinate. For the time steps at  $t = 58.7\text{ms}$  (orange contour) and  $t = 59.2\text{ms}$  (teal contour) the depth of the droplet deviates significantly from the respective frames directly before and after, thus explaining the high-frequency volume oscillation. Furthermore, it becomes apparent from the comparison of the reconstruction results of both networks that *DFS* underestimates the out-of-plane extent for all time steps, as can be seen by the smaller width of the contours in Figure 6.16b in comparison to *DFS-GP*. Consequently, the underestimation of the droplet's depth results in the observed bias error in the integral volume of the reconstruction (see Figure 6.15).

The higher error in the out-of-plane reconstruction was an expected result as the available features in the image are much more sparse in the case of images with glare points, or even missing in the case of a pure shadowgraph input in comparison to the in-plane reconstruction, for which the two-dimensional droplet contour is available. There are two key findings that can be derived from this observation. First, the neural network can be trained for a depth estimation that relies completely on the two-dimensional contour of a shadowgraph. Second, the glare points successfully encode additional three-dimensional information that is considered by the neural network during


reconstruction, which leads to a significant improvement in the depth estimation. Thus, the presented results demonstrate the effectiveness of glare points for constraining the global and local reconstruction of the three-dimensional gas-liquid interface, which leads to a higher reconstruction accuracy.” from [DKS24, Section 3.5]

### 6.3.3 Reconstruction of highly deformed gas-liquid interfaces



**Figure 6.17:** Reconstructed three-dimensional gas-liquid interfaces (bottom) and corresponding experimental input images (top) for droplet impingement on the structured hydrophobic PDMS substrate at  $\alpha = 0^\circ$  observation angle. Figure adopted from [DKS24].


“In the second series of experiments droplet impingement on the square-groove structured PDMS substrate was reconstructed volumetrically. The experiments featured the impingement of water droplets with an equivalent diameter of  $d_0 = 2.26$  mm on the PDMS substrate at an impact velocity  $u_0 = 0.88$  m/s that was recorded at  $0^\circ$ ,  $45^\circ$  and  $90^\circ$  orientation angles. Since PDMS is a hydrophobic material, its hydrophobicity is further increased by the surface structure, due to the increased surface area [247]. The equilibrium contact angle was measured to be  $\Theta_{\text{eq,p}} = 107^\circ$  in parallel and  $\Theta_{\text{eq,t}} = 97^\circ$  in transversal direction. The dynamic contact angles in both directions can be found in Table C.1 in the appendix. The pronounced hydrophobicity of the structured PDMS sample led to a partial rebound of the droplet in some of the experiments. The resulting complex three-dimensional deformation of the droplet is illustrated by exemplary images in Figure 6.17, alongside their respective volumetric reconstruction.



**Figure 6.18:** Temporal evolution of the normalized integral volume of the reconstruction for droplet impingement on the structured PDMS substrate at  $\alpha = 90^\circ$  observation angle with *DFS-GP* and *DFS*. Figure adopted from [DKS24].

Figure 6.18 shows the temporal evolution of the integral volume of the droplet during impingement on the PDMS substrate reconstructed from images taken at  $90^\circ$  orientation angle. Both *DFS-GP* and *DFS* successfully reconstruct the three-dimensional dynamics of the significantly deformed gas-liquid interface well, as indicated by the good agreement of the integral volume with the ground truth for most of its temporal evolution. However, both networks underestimate the volume. Furthermore, the fluctuations of the reconstructed volume during the early stages of the impact, characterized by a rapid droplet deformation, have a larger magnitude compared to the reconstruction of the lesser deformed droplets during impingement on the PLA substrate, in particular for the reconstruction by *DFS*. During the later stages of the droplet impact *DFS-GP* correctly estimates a constant volume, while *DFS* exhibits a low-frequency oscillation in the reconstructed volume.


*DFS-GP* reached an uncertainty of 6.6% and a bias error of 6.2% averaged over all tested orientations, while *DFS* reached an uncertainty of 6.5% and bias error of 8.2%. For both versions of the network, the errors are higher in comparison to reconstruction results for the PLA substrate. Furthermore, it was found that the reconstruction accuracy is dependent on the observation angle. The  $45^\circ$  observation angle resulted in a significantly lower uncertainty, as well as a lower bias error for the reconstruction of both droplet impingement on PDMS and PLA. *DFS-GP* reaches a combined uncertainty of  $\sigma_V = 3.4\%$  and bias error of  $\delta_V = 4.1\%$ , which is significantly lower than the errors in the  $0^\circ$  orientation with  $\sigma_V = 8.5\%$  and  $\delta_V = 6.7\%$ , and  $90^\circ$  orientation with  $\sigma_V = 5.5\%$  and  $\delta_V = 5.9\%$ . A similar behavior was observed for the reconstruction with *DFS*. A detailed summary of the uncertainties and bias errors is listed in Table C.3.



**Figure 6.19:** Comparison of the depth estimation for a time step with a high volumetric error during the reconstruction of droplet impingement on the structured PDMS substrate by *DFS-GP* (left) and *DFS* (right). Figure adopted from [DKS24].

The reconstruction results of both *DFS-GP* and *DFS* are shown in Figure 6.19 for the same image frame recorded in the experiments. The reconstruction is rotated by  $90^\circ$  with respect to the image plane and therefore shows the depth estimation of the two neural network states. As can be seen, the complicated shape of the gas-liquid interface was inferred in a physically reasonable way by *DFS-GP*, while *DFS* estimates a physically incorrect depth. This error in the depth estimation can be related to a locally high error in the reconstructed volume (see Figure 6.18), which indicates that the reconstruction of the particular frame proved difficult for the neural network trained without glare points.


Figure 6.20 shows the reconstruction of an image frame recorded during droplet impingement on the PDMS substrate at an observation angle of  $0^\circ$ . The reconstructed droplet shape is rotated,



**Figure 6.20:** Reconstruction of the droplet during the wetting of the structured surface by *DFS-GP* (right) and the respective input glare-point shadowgraph recorded at  $\alpha = 0^\circ$  observation angle. Note that the reconstruction is rotated. Figure adopted from [DKS24].

so that the contact area between the liquid and the solid substrate is visible. As can be seen, the grooves that are also visible in the input image are extended over the whole depth of the reconstructed volume, thus revealing that small-scale features are accurately reconstructed by the neural network. However, the reconstruction of the wetted state was only possible if the surface structure was visible in the input images and consequently only for the  $\alpha = 0^\circ$  observation angle. Under other observation angles the surface appears to be flat in the images and the liquid-solid interface is reconstructed as a flat plane by the neural network.

It should be noted that the training dataset  $[\mathcal{D}_{\text{DFS}}]$  contained both flat and structured surfaces, which could be a source of confusion for the neural network that leads to the reconstruction of flat surfaces if the grooves are not visible. In order to determine the effect of training data on the reconstruction of structured surface wetting, a benchmark version of the network *DS-GP* is trained solely on the part of the training data  $[\mathcal{D}_{\text{DS}}]$  that considers the droplet impingement on structured surfaces. It was found that the *DS-GP* network also only reconstructs the wetted state accurately, if the grooves of the substrate are visible in the input image, thus indicating that the effect is independent from the training data.



**Figure 6.21:** Reconstruction (right) by *DFS-GP* and the respective input images (left) of droplet impingement with partial rebound leading to multiple droplets (top) and a Worthington jet (bottom). Note that the reconstruction is rotated by  $90^\circ$  relative to the input. Figure adopted from [DKS24].

The upper row of Figure 6.21 shows a recording from the experiment and the respective volumetric reconstruction obtained by the *DFS-GP* network for a drop impact that resulted in a partial rebound. As can be seen, a successful volumetric reconstruction of multiple droplets can be carried out by the model. Furthermore, the separation into two droplets, as well as the subsequent coalescence were both reconstructed precisely by the neural network. These results demonstrate that topological changes can be accurately represented by the trained neural network. Other experiments show the emergence of a Worthington jet that leads to the detachment of a tiny droplet at high velocities, which was accurately reconstructed by the network as well as can be seen in the bottom row of Figure 6.21.

The ability to reconstruct these topological changes demonstrates the high flexibility of the three-dimensional representation by a level-set method that underlies the PIFu algorithm. Moreover, the discussed topological changes were not represented in the training dataset, which indicates a certain capability for an extrapolation beyond the training data cases in regard to fluid mechanical regimes. It should be noted that this does not directly imply an extrapolation in terms of the training data distribution. It can be concluded that the strict abidance of the network to the image features – in particular the contour of the shadowgraph – allows for a reconstruction of unknown shapes of the gas-liquid-interface. This capability takes immediate advantage of the presence of additional glare points, as saliently demonstrated in Figure 6.19.” from [DKS24, Section 3.6]

### 6.3.4 Reconstruction of obscured and hidden shapes

“The results show that the neural network was able to reconstruct hidden regions of the gas-liquid interface in a physically meaningful manner. This includes input images with self-occlusion, as seen in Figure 6.13, as well as the wetting state of structured surfaces, indicated by Figure 6.20 and finally, the learned depth estimation, as discussed previously. These results suggest that the neural network is able to learn an approximation of the underlying droplet dynamics that is applied to fill in missing information in the experimental data in a physically reasonable way during reconstruction. Consequently, the data-driven approach can be conditioned for a physically correct reconstruction by training on data that accurately represents the underlying physics of the problem, such as the direct numerical simulation that was used in this study. However, it should be noted that the quality of the reconstruction is reduced in the occluded regions due to artifacts and small-scale errors of the reconstruction, which falls in line with previous reports [191].


Furthermore, the reconstruction of the wetting state of structured surfaces was only successful if the surface structure was visible in the experiments at  $\alpha = 0^\circ$  observation angle, as seen in Figure 6.20. This effect also appeared for a version of the neural network that was exclusively trained on simulations of droplet impingement on structured surfaces (*DS-GP*) and thus was found to be independent of ambiguities due to training examples of the wetting of flat surfaces in the training dataset of *DFS-GP*. These results indicate that features in the image are prioritized by the network over the learned knowledge from the training data and, more specifically strictly abided, as evident from the straight contour of liquid-solid contact area in the reconstruction (see *e.g.* Figure 6.16b).

Moreover, the results highlight the relevance of pre-processing the input images with a binary mask. The mask, which is composed of the shadowgraph contour of the droplet and the contour of the substrate determined from the images recorded in the experiments prior to the impact of the droplet, as described in Subsection 6.1.3, imposes this information on all later frames, where the contour of the droplet would conceal the liquid-solid boundary. Thereby the masking furthers the reconstruction as only relevant information from the experiment is pre-selected and passed onto the




neural network. Consequently, the network does not need to learn to differentiate substrate from droplet, resulting in a simplified reconstruction. In the case of structured surfaces, the masking can be used beyond global shape estimation to additionally reconstruct the local wetting of the grooves. However, care has to be taken for the detection of the substrate in the image, as errors from the position of the ground translate into an erroneous mask, which in turn introduces an error to the volumetric reconstruction that is related to the extent and shape of the droplet close to the contact area. This source of error might be the cause for the low-frequency oscillation of the reconstructed integral volume observed in Figures 6.15 and 6.18.” from [DKS24, Section 3.7]

### 6.3.5 Adhering droplets in external shear flow



**Figure 6.22:** Envelopes of the droplet motion in the numerical simulation (left) and in the experiments (middle left to right) at different bulk velocities of external flow. Main flow direction from right to left. Figure adopted from [DHK<sup>+</sup>25].

“The network trained on synthetic data was applied for the volumetric reconstruction from images obtained by experiments at different external flow velocities, in particular at the measured bulk velocities of  $u_B = 5.85$  m/s,  $u_B = 7.58$  m/s and  $u_B = 8.32$  m/s. Figure 6.22 shows the envelopes of the droplet motion throughout the experiments for one representative case at each velocity in comparison to the numerical simulation. As can be seen, an increase in the external flow velocity resulted in a larger degree of droplet deformation. Furthermore, the comparison to the numerical data reveals that only at  $u_B = 5.85$  m/s a similar degree of deformation can be observed in the experiments. Consequently, the neural network trained on the numerical data has to generalize to significantly more deformed and unknown droplet shapes in order to successfully reconstruct the experiments at  $u_B = 7.58$  m/s and  $u_B = 8.32$  m/s. Note, that the same representative experiments (cases 1, 8, and 10) will be the subject of the following evaluation unless stated otherwise. The results of further experiments can be found in the appendix.



**Figure 6.23:** Temporal evolution of the reconstructed droplet volume normalized by the mean volume for experiments with adhering droplets at different bulk velocities of external flow. Figure adopted from [DHK<sup>+</sup>25].




**Table 6.2:** Uncertainty  $\sigma_V$  of the reconstructed integral volume in percent of the mean volume for different external flow velocities averaged over all experiments conducted at a certain velocity. Table adopted from [DHK<sup>+</sup>25].

cases	1 – 6	7 – 9	10, 11
$u_B$ [m/s]	5.85	7.58	8.32
$\sigma_V$ [% $\bar{V}$ ]	3.2	4.1	5.4

As no volumetric ground truth data is available for the reconstruction from images recorded in the experiments, the temporal evolution of the reconstructed droplet volume is considered for the evaluation of the reconstruction accuracy. In the experiments no significant evaporation occurred within the time frame of one recorded image series (1,000 images over 133.3 ms) and, consequently, the volume of the droplet can be assumed to be constant.


Figure 6.23 shows the temporal evolution of the normalized integral volume of the reconstructed droplet for one representative experiment at each evaluated velocity of the external flow. As can be seen, the volume of the reconstruction fluctuates periodically around the mean, with an increasing amplitude towards higher velocities of the external flow, *i.e.* higher degrees of droplet deformation. The measured uncertainties of the reconstruction, as detailed in Table 6.2, further underline that the reconstruction at higher degrees of droplet deformation is subject to higher uncertainty. As indicated by the low uncertainty for the reconstruction at  $u_B = 5.85$  m/s the neural network trained on synthetic image data generalizes well to the reconstruction of experiments with a similar degree of droplet deformation to the numerical simulation that underlies the training data. The successful reconstruction of the significantly more deformed gas-liquid interface in the experiments at  $u_B = 7.58$  m/s and  $u_B = 8.32$  m/s, which resulted in droplet shapes unknown to the neural network, indicate the trained network can extrapolate to different fluid mechanical conditions. Furthermore, these results demonstrate the robustness of the method, which, however, is constrained by a reduced accuracy for the reconstruction of different data.



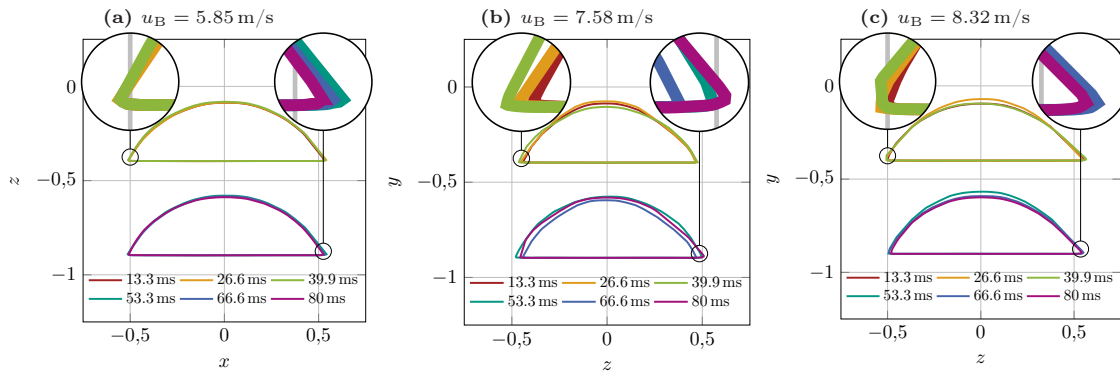
**Figure 6.24:** Input images (columns one and three) and out-of-plane projection (columns two and four) of the respective three-dimensional interfaces for the reconstruction at  $u_B = 8.32$  m/s. The top row shows the experimental recordings and the corresponding reconstructed interfaces and the bottom row shows matching samples from the synthetic training dataset [ $\mathcal{D}_{DSII}$ ]. Figure adapted from [DHK<sup>+</sup>25].

The periodical oscillation of the reconstructed volume was found to coincide with the oscillation of the droplet contour. Moreover, the input images to the volumetric reconstruction at the extrema of the volumetric deviation appear similar. Figure 6.24 shows examples of the input images from the experiment that resulted in a local maximum [top row left] and minimum [top row right] of the volume and the respective out-of-plane projection of the reconstructed droplet shape. In particular, the position of the glare points relative to each other and to the contour of the droplet was similar between the groups of input images that resulted in either the maxima or minima of reconstructed volume, which indicates a similar shape of the droplet interface. Most saliently, a low position of the blue  $p_{gp} = 1$  glare point in the input images resulted in a minimum volume, while a maximum in the reconstructed volume correlated with a high position of the  $p_{gp} = 1$  glare point.

For both cases, synthetic images with similar relative glare point positions can be found in the training dataset  $[\mathcal{D}_{\text{DSh}}]$ , which are shown in [the bottom row] in Figure 6.24. Note, that Figure 6.24 shows sample images from the experiments at  $u_B = 8.32 \text{ m/s}$  for better visibility of the differences between the minimum and maximum case, as this experiment resulted in the highest deviations in the reconstructed volume. Consequently, the contour of the droplet in the reconstruction is more deformed than the synthetic training data. However, the reconstruction at lower velocities follows the same behavior. The comparison of the out-of-plane projection between the reconstruction and the training data reveals a very similar shape of the contour. These results indicate that the neural network successfully learned the relation of the glare point positions to the three-dimensional geometry of the droplet and that this encoding of 3D information can be applied for depth estimation during reconstruction. However, as the neural network has been trained for the reconstruction of different droplet volumes by an augmentation of the training data, as described in Subsection 6.1.3 it is agnostic to the volume of the droplet. Consequently, the shape of the contour is reconstructed with disregard to the integral volume of the droplet and thus the volume of the reconstruction is not conserved in time.

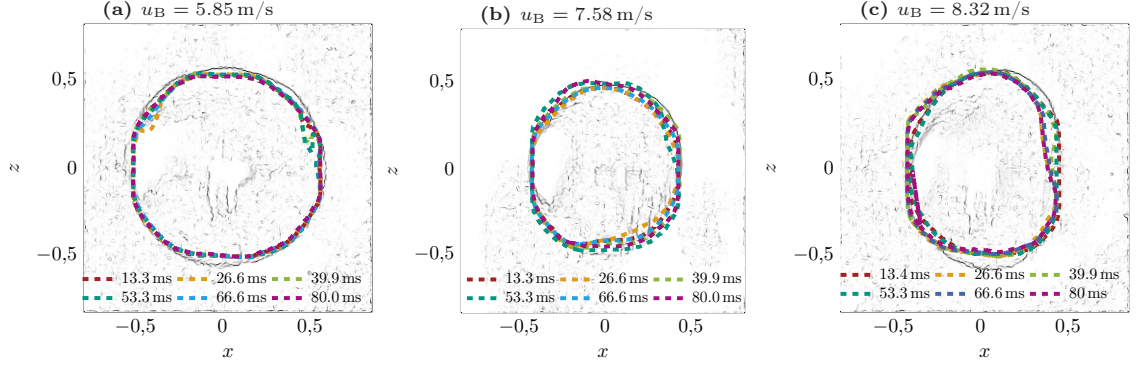


**Figure 6.25:** Temporal evolution of the reconstructed in-plane contour for different bulk velocities of external flow. The dashed colored lines indicate the reconstructed contours at different times and the black solid lines indicate the corresponding contours observed in the experiments. The main flow direction is from right to left. Figure adapted from [DHK<sup>+</sup>25].



**Figure 6.26:** Temporal evolution of the reconstructed out-of-plane contour for different bulk velocities of external flow. The solid colored lines indicate the reconstructed contours at different times. The image plane is aligned with the main flow direction. Figure adapted from [DHK<sup>+</sup>25].

In order to evaluate the in-plane and the out-of-plane accuracy of the reconstruction in more detail, two orthogonal projections of the reconstructed droplet interface are compared to the side- and top-view contours extracted from the corresponding images recorded in the experiments. The results of the in-plane reconstruction for different external flow velocities are presented in Figure 6.25 and the results of the out-of-plane reconstruction are shown in Figure 6.26. As can be seen, there is an



**Figure 6.27:** Temporal evolution of the reconstructed wetted area for different bulk velocities of external flow. The dashed colored lines indicate the reconstructed contours at different times and the grayscale image in the background indicates the corresponding contour observed in the experiments. The main flow direction is from right to left. Figure adapted from [DHK<sup>+</sup>25].

almost identical agreement of the reconstructed in-plane contours with the respective contours from the experiment, which underlines the findings of a high in-plane accuracy for the reconstruction of synthetic data. The planar symmetry in the out-of-plane direction, which was expected due to the symmetry of the flow, was also reconstructed well for the experimental data. However, the reconstruction exhibits a fluctuation in the out-of-plane direction that is particularly noticeable in Figure 6.26(b). This out-of-plane fluctuation is identified as the sole cause for the uncertainty of the volumetric reconstruction since the in-plane reconstruction is highly accurate. The out-of-plane reconstruction relies on the learned droplet geometry and the encoding of three-dimensional information by the glare points and, therefore, the accuracy of the reconstruction is reduced the more the input data deviates from the training data distribution. The in-plane reconstruction is based on the image features, most importantly the contour of the shadowgraph, which provides a significantly larger amount and more direct information for the reconstruction. As the in-plane reconstruction adheres closely to the contour in the images, the neural network adapts well to unknown shapes outside of the training data distribution. These results fall in line with the observed uncertainty of the out-of-plane reconstruction for synthetic data in Subsection 6.2.1 and previous observations of impinging droplets in Subsection 6.3.2.

The accuracy of the depth estimation is further evaluated by the comparison of the top-view projection of the reconstruction to the footprint of the droplet observed in the experiments through the top-view camera. As can be seen in Figure 6.22, the contact line of the droplet stays pinned during the duration of the experiments and the wetted area is therefore constant over time. Figure 6.27 shows the contour of the contact lines extracted from the experimental recordings at different external flow velocities. The contours from the experiments are overlayed with the projected contours of the reconstructed droplet shapes for six different time instances. As can be seen, the shape of the wetted area in the reconstruction is similar to the experiment and self-similar over time for  $u_B = 5.85$  m/s, while there is a larger deviation to the ground truth at higher external flow velocities. The aspect ratio of the footprint changes significantly towards higher velocities as the droplet is deformed by the external flow. Remarkably, the neural network successfully reconstructs these different aspect ratios, which is a significant extrapolation from the training data, which only contained one simulated case at the external flow velocity  $u_B = 5.85$  m/s, for which the footprint had an aspect ratio close to one. For the same velocity in the experiments ( $u_B = 5.85$  m/s) a very good agreement of the reconstruction with the ground truth was reached. At higher velocities ( $u_B = 7.58$  and  $u_B = 8.32$ ) the out-of-plane extent of the droplet was generally underestimated and reconstructed with a higher uncertainty, as indicated by the larger variation in the reconstructed footprints. Furthermore, the unknown droplet shapes in the experiments related to higher aspect


ratios of the wetted were reconstructed close to the distribution of shapes in the training data, thus revealing a bias of the trained model. For all velocities the contour of the reconstruction is more angular in comparison to the experiment, but similar to the training data, as seen in Figure 6.9, which is a further indication of model bias. However, these results also demonstrate that the geometry of the droplet was learned faithfully to the training dataset  $[\mathcal{D}_{\text{DSh}}]$  by the network. Therefore, the accuracy of the reconstruction is dependent on the truthful representation of the droplet dynamics in the training data. The successful reconstruction of unknown droplet shapes with higher aspect ratios indicates that the network learned to utilize the glare points for the depth estimation.” from [DHK<sup>+</sup>25, Section 3.2]

### 6.3.6 Extrapolation to unknown droplet dynamics

**Evaluation procedure** To further investigate the generalization capability of the trained models of droplet dynamics, the reconstruction framework is employed for the reconstruction of droplet impingement experiments with significantly different wetting properties in comparison to previous experiments. Specifically, droplet impingement experiments involving patterned substrates with gradients in wettability are conducted, resulting in a complex three-dimensional deformation of the gas-liquid interface, as demonstrated in Section 6.1.2. The resulting observed interface shapes significantly differ from both the previous experiments, as well as the previously used training data distributions. The PIFu neural network is trained on the datasets involving droplet impingement on flat and structured substrates ( $\mathcal{D}_{\text{DFS}}$ ), as well as adhering droplets in shear flow ( $\mathcal{D}_{\text{DSh}}$ ). Subsequently, the trained variants of the network are employed for the reconstruction of the novel experimental data, to reveal the generalization capability of the data-driven reconstruction approach. Furthermore, it is investigated if the combination of different training datasets and, thus, an increase in both the variance and size of the training data distribution can be beneficial for the reconstruction of unknown droplet shapes. In order to allow for the reconstruction on the basis of experimental recordings that feature green glare points from both lateral light sources instead of differentiated red and green glare points, alternative versions of the PIFu neural network are trained on synthetic datasets that contain only green glare points.

In the following, models trained with data  $\mathcal{D}$  containing only **green** glare points are indicated by  $X\text{-}g$  in order to differentiate them from previous models with red and green **G**lare **P**oints, which are indicated by  $X\text{-}GP$ , where  $X$  is a placeholder for the respective training dataset  $\mathcal{D}_X$ . The datasets  $\mathcal{D}_X\text{-}g$  are obtained by merging the red and green image channels of the datasets  $\mathcal{D}_X\text{-}GP$  into a new green image channel while setting the red image channel to zero. Note, that the addition of image channels is trivial for the synthetically rendered images since the information of the glare points and the shadowgraph are confined to the respective channels of the RGB images. The same procedure for images obtained in the experiments would require a careful pre-processing of the images through the color correction techniques introduced in Subsection 4.2. The suitability of different training datasets is evaluated in a comparative study between a model trained on the adhering droplet dataset  $DSh\text{-}g$ , a model trained on the dataset for impinging droplets  $DFS\text{-}g$ , and a model trained on the combined dataset  $DFSSh\text{-}g$ . To ensure consistency with the experiments, a new droplet impingement dataset  $\mathcal{D}_{\text{DFS}}$  is generated, which consists of synthetic images rendered at a scattering angle of  $96.3^\circ$ . All versions of the PIFu neural network are trained with the same set of hyperparameters, as detailed in Section 3.4.

**Validation on synthetic data** The approach of using only a single color of glare points is first validated by the comparison of the reconstruction accuracy on synthetic data between the respective  $X\text{-}g$  and  $X\text{-}GP$  models. As can be seen in Figure 6.28, the three considered  $X\text{-}g$  models reach a




**Figure 6.28:** Average 3D-IOU of the reconstruction for the models trained on data with only green glare points *DFS-g*, *DSh-g*, and *DFSSh-g* on the respective validation datasets during training of the network in comparison to models trained on the same datasets with red and green glare points *DFS-GP*, *DSh-GP*, and *DFSSh-GP*.

similarly high accuracy at the end of the training in comparison to their *X-GP* counterparts, thereby validating the approach of employing only one color of glare points. The rendering setup with a scattering angle of  $96.3^\circ$  for the lateral lights produces only  $p_{gp} = 0$  glare points, which are confined to either the left or right hemisphere of the droplet. Therefore, the synthetic training dataset facilitates the learning of the differentiation between the glare points of the same color by the neural network. However, previous experiments revealed that in practice spurious higher-order glare points are produced in the experiments, in particular for highly deformed gas-liquid interface shapes and reflective substrates. Therefore, the *X-g* models are expected to have a reduced accuracy in comparison to the *X-GP* models for the reconstruction of recordings obtained in the experiments, as the images carry a reduced amount of information. Consequently, the *X-g* models can be considered as a lower estimate for the reconstruction performance of the previous *X-GP* models. However, the *X-g* models trained on different datasets are comparable among themselves, which allows to evaluate the influence of the training data on the generalization capability of the trained droplet models.


**Reconstruction of patterned substrates with gradients in wettability** The networks trained on the different synthetic datasets are employed for the reconstruction of images recorded in droplet impingement experiments involving patterned substrates with gradients in wettability, which have been discussed in Section 6.1.2. Out of the three candidate models, *DFS-g* yielded the best reconstruction results, as indicated by comparatively low uncertainty of  $\sigma_V = 8.1\%$  and bias error  $\delta_V = 12.3\%$  averaged over all experiments in combination with a high quality of the reconstructed three-dimensional gas-liquid interface. The reconstruction accuracy of *DFSSh-g* was considerably lower ( $\sigma_V = 9.5\%$  and  $\delta_V = 11.6\%$ ), while *DSh-g* resulted in a largely unsuccessful reconstruction ( $\sigma_V = 28.7\%$  and  $\delta_V = 16.7\%$ ). These results are consistent over all tested substrates (see Appendix Table ??). Furthermore, it was observed that a greater deformation of the gas-liquid interface due to a larger displacement of the droplet was related to a higher uncertainty of the reconstruction. These results confirm the earlier observation that more severely deformed gas-liquid interfaces of adhering droplets are more challenging to reconstruct (see Section 6.3.5).

A visual inspection of the reconstructed three-dimensional gas-liquid interface shapes for the three models revealed that *DFS-g* was able to reconstruct all degrees of gradiented surfaces successfully,



**Figure 6.29:** Temporal evolution of the normalized integral volume of the reconstruction for droplet impingement on the patterned PDMS substrate *SMG3* at  $\alpha = 0^\circ$  observation angle with *DSh-g*, *DFS-g*, and *DFSSh-g*.

resulting in high-quality meshes with watertight surfaces. The reconstruction results for *DFSSh-g* exhibited a lower quality, with the surface of the reconstructed meshes appearing jagged and containing cavities at certain time steps, particularly for frames involving a high deformation of the interface immediately after impact. This leads to a systematic underestimation of the droplet volume, which is reflected in Figure 6.29, showing the volume of the reconstructed gas-liquid interface at each frame normalized by the ground truth volume measured from the experimental recordings. Overall, the reconstruction results of *DSh-g* showed the lowest quality, as the model was not able to infer the depth accurately. This was especially pronounced for frames that deviated most from the shape distribution of adhering droplets, such as the ellipsoidal shape of the falling droplet and the elongated droplet shapes observed in the retraction phase (see *e.g.* Figure 6.4, middle frame). It can be concluded that the training data distribution of *DFS-g* represents the dynamics of the gas-liquid interface in the unknown experimental data sufficiently to allow for successful reconstruction, while the training data for *DSh-g* and the unknown experimental data do not sufficiently overlap. As a result, the increased variance of the training data for *DFSSh-g* offers no benefit. The additional data from *DSh-g* does not increase the overlap with the test dataset's distribution, but apparently shifts the bias of the trained droplet model toward the shape distribution of adhering droplets, moving it away from the distribution of impinging droplet shapes.



**Figure 6.30:** Temporal evolution of bottom view projection of the gas-liquid interface during droplet impingement on the *SMG2* substrate. The dashed colored lines indicate the interface reconstruction by *DFS-g* and the solid black lines indicate the contours extracted from the bottom view recordings in the experiments.




The reconstruction results of *DFS-g* and *DFSSh-g* show a good agreement of the in-plane contour to the experimental input data and even *DSh-g* was able to reconstruct the in-plane contour on most frames. These results align with the earlier observation that the directly available information on the shadowgraph contour ensures an accurate reconstruction of the in-plane contour. The out-of-plane reconstruction proved to be more sensible towards training data effects as the depth estimation relies more heavily on the trained droplet model in order to compensate for the lower amount of available information in the input images. Figure 6.30 shows the temporal evolution of the gas-liquid interface contour in the bottom view projection for droplet impingement of the *SMG2* substrate, as observed in the experiments (black solid line) and reconstructed by *DFS-g* (colored dashed lines). The projected area is compared in order to ensure the consistency of the results, as the high contact angles of the superhydrophobic surface result in the self-occlusion of the contact line and, therefore, the wetted area cannot be directly observed (*cp.* Figure 6.4). As can be seen in Figure 6.30 the projected area of the reconstructed droplets has the shape of an ellipse in all time steps, while the real shape of the droplets interface at  $t = 8.0$  ms significantly deviates from elliptical cross-section, which indicates a bias of the trained model of droplet dynamics. The training dataset for *DFS-g* featured the impingement of droplets on structured substrates, where the deformation of droplets occurred within two planes of reflective symmetry, and the impingement on flat substrates, which yielded an axisymmetric deformation. Therefore, the distribution of 3D ground truth shapes in the training data predominantly consisted of droplets with elliptical cross-sections or otherwise circular cross-sections. The distribution of ground truth shapes with elliptical cross-sections in the training data apparently introduces a bias to the neural network, which is especially noticeable for the depth estimation that relies more on the trained droplet dynamics. Consequently, the reconstruction exhibits increasing errors the more the real gas-liquid interface deviates from an elliptical cross-section. However, the absolute depth of the droplet and in turn its volume was approximated well for different aspect ratios as evident by the reconstructed contours at the time steps 1.1 ms, and 19.1 ms. These results fall in line with earlier observations in Section 6.3.5 that demonstrated the successful reconstruction of adhering droplets with different aspect ratios by a model that was trained on less severely deformed droplet shapes with a different aspect ratio.

### 6.3.7 Surface quality of the reconstructed interfaces

“The implicit representation of the three-dimensional droplet geometry by the PIFu neural network [191] allows for the reconstruction at an arbitrary resolution. As higher resolutions are expected to result in a higher quality of the surface, but also increased computational costs, in the following the results of the reconstruction at different resolutions were evaluated. Table 6.3 details the average time required for the reconstruction of one time step at different resolutions. The domain for the reconstruction was discretized by an equidistant 3D grid with the same amount of grid nodes in all directions. The trained PIFu network was sampled on this grid to predict the scalar occupancy field of the phase distribution. Subsequently, the isosurface extraction algorithm marching cubes [130, 136] was employed to reconstruct the surface mesh from the occupancy field. All calculations were performed on a single Nvidia RTX A5000 graphics processing unit. As expected the computational costs grow quickly towards higher resolutions. However, the octree [146] structure used during reconstruction significantly reduces the reconstruction time at higher resolutions up to  $512^3$  grid nodes, which becomes obvious by the deviation from the cubic growth rule that would be expected otherwise. Note that this trend only holds true until  $512^3$  grid nodes due to given hardware limitations.

Figure 6.31 shows the results of the volumetric reconstruction at different resolutions for the same input image obtained in the experiments [involving adhering droplets] at an external flow velocity






**Figure 6.31:** Volumetric reconstruction at different resolution for the same input image from the experiment at  $u_B = 8.32$  m/s, from left to right:  $64^3$ ,  $128^3$ ,  $256^3$ ,  $512^3$ ,  $1024^3$  grid nodes resolution. Figure adopted from [DHK<sup>+</sup>25].

**Table 6.3:** Average time required for the reconstruction of one time step at different output resolutions. Table adopted from [DHK<sup>+</sup>25].

grid nodes	$64^3$	$128^3$	$256^3$	$512^3$	$1,024^3$
time [s]	0.2	2.5	4.4	17.0	266.5

$u_B = 8.32$  m/s. As can be seen, there is a significant visual improvement in the surface smoothness as the resolution increases from  $64^3$  to  $512^3$ . However, with the increase in the resolution from  $512^3$  to  $1,024^3$  grid nodes, the quality of the reconstruction did not increase further. In fact, a marginal decrease of the surface quality can be observed, as small-scale ripples appear on the surface for a resolution of  $1,024^3$  nodes, which are likely smoothed out by lower resolutions. Considering the total time of 4.7 h required for the reconstruction of one experiment with 1,000 frames, a resolution of  $512^3$  grid nodes was found to be optimal within the scope of obtaining accurate and smooth surface meshes for distortion correction. Note, that all reconstruction results presented in this work have a resolution of  $512^3$  grid nodes.

It should be noted, that the presented reconstruction method is able to reconstruct the gas-liquid interface of the droplet at a significantly higher resolution compared to the training data. The spatial resolution of the numerical simulation that underlies the training data was 20 cells in the vertical direction and 17 cells in the streamwise and spanwise direction at  $t = 0$ . These results demonstrate that the neural network has the capability to learn a highly accurate representation of the droplet geometry even from data much coarser than the targeted reconstruction resolution. Furthermore, the approach of smoothing the ground truth gas-liquid interfaces that were extracted from the numerical simulation is validated by the positive reconstruction results. In order to further enhance the surface quality of the reconstruction smoothing by means of the Taubin filter (see Subsection 5.2) is employed.



**Figure 6.32:** Smoothed volumetric reconstruction at  $512^3$  grid nodes resolution. Figure adopted from [DHK<sup>+</sup>25].

As shown in Figure 6.32, the proposed process yields a smooth surface mesh of the droplet’s gas-liquid interface that can be used for the intended distortion correction of PIV measurements. Such distortion correction techniques require surface meshes that exhibit both a high fidelity to the true interface shape in the experiments and a smooth curvature, in order to allow for the accurate calculation of light refraction at the gas-liquid interface. The prediction of the scalar occupancy field for the phase distribution by the PIFu neural network serves as an accurate basis for the surface reconstruction. In this study, the efficient implementation [130] of the marching cubes algorithm [136] is employed to reconstruct the surface meshes. However, alternative isosurface extraction algorithms, including extensions of the marching cubes algorithm [248, 118], methods based on Delaunay triangulation [14], or methods used in the numerical simulation of two-phase flows, such as the Piecewise-Linear Interface Calculation (PLIC) scheme [259], may offer improved reconstruction performance. Therefore, further research is required to assess the suitability of the reconstructed surfaces for distortion correction in the experiments. This includes further development of surface reconstruction methods and smoothing operations, as well as experimentally validating the reconstructed meshes in the ray tracing approach to confirm their feasibility for distortion correction.” from [DHK<sup>+</sup>25, Section 3.3]

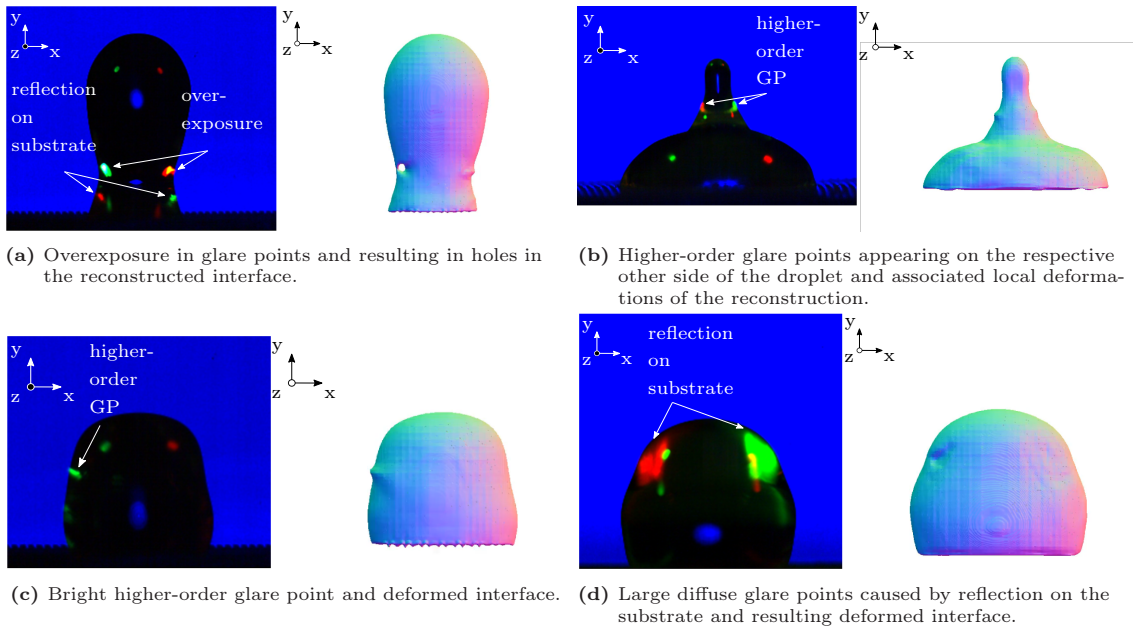
## 6.4 Limitations of the method

“The comparison of the images recorded in the experiments with droplet impingement on the structured PLA and PDMS substrates reveals that a reflection of the incident light from the lateral illumination on the substrate occurs predominantly for the PDMS substrate. The lateral illumination has a  $45^\circ$  incidence angle. As such, light can be reflected at the liquid-solid contact area after passing through the droplet and subsequently produce additional glare points on the gas-liquid interface while exiting the droplet again. Additionally, the light can reflect on the substrate outside of the droplet and subsequently enter the droplet at a  $-45^\circ$  incidence angle. Both mechanisms of reflection were confirmed to occur in the experiments. The PLA substrate was produced from a low-reflective material in order to suppress spurious glare points from reflection and allow for an evaluation of their influence on the quality of the reconstruction. Furthermore, impingement on the hydrophobic PDMS substrate led to a higher degree of droplet deformation in comparison to the PLA substrate. Photographs of either substrate are shown in Figure C.3. The corresponding wettability properties are listed in Table C.1. The resulting complex shapes of the gas-liquid interface allowed for the emergence of higher-order glare points [232, 233] that are created by internal reflection of the light on the gas-liquid interface. The consequence of the internal reflection from higher-order glare points and reflection on the substrate within the droplet is that the glare points change their position to the other hemisphere of the droplet, as can be seen for the lowest pair of glare points in Figure 6.33a.

A second consequence of the internal reflections is the focusing of the incident light by the curved droplet contour, which can result in a magnification of the light intensity for higher-order glare points, as obvious from Figure 6.33a by the second row of glare points. Overexposure in the images leads to color clipping and a loss of information in the experiments, as the clipped image channel cannot register any further increase in intensity [165]. As evident from the reconstruction result in Figure 6.33a, color clipping in the regions of the bright glare points causes nonphysical artifacts in the reconstructed gas-liquid interface, such as the hole on the left and the dimple on the right of the reconstructed geometry. Further examples of images with reflection and higher-order glare points and their respective volumetric reconstruction can be found in Figure 6.33.

The comparison of the uncertainty and bias errors of the reconstructed integral volume between droplet impingement on the PLA and PDMS substrate, as detailed in Table C.3, reveals significantly lower errors for the reconstruction of the PLA case. Furthermore, the specific inspection of frames associated with a particularly high volumetric error as indicated by Figure 6.18 points to an association with images that are affected by a higher degree of disturbance through reflections and overexposure. These results indicate that reflection and higher-order glare points lead to higher errors in the reconstruction of the gas-liquid interface. This effect has to be expected since the input images to the neural network have to be similar to its training data in order to reach a high quality in the reconstruction. This aspect has been thoroughly elaborated in Subsection 3.3 and, in particular, reflection on the substrate, as well as high-order glare points were not modeled for synthetic training data generation. Therefore, the disturbed input data causes a high uncertainty in the reconstruction and can even lead to the erroneous prediction of the local gas-liquid interface, as seen in Figure 6.33a.

Conversely, these results further underline that the neural network considers the glare points during reconstruction in order to guide the reconstruction. Furthermore, the results indicate a certain robustness towards disturbance outside of local errors, as glare in unexpected regions (see Figure 6.33a) and unexpectedly large glare points (Figure 6.33d) appear to be mostly disregarded by the neural network for the reconstruction of the global shape.” from [DKS24, Appendix A]



**Figure 6.33:** Input images with overexposure and higher-order glare points (left) and the respective reconstruction (right). Figures adapted from [DKS24].

## 6.5 Concluding remarks regarding the volumetric reconstruction of droplet flows by means of deep learning

“The positive results for the spatio-temporal reconstruction of the gas-liquid interface [...] from monocular experimental recordings demonstrate the success of the proposed approach based on neural networks and synthetic training data generation. The employment of synthetic image rendering from the results of direct numerical simulation proved to be an effective method for the

generation of suitable training data while eliminating potential errors due to the inherent discrepancy between experimental and numerical results.” from [DKS24, Section 4] “Furthermore, the positive results for the reconstruction based on images recorded in the experiments demonstrate the applicability of the neural network trained on synthetic data to the real-world use case and thus validate the approach of training on synthetic data.” from [DHK<sup>+</sup>25, Section 4] “The single-camera setup required for the proposed method is both cost-effective and easy to calibrate and, therefore, accessible to a wide field of applications.

It was shown that the additional information on the three-dimensional shape encoded in the glare points is successfully exploited by the neural network during reconstruction, which results in a higher accuracy of the reconstruction. In particular, color-coded glare points improve the depth estimation of the neural network and, thus, allow for the reconstruction of the complex non-axisymmetric shape of the gas-liquid-interface for droplet impingement on structured substrates from any arbitrary azimuth angle. These findings recommend the proposed method for the reconstruction of the three-dimensional droplet dynamics for an impact on substrates for which the orientation angle is not known a-priori, which consequently allows for an efficient characterization of these substrates. It was demonstrated that the proposed method can reach a high accuracy on synthetic data, which confirms that the chosen neural network architecture, as well as the implicit three-dimensional representation through a level-set function, are well suited for the given task. Furthermore, by the reconstruction of experimental cases from different fluid mechanical regimes, it was shown that the neural network was able to learn a versatile model of the involved two-phase flow phenomena, that even allows for topological changes.” from [DKS24, Section 4]

“In particular, the in-plane component of the reconstruction reaches an almost identical agreement to the ground truth, as the contour of the shadowgraph presents a strong basis for the in-plane reconstruction. The out-of-plane reconstruction relies on the learned droplet geometry, in addition to the three-dimensional information encoded in the glare points, which is more sparsely distributed across the droplet interface. Consequently, the uncertainty of the reconstruction is higher in the out-of-plane direction. The successful reconstruction of adhering droplets at much higher flow velocities and greater deformation than in the training data shows that the method can effectively handle and predict new, unknown shapes.” from [DHK<sup>+</sup>25, Section 4] The accurate in-plane reconstruction for droplet impingement on unknown substrates with gradients in wettability that lead to noticeably different shapes of the gas-liquid interface elucidates that the image features, in particular the shadowgraph contour, facilitate the generalization to unknown shapes of the gas-liquid interface. The prior knowledge obtained from the numerical training data leads to a bias of the trained droplet model towards the reconstruction of droplet geometries similar to the training data distribution in the out-of-plane reconstruction. Therefore, specific training data that considers the expected symmetries and variation of the gas-liquid interface in the experiments yields the most accurate results, while the extrapolation to different droplet geometries is possible, yet at a reduced accuracy. “It can be concluded that the combination of the learned geometry and the depth encoding by the glare points results in a robust and flexible model of the gas-liquid interface dynamics that can be used for the extrapolation to different flow conditions outside of the training data distribution.” from [DHK<sup>+</sup>25, Section 4]

“The successful reconstruction of obscured regions in the input images indicates that the neural network leverages the learned knowledge of the droplet dynamics from training data that is based on numerical simulation to fill in missing information in a physically reasonable way. It can be concluded that training the neural network with limited numerical data – in our case two direct numerical simulations [for the reconstruction of droplet impingement and only one simulation case for the reconstruction of adhering droplets] – already enables the network for a physically correct reconstruction. This highlights the advantage of the proposed data-driven method over

conventional approaches, which furthermore can be improved by more data and expanded to a broader range of applications by training on new regimes.” from [DKS24, Section 4]

“The proposed method leverages the planar symmetry of the setup [for adhering droplets], simplifying the depth prediction along the streamwise direction. Concurrently, the more significant deformation of the droplet in spanwise direction is imaged directly via the shadowgraph contour, thus reserving the previously mentioned high in-plane accuracy for the direction that undergoes more significant deformation. The implicit representation of the interface by the neural network allows for the reconstruction at a fine resolution, while the training on the results of numerical simulation facilitates a high fidelity of the reconstructed contours to the underlying physics. Therefore, the reconstructed gas-liquid interfaces are both spatially accurate and smooth, which makes the proposed method well-suited for distortion correction of PIV images, as the accurate reconstruction of the local curvature of the interface is important for the correct calculation of the refraction that causes the distorted velocity fields. In comparison to the current state-of-the-art approach for distortion correction, in which rotational symmetry of the gas-liquid interface is assumed, the presented approach becomes increasingly beneficial towards higher external flow velocities that result in larger degrees of droplet deformation and thus non-axis-symmetrical droplet shapes.” from [DHK<sup>+</sup>25, Section 4] The suitability of the reconstructed gas-liquid interfaces for distortion correction in PIV measurements should be investigated experimentally in future works. The ray tracing approach proposed by Kang et al. [107] and Minor et al. [150] shows promise for the reconstruction of the original particle distribution by tracing the approximately reciprocal light path from the distorted images, through the reconstructed gas-liquid interface, to the position of the light sheet. Preliminary testing within the developed synthetic image generation setup, utilizing physically-based rendering, has confirmed the feasibility of this correction approach.

“Further applications for the proposed monocular reconstruction approach are experiments where a complete and continuous imaging of the gas-liquid interface cannot be guaranteed at all times. The volumetric reconstruction of the gas-liquid interface for droplets impacting at an angle [...] appears straight-forward, while splashing droplets and sprays appear possible, as indicated by the successful reconstruction of secondary droplets. Furthermore, it is hypothesized on the grounds of the present findings that the described approach is likewise feasible for the reconstruction of gas bubbles in liquids since theoretical considerations according to van de Hulst [232] identified  $\theta \approx 78.5^\circ$  as a possible scattering angle for unambiguous  $p_{gp} = 0$  glare points.” from [DKS24, Section 4]

“The addition of training data that represents the deformed interface of adhering droplets at higher external flow velocities is expected to increase the accuracy of the reconstruction. Furthermore, the bias of the network towards certain geometries can most likely be reduced by a greater variation in the training dataset. Thereby, the flexibility and robustness of the trained model could be further enhanced. Further improvements in the reconstruction accuracy can be expected from an increased resolution of the numerical simulation that is used to source the training data. The effect is two-fold, on the one hand, more accurate three-dimensional ground truth data facilitates the learning of the droplet geometry and on the other hand, more finely resolved surface meshes would increase the quality of the synthetic images rendered by ray tracing.” from [DHK<sup>+</sup>25, Section 4]

Currently, the temporal dynamics of the numerical training data are not yet exploited. For this purpose, neural networks designed for processing image sequences, which leverage the temporal relationships in data structures, could be implemented in the neural network architecture. For example, the backbone feature extraction network could be expanded with convolutional long short-term memory networks [99] or spatio-temporal transformer networks [262]. Thereby, the additional

information that the three-dimensional shape of the gas-liquid interface changes smoothly in time can be used to regularize the reconstruction, which potentially allows for a higher accuracy.





## 7 Physics-informed learning for spatio-temporal droplet reconstruction

In this chapter a physics-informed reconstruction technique based on the PIFu approach [191] and the integration of prior knowledge by the governing equations of two-phase flows is developed aiming to improve the accuracy of the reconstruction and obtain additional 3D field information on velocity and pressure. For that purpose, first, related literature on the application of physics-informed neural networks (PINNs) [176] for forward and inverse two-phase flow problems is reviewed. Subsequently, the methodology for the incorporation of physics-informed learning techniques into the reconstruction framework is developed and evaluated in a comparative study with the baseline data-driven approach. The results discussed in this chapter have previously been published in the preprint *PINNs4Drops: Convolutional feature-enhanced physics-informed neural networks for reconstructing two-phase flows* by Dreisbach *et al.* [DKK<sup>+</sup>24].

### 7.1 Physics-informed neural networks

As elaborated in further detail in Section 2.4.4, “PINNs have emerged as a powerful framework for solving complex problems by seamlessly integrating governing equations into the data-driven approach. Their expressivity in approximating non-linear functions, combined with their unified framework for solving both forward and inverse problems, makes the PINN approach particularly well-suited for tackling fluid mechanical challenges [28, 204, 141, 116, 59].” from [DKK<sup>+</sup>24, Section 1] In the following section, the most significant advances in PINNs for solving fluid mechanical problems and recent progress in two-phase flows will be discussed to lay the theoretical grounds for the development of physics-informed neural networks for the spatio-temporal reconstruction of gas-liquid interfaces.

#### 7.1.1 PINNs for fluid mechanics

“In the introductory publication of PINNs, Raissi *et al.* [176] already investigated their application for encoding the continuity equation and the non-dimensional incompressible Navier-Stokes equations (NSE) to solve two types of inverse problems related to the flow in the wake of a circular cylinder. The authors showed that PINNs can be employed for the data-driven discovery of the viscosity in unknown fluids and the inference of the unknown hidden states, particularly, the continuous pressure field from scattered velocity observations.

Jin *et al.* [103] employed PINNs to simulate a 3D turbulent channel flow at  $Re_\tau = 1,000$  by training a neural network on the NSE with specified velocity boundary and initial conditions. The authors addressed the challenge of ill-posed problems arising from partially missing or noisy boundary conditions, demonstrating that PINNs can still yield accurate solutions under such circumstances. Similarly, Xu *et al.* [254] extended the application of PINNs to identify the artificial viscosity in the parametric turbulent incompressible NSE, leveraging a dataset from turbulent flow simulations. Furthermore, the authors demonstrated that PINNs can accurately reconstruct both

laminar and turbulent flow fields in regions of the domain where no measurement data is available, addressing challenges such as low fidelity or experimental limitations. Expanding on this capability, Cai et al. [28] employed PINNs to reconstruct the 3D velocity and pressure fields in the unsteady wake flow past a circular cylinder at  $Re = 200$  from 2D two-component (2D2C) velocity observations. Motivated by the potential application in the post-processing of particle image velocimetry measurements, the authors showed that as few as four cross-planes of 2D2C data are sufficient to accurately reconstruct the continuous 3D velocity and pressure fields, even without initial or boundary conditions. Raissi et al. [177] demonstrated the ability of PINNs to predict quantitatively accurate continuous 3D velocity and pressure fields from measurements of the tracer particle concentration. This was achieved without requiring any data, boundary, or initial conditions for the quantities of interest by training PINNs on the NSE, the continuity equation, and the transport equation for the passive scalar of the concentration. Additionally, the authors showed that PINNs could estimate derived quantities, such as shear forces and vorticity fields since exact gradients of the solution can be calculated at all points in the domain and at any desired resolution. Building on this concept, Cai et al. [29] extended the hidden fluid mechanics approach [177] to reconstruct continuous 3D velocity and pressure fields from 3D temperature measurements. By training PINNs on the Boussinesq approximation of the incompressible Navier-Stokes equations, the heat transfer equation, and the measured temperature data, they successfully reconstructed the buoyancy-driven flow over an espresso cup, showcasing the versatility of PINNs in handling diverse fluid mechanics problems. Furthermore, the authors demonstrated that PINNs are robust to noise and low-resolution training data in both the spatial and temporal domains, enabling the reconstruction of latent variables at a higher spatio-temporal resolution than the experimental data. Mao et al. [141] applied PINNs to compressible 2D flows governed by the Euler equations to reconstruct velocity, pressure, and density field in high-speed aerodynamic flows, using measurements of the density gradient in analogy to Schlieren experiments. The authors highlighted the importance of the spatial distribution of the sampling points, particularly for problems involving discontinuities in the flow fields. To address this, the authors proposed adaptive sampling strategies that dynamically track the discontinuities during training, increasing sampling density in critical regions based on the residuals of the PDEs or the gradients of the predicted solution. Building on these advancements, Wang and Perdikaris [242] introduced a multi-network PINNs framework to tackle the partial differential equations (PDEs) involving dynamic interfaces and phase change phenomena. In their approach, one neural network is trained to approximate the temperature field, while a second network parameterized the unknown and moving interfaces. These networks were coupled through a physics-informed network that encodes the corresponding heat transfer equation.” from [DKK<sup>+</sup>24, Section 1]

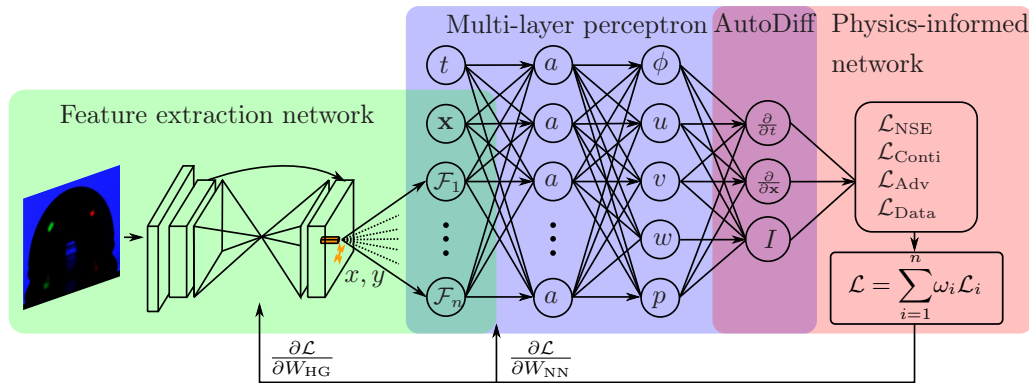
## 7.1.2 PINNs for two-phase flows

“The application of PINNs for two-phase flows necessitates a dedicated treatment of the interface, as the approximation of the discontinuity at the interface can result in steep gradients in the solution and locally high errors, potentially hindering optimization and even preventing convergence of the method. To address these challenges, Buhendwa et al. [21] extended the concept of the hidden fluid mechanics framework to predict continuous 2D velocity and pressure fields for incompressible two-phase flows from auxiliary information about interface position. Following the Volume of Fluid (VoF) approach, the PINNs are trained on the residual of the dimensionless single-field Navier-Stokes equations, the continuity equation, and the advection equation for the volume fraction  $\alpha$ , alongside a data loss term for the interface position. By implementing local refinement of the residual and data points near the interface, the optimization process was effectively focused on the region of the domain with the highest residuals, significantly stabilizing learning and reducing

solution errors [138]. Moreover, appropriate weighting of the loss terms, particularly the momentum equation residuals, was crucial to prevent divergence of the optimization caused by abrupt changes in the magnitude of the surface tension term. Qiu et al. [173] developed PINNs using the phase-field method to simulate 2D incompressible two-phase flows with large density ratios up to 1,000 between the two phases. The authors noted that the VoF approach as proposed by Buhendwa et al. [21] has potential stability issues at high density ratios due to the non-zero divergence of the velocity field at the interface. Previously, Buhendwa et al. [21] successfully addressed similar two-phase flows with density ratio ratios up to 10. Building on this, Qiu et al. [173] trained PINNs on a physics-informed loss incorporating the residuals of the continuity equation, the momentum equations as well as the Cahn-Hilliard equation to capture the interface evolution. This was combined with loss terms for the initial and boundary conditions for both the velocity and phase indicator variable. Expanding the scope of interfacial problems, Chen et al. [33] employed PINNs to solve the coupled Cahn-Hilliard and Allen-Cahn equations within the context of the phase-field method to simulate interfacial problems related to various phenomena related to electrochemical corrosion. Dynamic interfaces were addressed using adaptive sampling strategies based on the residuals of the PDEs or the gradients of the solution, as suggested by Mao et al. [141], which significantly enhanced prediction accuracy. Furthermore, an adaptive weighting scheme based on the convergence rate of the individual loss terms was employed to balance the contributions of the coupled PDEs and the initial and boundary conditions [244]. This approach significantly improved the accuracy for predicting the location of the dynamic interface.” from [DKK<sup>+</sup>24, Section 1]

## 7.2 Methodology

The previously employed data-driven reconstruction technique based on convolutional neural networks [191] allows for the accurate prediction of the three-dimensional gas-liquid interface of various droplet flows from monocular optical experiments as elaborated in Chapter 6. While the supervised learning approach facilitates the inclusion of prior knowledge from the numerical simulation, the underlying laws of physics governing the two-phase flow dynamics are only implicitly learned from



**Figure 7.1:** Schematic of the Convolutional Feature-Enhanced PINN Framework architecture. The proposed *PINNs4Drops* framework predicts the three-dimensional gas-liquid interface and the velocity and pressure distributions. Initially, glare-point shadowgraphy images are processed using a convolutional hourglass network, which extracts pixel-aligned features  $\mathcal{F}_i(x, y)$  from the input image at the pixel location  $(x, y)$  on the image plane. These extracted features, along with the temporal coordinate  $t$  and the spatial coordinates  $\mathbf{x}$ , are provided as inputs to an MLP. The MLP predicts the phase distribution  $\phi$ , the three components of the velocity vector  $\mathbf{u} = (u, v, w)^T$ , and the pressure  $p$  at the spatio-temporal coordinates  $(x, y, z, t)$ . The loss function comprises a data loss term and physics-informed loss terms. The physics-informed loss terms enforce the governing equations and are defined as the MSE of the following residuals:  $\mathcal{L}_{\text{Conti}}$ , enforcing the continuity equation;  $\mathcal{L}_{\text{Adv}}$ , representing the advection equation for phase distribution  $\phi$ ; and  $\mathcal{L}_{\text{NSE},j}$ , representing the Navier-Stokes equations for momentum conservation, with  $j = (x, y, z)$  indicating the spatial components. Figure adapted from [DKK<sup>+</sup>24].

the numerical data by the neural network. Furthermore, the reconstruction of individual snapshots does not consider the temporal coherence of the gas-liquid interface dynamics, which leads to an unphysical fluctuation of the reconstruction in time. In order to incorporate both these untapped sources of prior knowledge into the reconstruction framework, the unsteady governing equations are integrated with the data-driven reconstruction technique following the PINN approach with the aim of further improving the accuracy of the reconstruction. Such an approach is anticipated to enhance the learning of spatio-temporal droplet dynamics, as the optimization of the neural network is constrained to a physically reasonable solution space. Specifically, “the Pixel-Aligned Implicit Function (PIFu)-approach [191] for the volumetric reconstruction of the gas-liquid interface is extended by the integration of physics-informed losses by which the single-field two-phase formulation of the Navier-Stokes equations, the continuity equation and the equation for the interface evolution are encoded in the neural network. The continuity and Navier-Stokes equations are harnessed to learn the dynamics of the two-phase flow, while the interface evolution equation is employed to couple the velocity field to the distribution of the two phases and, consequently, the location of the gas-liquid interface. The representation of the gas-liquid interface as a continuous implicit function through a neural network in the PIFu-approach is more akin to diffuse interface methods than sharp interface methods [*cp.* Section 2.2.4]. Moreover, the physically sound modeling of the interface dynamics in the phase-field approach renders the method promising for accurate predictions of two-phase flows at high density and viscosity ratios, as demonstrated recently by Qiu et al. [173]. Therefore, the convective Cahn-Hilliard equation (eq. 2.29) appears as a highly suitable choice for representing the interface evolution. However, computing the physics-informed loss derived from the Cahn-Hilliard equation involves fourth-order derivatives, which coupled with the complexity of the three-dimensional domain, lead to high computational demands. Furthermore, the repeated calculation of gradients through automatic differentiation to obtain the fourth-order derivatives leads to the accumulation of potential errors. In contrast, the VoF approach offers a simplified alternative by representing the interface evolution through the transport equation for the volume fraction, which is purely convective and thus involves only first-order derivatives. The successful application of PINNs based on the VoF method to inverse two-phase flow problems, as demonstrated by Buhendwa et al. [21], highlights the feasibility and efficiency of this approach for tackling such problems. However, the applicability of PINNs in the VoF formulation for two-phase flows at high density ratios remains unverified. To address this, we will develop two variants of PINNs, one based on the phase-field method (*PF-PINNs-v1*) and the other on VoF (*VoF-PINNs*). These variants will share the same neural network architecture and will be comparatively evaluated to assess their performance.

The computation of the physics-informed loss terms using automatic differentiation requires predictions of the three-dimensional velocity and pressure fields, necessitating modifications to the network architecture. Specifically, the output layer of the network must be extended to predict three velocity components  $u, v$ , and  $w$ , as well as the pressure  $p$ . Additionally, the input layer must be adapted to include the spatio-temporal coordinates  $(x, y, z)$ , and time  $t$  to enable the computation of the derivatives of the predicted quantities.

Figure 7.1 illustrates the network architecture of the proposed PINN, which consists of three major components, namely, the feature extraction network (a CNN), the neural network (an multi-layer perception (MLP)), and the physics-informed network. First, the glare-point shadowgraphy images are processed using a convolutional hourglass network [160], which extracts pixel-aligned features  $\mathcal{F}_i$  from the input image at the location  $x, y$  on the image plane. The extracted image features  $\mathcal{F}(x, y)$  along with their corresponding spatial coordinates  $x, y$  are forwarded to the MLP. Additionally, the temporal coordinate  $t$  and the spatial coordinate  $z$  are given as inputs to the MLP. On this basis, the MLP predicts the phase distribution  $\phi$ , the three components of the velocity vector  $\mathbf{u} = (u, v, w)^\top$ ,

as well as the pressure  $p$  at  $(x, y, z, t)$ . The residuals of the single-field two-phase Navier-Stokes, continuity, and interface evolution equations are computed using automatic differentiation applied to the predicted output with respect to the spatio-temporal input coordinates. These residuals form the basis for physics-informed loss terms, *i.e.*, the continuity equation ( $\mathcal{L}_{\text{Conti}}$ ), the advection equation governing the interface ( $\mathcal{L}_{\text{Adv}}$ ) and the Navier-Stokes momentum equations ( $\mathcal{L}_{\text{NSE},j}$ ), with  $j = (x, y, z)$ . Each loss term is defined as the mean squared error (MSE) of the respective residuals. In addition to these physics-informed losses, ground truth labels for the velocity and pressure fields, as well as surface meshes of the gas-liquid interface, are extracted from numerical simulations. These labels are used to define data loss terms for the predicted quantities  $(\phi, u, v, w, p)$ , computed as the MSE between the predictions and ground truth. The joint neural networks are trained on a composite loss representing the weighted sum of the physics-informed loss terms and the data loss terms

$$\begin{aligned} \mathcal{L} = \sum_{i=1}^n \omega_i \mathcal{L}_i = & \omega_\phi \mathcal{L}_\phi + \omega_u \mathcal{L}_u + \omega_v \mathcal{L}_v + \omega_w \mathcal{L}_w + \omega_p \mathcal{L}_p \\ & + \omega_{\text{Conti}} \mathcal{L}_{\text{Conti}} + \omega_{\text{Adv}} \mathcal{L}_{\text{Adv}} + \omega_{\text{NSE},x} \mathcal{L}_{\text{NSE},x} \\ & + \omega_{\text{NSE},y} \mathcal{L}_{\text{NSE},y} + \omega_{\text{NSE},z} \mathcal{L}_{\text{NSE},z}. \end{aligned} \quad (7.1)$$

The formulation of governing equations, dedicated sampling schemes, and particular design choices for the neural network architecture play a critical role in the successful application of PINNs for two-phase flow problems. The methods employed for the proposed PINNs will be detailed in the following sections.” from [DKK<sup>+</sup>24, Section 2]

### 7.2.1 Governing equations

“It has been shown that the computation of the physics-informed losses based on the dimensionless formulation of the Navier-Stokes equations results in moderately improved accuracy of the predicted flow quantities compared to the dimensional formulation [29]. Furthermore, the similar magnitude of the dimensionless PDE terms facilitates the balancing of different physics-informed loss components [41]. Therefore, the dimensionless single-field Navier-Stokes equations are employed

$$\begin{aligned} \rho^* \left( \frac{\partial \mathbf{u}^*}{\partial t^*} + (\mathbf{u}^* \cdot \nabla^*) \mathbf{u}^* \right) = & -\nabla^* p^* + \nabla^* \cdot \left( \frac{1}{Re} (\nabla^* \mathbf{u}^* + (\nabla^* \mathbf{u}^*)^\top) \right) \\ & + \frac{1}{We} \frac{\mathbf{f}_\gamma}{\gamma} + \rho^* \frac{1}{\mathbf{Fr}^2}, \end{aligned} \quad (7.2)$$

with the following relationships between dimensional and non-dimensional quantities for the velocity  $\mathbf{u} = \mathbf{u}^* u_r$ , the pressure  $p = p^* \rho_r u_r^2$ , the density  $\rho = \rho^* \rho_r$ , the spatial coordinates  $\mathbf{x} = \mathbf{x}^* L_r$ , and the time  $t = t^* L_r / u_r$ , as well as the Reynolds number  $Re = \rho_r u_r L_r / \mu$ , Weber number  $We = \rho_r u_r^2 L_r / \gamma$ , and Froude number  $\mathbf{Fr} = u_r / \sqrt{\mathbf{g} L_r}$ . The dynamic viscosity, surface tension coefficient, and gravitational acceleration are indicated by  $\mu$ ,  $\gamma$ , and  $\mathbf{g}$ , respectively. The reference quantities are the impact velocity  $u_r = u_0$  for the considered case of droplet impingement, the density of the liquid phase  $\rho_r = \rho_L$  and the reproduction scale  $L_r = R_p$  of the computational domain to the experiments, which ensures the correct scaling of the surface curvature required for the calculation of surface tension. Both versions of the PINNs employ the continuity equation for incompressible fluids [eq. 2.3].

Volume-of-Fluid-based Physics-Informed Neural Networks (*VoF-PINNs*) represent the interface evolution in two-phase flows using the Volume of Fluid (VoF) method, a well-established approach for capturing fluid interfaces. The transport equation for the volume fraction  $\alpha$  [eq. 2.21] is employed to describe the interface evolution, following the formulation of algebraic VoF approaches [86]. [...] The surface tension  $f_\gamma$  is modeled using the Continuum Surface Force (CSF) model [16] as a localized body force within the transition region of finite thickness at the interface [and represented by Equation 2.25] [...]. The mixture density  $\rho$  and viscosity  $\mu$  are determined by the arithmetic mean of the fluid properties in both phases [eq. 2.23] [...].

The phase-field version employs the convective Cahn-Hilliard equation [eq. 2.29] [27] to represent interface evolution [...]. According to the diffuse-interface theory [234] the phase separation and diffusion in two-phase flows are driven by the chemical potential at the interface  $\psi$ , which is derived as the variational derivative of the mixing energy with respect to the order parameter  $C$  [eq. 2.28] [...]. The value for [the capillary width]  $\epsilon$ , [which is proportional to the interface thickness] [...], needs to be chosen and is typically defined in relation to the characteristic macroscopic length scale of the flow. In the phase-field version, the continuum surface tension in the potential form [100] is employed [eq. 2.32] [...]. The mixture density  $\rho$  and viscosity  $\mu$  are determined [by eq. 2.31] [...].

In the phase-field PINNs, referred to as *PF-PINNs-v1*, the capillary width  $\epsilon$  is treated as a learnable parameter. In order to promote the convergence toward a thin interface, an additional loss term is introduced, defined as the Huber loss [96] between  $\epsilon$  and the value of  $\epsilon_1 = 2.2 \cdot 10^5$ , determined by Fink et al. [64]. The capillary width  $\epsilon$  is initialized with values ranging from 0.01 to 0.05 to study the influence of the parameter and the weight for the additional learnable interface loss is set to  $\omega_\epsilon = 100$ . To address potential issues arising from the fourth-order derivatives in the Cahn-Hilliard equation employed for the phase-field version of the PINNs, a second variant of phase-field PINNs is proposed. In this variant, referred to as *PF-PINNs-v2*, the chemical potential  $\psi$  is directly predicted using the neural network, in contrast to the computation of  $\psi$  from  $C$  in *PF-PINNs-v1*, reducing the highest order of derivatives to the second order. An additional loss term is derived from the residual of Equation 2.28, to ensure the consistency between the predicted phase distribution and the distribution of the chemical potential. This provides additional supervision for the learning of  $\psi$ . The weight for this additional identity loss of the chemical potential is set to  $\omega_\psi = 0.1$ ” from [DKK<sup>+</sup>24, Section 2.1]

## 7.2.2 Loss weighting

“Previous research has demonstrated the critical role of proper loss weighting to achieve the simultaneous convergence of all loss terms [21, 41, 244]. To address this, a combination of fixed loss weights and adaptive loss weighting is used to balance the data-driven and physics-informed loss terms during training. For the inverse problem of the flow field reconstruction in a flow with natural convection, larger relative weights for the data loss terms compared to the physics-informed losses have been shown to be beneficial [29]. Furthermore, the accurate prediction of the phase distribution is crucial for correctly computing the momentum equation, as the mixture density  $\rho_M$  significantly affects the momentum balance. Optimal weights for achieving the simultaneous convergence of all loss terms were determined experimentally as  $\omega_{\{u,v,w,p\}} = 10$ ,  $\omega_{\{\text{Conti}, \text{Adv}, \text{NSE}, x\}} = 1,000$ ,  $\omega_{\text{NSE}, y} = 1$  and  $\omega_{\text{NSE}, z} = 100$ . These weights yielded weighted physics-informed and data loss terms for  $(u, v, w, p)$  that ranged between one and three orders of magnitude lower than the weighted loss for the phase distribution. In addition to using fixed loss weights, the adaptive loss weighting scheme *SoftAdapt* [84] was employed to dynamically balance the loss terms during training. This simple loss weighting scheme is based on the relative convergence rate of the different loss terms, with inversely proportional loss weights determined dynamically. *SoftAdapt*



significantly improves the training dynamics and enhances the accuracy of the trained model, all while incurring minimal computational costs. Both fixed and adaptive loss weights were applied in a multiplicative manner. Given the importance of the phase distribution for the physics-informed loss terms, it is essential that the optimization process ensures the accurate convergence of the phase distribution prediction before introducing the physics-informed losses. Therefore, the neural network was initially trained only on the data loss terms until the data  $\mathcal{L}_\phi$  decreased below a specified threshold  $\mathcal{L}_{\phi,T}$ . Once this threshold was reached, the physics-informed losses were incorporated into the weighted composite loss. A threshold value  $\mathcal{L}_{\phi,T} = 0.03$  was determined to provide a sufficient convergence of the prediction for the phase distribution to obtain adequate physics-informed losses. Additionally, the data loss terms for  $(u, v, w, p)$  were gradually increased during the first 5,000 training iterations by an additional weighting factor, which scaled linearly from zero to one from training iteration zero to 5,000 to prioritize the learning of an accurate interface prediction in the early stages of training.” from [DKK<sup>+</sup>24, Section 2.2]

### 7.2.3 Sampling methods

“Previous studies have demonstrated that an adaptive refinement of the sampling points for the loss computation at the gas-liquid interface is essential for the successful application of PINNs to two-phase flow problems [21, 33]. This necessity arises from the steep gradients of the solution near the interface [141, 138], which require higher sampling density to accurately capture the physics in these regions. Therefore, a combination of adaptive and random sampling at a ratio of 16 : 1 is employed. The dense adaptive sampling near the interface facilitates the learning of an accurately localized prediction for the gas-liquid interface, while few random sampling points across the rest of the domain are required to prevent overfitting. Following the adaptive surface sampling method introduced by Saito et al. [191], random sampling from a normal distribution centered on the interface with a standard deviation  $\sigma = 3.9\%$  of the domain size is used to refine the sampling points near the interface. This corresponds to an adaptive sampling at a thickness of 0.107 mm near the interface in physical space or 5% of droplet diameter  $d_0 = 2.1$  mm, respectively. The data loss terms and the physics-informed loss terms are calculated on different sets of sampling points in order to effectively exploit the continuous nature of PINNs, which allows for the computation of residuals at any sampling point in the spatio-temporal domain. Thereby, prior knowledge from the numerical simulation is incorporated into the PINNs through supervised learning, while further residual sampling points at different locations in the domain provide supplementary supervision by the physics-informed losses. To ensure a fine resolution of the three-dimensional domain, 5,000 data points and 10,000 residual points are sampled at each time step. These points are re-sampled at every epoch to provide additional coverage of the spatio-temporal domain during the training.

Motivated by the residual-based attention scheme [1] and adaptive sampling methods based on the residuals of the PDEs [33, 141], a residual-based weighting of the residual points is employed to guide the optimization of the PINNs towards regions in the domain that are challenging to optimize. In the proposed approach, the residual points are additionally weighted by a factor  $\omega_k$  ranging from 0.8 to 1.2 based on the relative magnitude of the residual, calculated as

$$\omega_{k,i} = 0.8 + 0.4 \frac{|r_i|}{\max_i(|r_i|)}, \quad (7.3)$$

where  $r_i$  is the residual at sampling point  $i$ .” from [DKK<sup>+</sup>24, Section 2.3]



## 7.2.4 Neural network architecture

“As illustrated in Figure 7.1 the architecture of the PIFu neural network is adapted for the proposed PINNs. While the hourglass network for feature extraction can remain unchanged, the input and output layers, as well as the activation functions of the MLP are modified. Specifically, the MLP is comprised of 260 input nodes, with 256 nodes receiving the pixel-aligned features from the hourglass network and four nodes receiving the spatio-temporal coordinates  $(x, y, z, t)$ , four hidden layers with 1024, 512, 256, 128 neurons, respectively, and five output neurons for the prediction of the flow quantities  $(\phi, u, v, w, p)$ . The leaky ReLU activation function [255] used as the non-linearity for the hidden layers in the MLP is replaced by the infinitely differentiable hyperbolic tangent activation function to enable the computation of higher-order derivatives through automatic differentiation. Adaptive activation functions [101] introduce an additional learnable scaling coefficient  $n_s \cdot a_s$  to the activation function that regulates its slope and thus the sensibility to its inputs. The scaling coefficient consists of the fixed scale factor  $n_s$  and the adaptive activation coefficient  $a_s$ , which is optimized alongside the parameters of the neural network. Layer-wise adaptive activation functions have been found to improve the convergence rate and accuracy of PINNs across various types of problems [101], including two-phase flows [21]. Therefore, layer-wise adaptive activation functions are employed in the hidden layers with a scale factor  $n_s = 2$  and an initial value for the adaptive activation coefficient of  $a_s = 0.5$ .

In the output layer, different activation functions are chosen for each predicted quantity according to the range of possible values. The sigmoid activation function is used for the output neuron of the phase distribution  $\phi$  to confine the prediction within the physical bounds  $\alpha \in [0, 1]$  or  $C \in [-1, 1]$  for the volume fraction  $\alpha$  in the VoF version and the order parameter  $C$  in the phase-field version of the PINNs, respectively. The prediction of the order parameter  $C$  requires an additional scaling of the activation function by a factor of two. The output neurons for the prediction of the velocity components  $(u, v, w)$  employ linear activation functions, while the output neuron for the prediction of the pressure  $p$  features an exponential activation function, as suggested by Buhendwa et al. [21]. Skip connections [80] [(see Section 2.4.3)] are employed at each hidden layer of the MLP to propagate the information of the input feature vector and spatio-temporal coordinates to later layers in the network, which has been shown to improve the accuracy for both data-driven volumetric reconstruction [34, 191] and PINNs [35, 245].

The proposed PINNs are trained for eight epochs by the RMSProp optimizer [225] with a learning rate decay by a factor of ten at the beginning of epochs six and eight. The batch size is reduced to 1 due to the increased computational requirements of the PINNs and the initial learning rate is accordingly decreased to 0.0001. Data augmentation through random translation of the input images is employed in order to improve the generalization capability of the network. The augmentation of the input images by random scaling, previously used in the optimization of the data-driven reconstruction network cannot be used, as the developed PINNs are no longer scale invariant due to the surface tension term in the momentum equation, which depends on the curvature of the gas-liquid interface.” from [DKK<sup>+</sup>24, Section 2.4] It should be noted that the calculation of the derivatives in the physics-informed network does not consider the implicit dependence of the derivatives on the pixel-aligned feature embeddings computed by the hourglass network. The local pixel-aligned features are obtained by bilinear interpolation of the final feature maps in the hourglass network, as elaborated in Section 3.4. This interpolation is dependent on the location  $(x, y)$  of the sampling point in the image plane, which introduces an implicit dependence of the derivatives of the solution  $\mathbf{u}(\mathbf{x}, t)$  with respect to the spatial input coordinates  $\frac{\partial^n \mathbf{u}}{\partial \mathbf{x}^n}$  on the feature embeddings.

### 7.2.5 Evaluation metrics

“The performance of the PINNs is evaluated considering the reconstructed three-dimensional interface geometries, velocity and pressure field, as well as the availability of ground truth data. The following metrics [in addition to those defined in Subsection 6.1.4] are used for the evaluation:

- The  $L_1$  error of the predicted quantities  $q \in \{u, v, w, p\}$  is calculated as

$$L_{1,q} = \frac{1}{n} \sum_{i=1}^n |q_{\text{GT},i} - q_{\text{pred},i}|, \quad (7.4)$$

over all sampling points  $i$  for one reconstructed snapshot and averaged over all samples in the the validation dataset.

- The  $L_2$  error is calculated as

$$L_{2,q} = \sqrt{\frac{1}{n} \sum_{i=1}^n (q_{\text{GT},i} - q_{\text{pred},i})^2}. \quad (7.5)$$


The  $L_1$  and  $L_2$  errors are given in absolute terms or relative to the maximum  $\max_i(|q_{\text{GT},i}|)$  for a given validation sample.” from [DKK<sup>+</sup>24, Section 2.6]

## 7.3 Results and Discussion

“In the following, the three proposed versions of the PINNs are first validated considering their training dynamics, with a particular focus on the convergence of the various loss terms. Afterward, the predictive accuracy for the three-dimensional gas-liquid interface and the extended capacity for flow topology prediction are investigated by means of synthetic validation data. Finally, the results obtained by the application of the developed PINNs for the reconstruction of images recorded in the experiments are discussed. The three candidate models, *VoF-PINNs*, *PF-PINNs-v1*, and *PF-PINNs-v2* were trained on the [DFS] dataset that featured droplet impingement on flat and structured substrates and synthetically rendered images with glare points [(see Section 5.4 for details)] with the same set of hyperparameters, aside from the additional learnable interface thickness and chemical potential of the phase-field versions.” from [DKK<sup>+</sup>24, Section 3] Additionally, the purely data-driven version of the PIFu neural network is re-trained using the same training hyperparameters as the PINN variants for validation purposes, which will be referenced as *DFS-GP2* in the following.

### 7.3.1 Validation on synthetic data

**Training dynamics** “The fixed weights of the data and physics-informed loss terms were iteratively tuned in order to achieve a simultaneous convergence of all loss terms, which yielded optimal training results for  $\omega_{\{u,v,w,p\}} = 10$ ,  $\omega_{\{\text{Conti}, \text{Adv}, \text{NSE}, x\}} = 1000$ ,  $\omega_{\text{NSE}, y} = 1$ , and  $\omega_{\text{NSE}, z} = 100$ . The evolution of the various unweighted loss terms during the training of the *VoF-PINNs* for the aforementioned combination of weights are plotted in Figure 7.2. The other two models, *PF-PINNs-v1* and *PF-PINNs-v2*, show similar dynamics of losses during training. As can be seen, all individual loss terms are converging, which indicates the successful simultaneous learning of the phase distribution, velocities, and pressure by the neural network. The similar convergence rate of




**Figure 7.2:** Evolution of the loss terms during training of the *VoF-PINNs*. The data loss terms are indicated by the solid lines, while the physics-informed losses are indicated by the dashed lines, and total loss is indicated by the dashed-dotted black line. Figure adopted from [DKK<sup>+</sup>24].

the loss terms validates the choice of loss weights. Furthermore, the difference in the magnitude of the data losses for the velocity components and pressure in comparison to the phase distribution and the even lower physics-informed loss terms highlights the necessity of loss weighting. During the initial phase of the training, the data loss for the phase distribution  $\mathcal{L}_\phi$  only decays slowly, which is followed by a sudden drop that can be attributed to the learning of the gas-liquid interface. A further drop can be observed at the start of epoch six, at which the first learning rate decay occurs. The first drop of  $\mathcal{L}_\phi$  is correlated with a sudden increase in the physics-informed loss terms. Similar training dynamics were already observed by Buhendwa et al. [21], who reported that the sudden learning of the interface location leads to a rapid increase in the magnitude of the gradients in the phase distribution at the interface and in turn, an increase in the magnitude of the surface tension term in the momentum equation. Furthermore, Buhendwa et al. [21] found that the sudden increase of the physics-informed losses destabilizes the optimization and can cause complete divergence. The reason is that the physics-informed loss terms offer only limited guidance when the interface is not accurately learned and may even mislead the optimization of the neural network. This is particularly problematic for the considered water-air flow, where an erroneous prediction of the occupancy field introduces large errors in the momentum equation due to the high fluid's density ratio. In the proposed PINNs, this issue is mitigated by ensuring that the prediction of the gas-liquid interface and the flow topology have sufficiently converged before considering the physics-informed loss terms. This is successfully achieved through purely data-driven training until the loss for the phase distribution reaches the threshold value  $\mathcal{L}_{\phi,T} = 0.03$ , as indicated by the monotonic decrease of all loss terms.

The gradual introduction of the data loss terms for  $u, v, w$ , and  $p$  further facilitates the learning of an accurate interface prediction by focusing the optimization on  $\mathcal{L}_\phi$ . The proposed loss weighting methods significantly reduced the initial plateau of the phase distribution loss, achieving faster convergence and a lower final  $\mathcal{L}_\phi$ . This in turn, significantly enhanced the overall training dynamics of the PINNs and has led to an increased accuracy of the prediction for the interface location, without impeding the learning of the flow topology. The second drop of  $\mathcal{L}_\phi$ , unlike the first, correlates with a sudden decrease of the physics-informed loss terms, which indicates that the location of the gas-liquid interface, as well as the velocity field, are learned accurately enough at this later stage of training so that further refinement of the interface location results in an improved adherence to the governing equations. Both phase-field versions of the PINNs converged

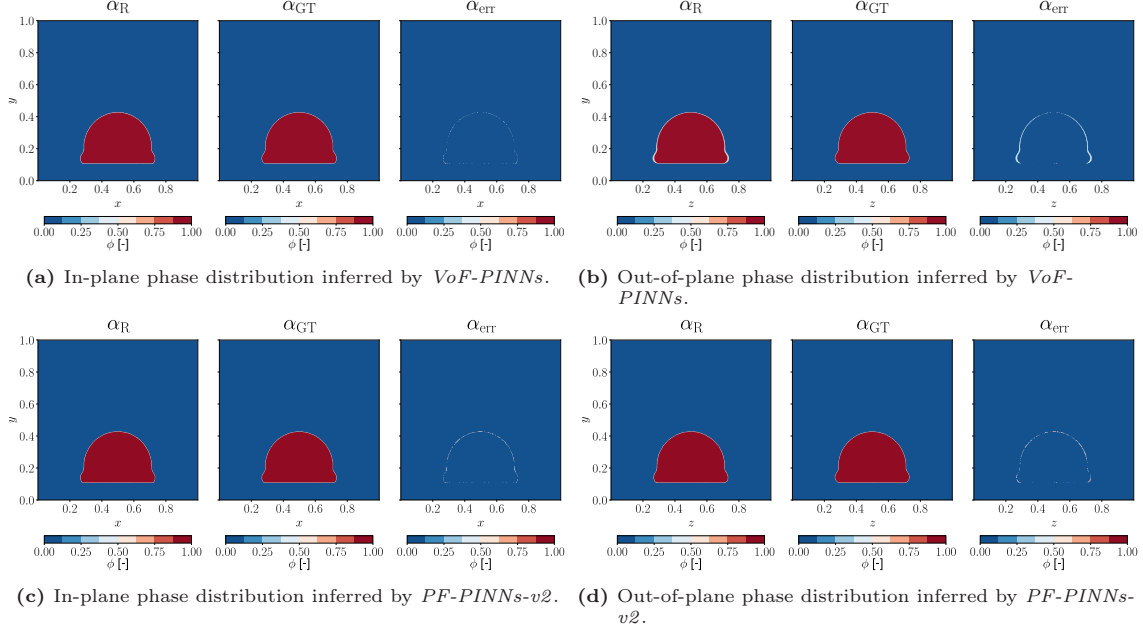
to similarly low losses in comparison to the VoF version except for  $\mathcal{L}_\phi$ . Specifically, *PF-PINNs-v1* with an initial value of  $\epsilon_0 = 0.01$ , remained at a considerably higher  $\mathcal{L}_\phi = 0.0358$  in comparison to *VoF-PINNs* with  $\mathcal{L}_\phi = 0.0092$ , while *PF-PINNs-v2* at  $\epsilon_0 = 0.01$  reached  $\mathcal{L}_\phi = 0.0135$ . These results suggest that the proposed VoF approach is more appropriate for training PINNs aimed at the accurate reconstruction of the gas-liquid interface in the considered two-phase droplet flows. Furthermore, the lower  $\mathcal{L}_\phi$  achieved by *PF-PINNs-v2* indicates that the separate prediction of chemical potential is beneficial for learning the phase distribution in phase-field PINNs. The cause for the better performance of *PF-PINNs-v2* might be an improvement of the training dynamics by avoiding the fourth-order derivative in the Cahn-Hilliard equation. The learnable interface thickness of *PF-PINNs-v2* remains close to the initial values of  $\epsilon_0 = 0.01$  and  $\epsilon_0 = 0.05$  throughout the training and only marginally decreases from  $\epsilon = 0.01$  to  $\epsilon = 0.0093$  for *PF-PINNs-v1* towards the end of the training. It was found that a higher weighting of the interface loss term leads to lower values of  $\epsilon$ , however, at the cost of significantly reduced convergence for the other loss terms, which resulted in a degraded accuracy of the interface reconstruction. Similarly, an increased initial value of  $\epsilon_0$  from 0.01 to 0.05 led to a significantly improved convergence of  $\mathcal{L}_\phi$  for *PF-PINNs-v2*, reaching similarly low values of  $\mathcal{L}_\phi = 0.0093$  at the end of the training in comparison to *VoF-PINNs*. Consequently, a more diffuse interface of the phase-field PINNs was found to be beneficial for their optimization.

As the physics-informed loss terms are calculated on a different set of sampling points than the data losses, the low residuals of the governing equations indicate that a continuous and accurate function approximation in 3D space was learned by the neural network for all predicted quantities. In order to further elucidate to which extent the governing equations were learned by the PINNs, a version of the neural network was trained only through the data loss terms, while the development of the residuals was tracked during training. It was found that the PINNs achieved MSE errors for the residuals of the governing equations that were multiple orders of magnitude lower (25 to 9,600 times) in comparison to the network only trained on the data loss terms. In particular, the residuals for the interface evolution equation were reduced by almost a factor of  $10^4$  by the PINNs. These results confirm the successful encoding of the underlying governing equations in the neural network by the physics-informed losses in the PINNs.” from [DKK<sup>+</sup>24, Section 3.1]




**Figure 7.3:** Average 3D-IOU of the reconstruction for the validation datasets during training of *VoF-PINNs*, as well as *PF-PINNs-v1*, and *PF-PINNs-v2* with different capillary width  $\epsilon_0$  in comparison to the baseline [*DFS-GP2*]. Figure adopted from [DKK<sup>+</sup>24].

**Gas-liquid interface prediction** “The reconstruction accuracy of the developed PINNs is evaluated quantitatively by the results for the reconstructed gas-liquid interfaces on the synthetic validation



**Figure 7.4:** Phase distribution predicted by *VoF-PINNs* (top) and *PF-PINNs-v2* with  $\epsilon_0 = 0.01$  (bottom) along the center planes of the droplet in the in-plane and out-of-plane directions (left) in comparison to the ground truth simulation data (middle), and the absolute error between the prediction and the ground truth (right). Figure adopted from [DKK<sup>+</sup>24].

dataset. Figure 7.3 shows the evolution of the 3D-IOU on the synthetic validation data during the training of the PINNs in comparison to the purely data-driven baseline network [*DFS-GP2*]. As can be seen, *VoF-PINNs* reaches a higher reconstruction accuracy in comparison to the data-driven baseline [*DFS-GP2*], while *PF-PINNs-v1* and *PF-PINNs-v2* do not yield an improvement over the baseline. More specifically, *VoF-PINNs* reaches an accuracy of 3D-IOU=0.967, which translates to a 0.9% improvement over the data-driven baseline *DFS-GP2* with 3D-IOU=0.958 [and a 1.24% improvement over *DFS-GP*], which is significant considering the proximity to optimal reconstruction results at 3D-IOU=1. By means of an ablation study for the residual-based weighting of the sampling points for the calculation of the physics-informed losses and the layer-wise adaptive activation functions, we found that both measures only marginally improved the accuracy of the interface prediction. Consequently, the gains in reconstruction accuracy can solely be attributed to the introduction of physics-informed losses to the reconstruction framework. Moreover, the different training dynamics between the VOF variant and the two phase-field versions of the PINNs are reflected in the reconstruction accuracy of the gas-liquid interface. The accuracy of *PF-PINNs-v2* with  $\epsilon_0 = 0.01$  sharply drops between training epochs two and three, which correlates to the convergence of  $\mathcal{L}_\phi$  below  $\mathcal{L}_{\phi,T} = 0.03$ , which marks the point in training at which the physics-informed losses are introduced. A potential cause might be the onset of the identity loss for the chemical potential that couples the phase distribution to the chemical potential, which at that point in the training is still in the condition of the random initialization. Consequently, the randomness of distribution and magnitude of the chemical potential might introduce an erroneous objective for the phase distribution through the identity loss, as long as the prediction for the chemical potential is not yet sufficiently converged. An increase of the capillary width from  $\epsilon_0 = 0.01$  to  $\epsilon_0 = 0.05$  helps to mitigate this issue, leading to a substantial improvement in reconstruction accuracy, as shown in Figure 7.3. However, the previously observed phenomenon remains to some extent, as evidenced by the plateau in accuracy between epochs one and four. Further improvement could be reached by an additional data loss for the chemical potential, as the prediction of the chemical potential could be optimized prior to the onset of the physics-informed losses.



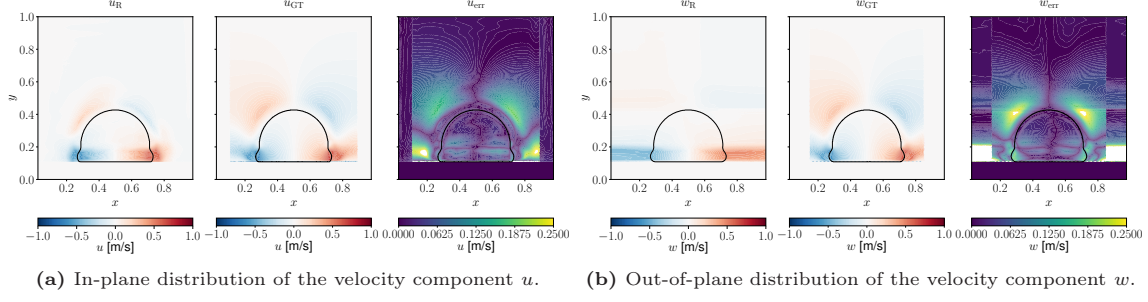
(a) Synthetic input image rendered from a simulated droplet during impingement. (b) Predicted velocity vector field and pressure contours along the center plane of the droplet (left) in comparison to the ground truth simulation data (right). The contour of the gas-liquid interface is indicated by the black solid line.

**Figure 7.5:** Snapshot from the validation data set (a) and respective in-plane reconstruction (b) by *VoF-PINNs*. Figure adopted from [DKK<sup>+</sup>24].

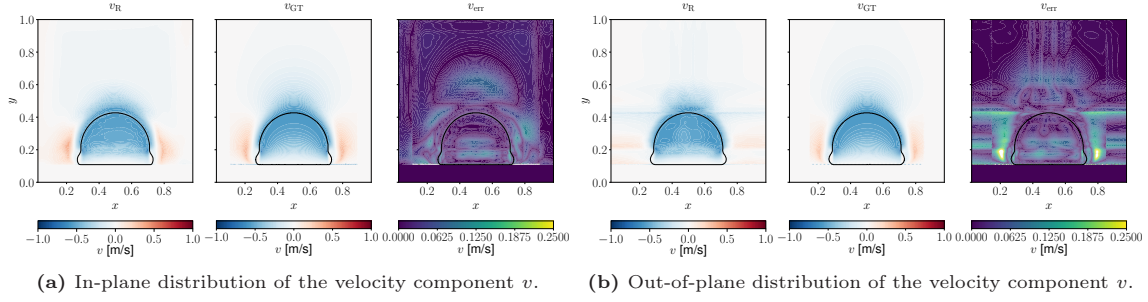
The impact of the different formulations of the governing equations of the VOF and phase-field PINNs on the prediction of the gas-liquid interface is further investigated by the comparison of the spatial distribution of the phases in the predictions returned by *VoF-PINNs* and *PF-PINNs-v2*. Figure 7.4 shows the predicted phase distribution in the in-plane and out-of-plane direction for one sample from the validation subset of the *DFS* dataset in comparison to the ground truth and the respective error distributions. The comparison of the error distributions reveals that the prediction of *PF-PINNs-v2* with  $\epsilon_0 = 0.01$  features a less diffuse edge in the out-of-plane direction compared to *VoF-PINNs*, while both models predict a rather sharp in-plane phase distribution. These results indicate that the in-plane reconstruction adheres closely to the shadowgraph contour, while the out-of-plane reconstruction heavily relies on the trained model of droplet dynamics. For  $\epsilon_0 = 0.05$  *PF-PINNs-v2* the out-of-plane prediction has a similar degree of diffusion compared to *VoF-PINNs*. The different diffusivity of the out-of-plane phase distribution elucidates the influence of the PINNs on the trained droplet model. A sufficiently small value for  $\epsilon$  in the phase-field approach encourages the learning of a more accurately localized gas-liquid interface in comparison to the algebraic VOF approach, in which the interface thickness is not explicitly considered, but instead, a sharp interface is assumed while a certain numerical diffusion is accepted.” from [DKK<sup>+</sup>24, Section 3.1]

**Velocity and pressure field prediction** The extension of the reconstruction framework by the proposed PINNs provides the additional capability for the prediction of the three-dimensional velocity and pressure fields in both phases of the two-phase droplet flow. “The accuracy of the velocity and pressure prediction by *VoF-PINNs* is evaluated on the validation dataset by the comparison to the ground truth velocity and pressure data obtained by direct numerical simulation in the phase-field method conducted by Fink et al. [64] [(see Section 5.4 for further details)]. Figure 7.5b shows the predicted in-plane pressure and velocity fields in the center plane of the droplet (left) in comparison to the ground truth (right) for one snapshot of the synthetic validation data set displayed in Figure 7.5a. As can be seen, there is a good topological agreement of the predicted pressure and velocity fields with the ground truth. Furthermore, the prediction of the pressure reaches a good quantitative agreement in both the in-plane and out-of-plane directions. The predicted velocity field was found to be accurately reconstructed in the liquid phase but





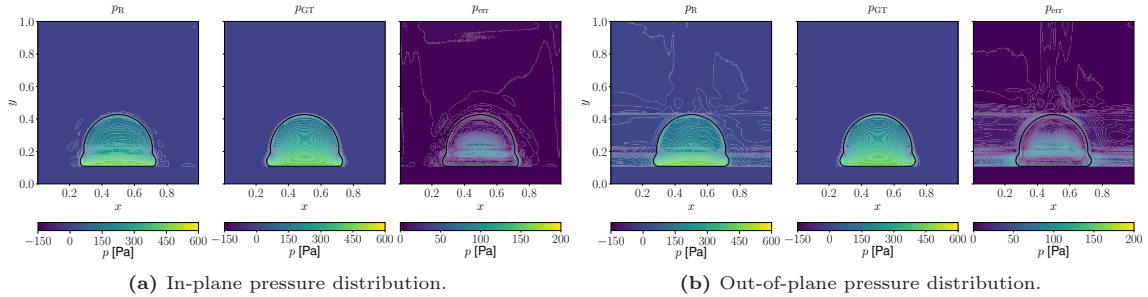
**Figure 7.6:** Prediction of *VoF-PINNs* for the horizontal components of the velocity vector along the center planes of the droplet in the in-plane and out-of-plane directions (left) in comparison to the ground truth simulation data (middle), and the absolute error between the prediction and the ground truth (right). Figure adopted from [DKK<sup>+</sup>24].



**Figure 7.7:** Prediction of *VoF-PINNs* for the vertical component of the velocity vector along the center planes of the droplet in the in-plane and out-of-plane directions (left) in comparison to the ground truth simulation data (middle), and the absolute error between the prediction and the ground truth (right). Figure adopted from [DKK<sup>+</sup>24].

exhibited larger deviations in the gaseous phase. Similar results were obtained for the prediction of the velocity in the out-of-plane direction, as illustrated by Appendix Figure C.5b.

The results for the velocity and pressure prediction are illustrated in more detail in Figures 7.6 to 7.8, showing the prediction for the horizontal velocity components  $u$  and  $w$  in Figure 7.6, the vertical velocity component  $v$  in Figure 7.7 and the pressure in Figure 7.8 in comparison to the ground truth data and the absolute error distribution for the previous sample of the validation data. As can be seen, all three components of the predicted velocity field show an overall good topological agreement with the ground truth data. Due to the symmetry of the problem, the out-of-plane prediction for velocity component  $u$  and the in-plane prediction for  $w$  are zero, which was learned accurately by the network. As indicated by the plots of the absolute errors, the prediction



**Figure 7.8:** Prediction of *VoF-PINNs* for the pressure distribution along the center planes of the droplet in the in-plane and out-of-plane directions (left) in comparison to the ground truth simulation data (middle), and the absolute error between the prediction and the ground truth (right). The predicted contour of the droplet gas-liquid interface is indicated by the black solid line overlayed on the prediction and analogously the ground truth contour is overlayed on the ground truth velocity and pressure distributions. Figure adopted from [DKK<sup>+</sup>24].



of the vertical velocity component  $v$  is marginally more accurate in comparison to the horizontal velocity components, while the in-plane velocity  $u$  is predicted marginally more accurate than the out-of-plane velocity  $w$ . These results are consistent over the whole validation data set and are, furthermore, reflected in the relative  $L_1$  and  $L_2$  errors of the predicted quantities. The averaged relative  $L_1$  errors for all predicted quantities lie between  $L_{1,v} = 3.4\%$  and  $L_{1,w} = 4.6\%$ , while the relative  $L_2$  errors lie between  $L_{2,v} = 5.4\%$  and  $L_{2,w} = 8.7\%$ . A detailed overview of the relative and absolute errors for the prediction of the velocity and pressure distribution can be found in Appendix Tables C.6 and C.7 for *VoF-PINNs* and *PF-PINNs-v2*, respectively. The errors for the velocity and pressure prediction by *VoF-PINNs* were consistently lower in comparison to the prediction by *PF-PINNs-v2*, which indicates that the proposed VOF approach, in addition to a more accurate prediction of the interface, also allows a better prediction of the velocity and pressure. The relative errors of all predicted quantities varied across the validation dataset, which follows the dynamic deformation of the droplet and exhibits decaying values for velocity and pressure due to the damped oscillation of the interface. No substantial outliers or clear trend were observed for the development of the errors, in contrast to the observation of Qiu et al. [173], who found that low absolute velocities correlated with high relative errors, due to the focus of the optimization on training samples with large velocities. These results indicate that the proposed PINNs successfully learned to predict the velocity and pressure consistently during the dynamic deformation of the droplet. A potential reason for the more consistent results might be the employed training dataset, which consists mostly of samples with low velocities due to the aforementioned nature of droplet impingement, which in turn could balance out the influence of the relatively high losses from high-velocity samples. The evolution of the relative and absolute errors for *VoF-PINNs* on the subset of the validation data that features droplet impingement on a structured surface are plotted in Appendix Figures C.6 and C.7, respectively.

The comparison of the in-plane and the out-of-plane predictions reveals a higher in-plane accuracy for all predicted quantities. While comprehensive image features from the shadowgraphy contour are available for the in-plane prediction, the out-of-plane prediction depends on the trained model of the droplet dynamics and the 3D information encoded in the glare points. However, in comparison to the prediction of the gas-liquid interface, the available information for the velocity and pressure reconstruction is even more limited. Only the temporal evolution of the interface contour and the glare points provide some cues for the reconstruction, as the bulk of the shadowgraph does not carry any information for the velocity and pressure, as opposed to the phase distribution. Consequently, only the in-plane velocity at the contour and the glare points can be predicted with the support of image features, while the rest of the flow topology and pressure distribution have to be inferred from the trained model of droplet dynamics. The capability of the proposed PINNs to infer the three-dimensional distribution of these hidden quantities from the limited information in the input images demonstrates the capability of the physics-informed learning approach.

As observed in Figures 7.6 and 7.7 the accuracy of the velocity prediction inside of the droplet is higher in comparison to the outside. A likely cause for this difference is the significantly greater contribution of the physics-informed losses for the momentum equation in the liquid phase in comparison to the gaseous phase, due to the higher density and viscosity, which leads to higher residuals in the liquid phase and, thereby, focuses the optimization of the physics-informed losses on the inside of the droplet. Another source of prediction discrepancies in the outer region, farther from the interface, might arise from the outer boundary condition applied in the simulation (zero-gradient for velocity), which was intentionally omitted in the PINN implementation. This omission is deliberate, aiming to better capture arbitrary realistic distributions observed at the boundaries in experimental shadowgraphs. Furthermore, the errors in the vicinity of the gas-liquid interface are significantly lower compared to the rest of the domain, which is most obvious for the prediction

of the pressure distribution (see Figure 7.8). These results are likely related to the local refinement of the sampling points at the interface and the residual-based weighting. The higher density of the sampling points for both the data and physics-informed loss calculation focuses the optimization on the interface, while the residual-based weighting further reinforces the physics-informed losses at the interface. As can be seen most saliently in Figure 7.7(a) the prediction by the PINNs completes the velocity field in a physically reasonable manner, which indicates a certain capability for the extrapolation of prediction outside of the training data domain.

The successful prediction of the three-dimensional flow topology in combination with an improvement in the accuracy of the gas-liquid interface reconstruction on the synthetic data validates the proposed approach based on PINNs. The optimization of the phase distribution and velocity field is mutually dependent, as they are coupled by physics-informed loss derived from the residuals of the interface evolution equation. Thereby, the learning of an accurate velocity distribution at the interface promotes the learning of the interface location and vice versa. The accurate prediction for both the phase distribution and the velocity field in combination with low residuals for the interface evolution equation at the end of the training indicates that the underlying physics were successfully encoded in the neural network during the optimization of the PINNs. The developed PINNs were optimized towards a high accuracy of the gas-liquid interface, which supports an improvement of the velocity and pressure prediction by a different weighting of the loss terms during optimization.” from [DKK<sup>+</sup>24, Section 3.1]


### 7.3.2 Reconstruction of experimental data

“The PINNs that were trained and validated on the synthetic data [are employed] for the reconstruction of images recorded in the experiments by means of the glare-point shadowgraphy technique in order to evaluate their effectiveness in real-world applications. The same experimental data of droplet impingement on structured [hydrophilic polylactide (PLA) and hydrophobic Polydimethylsiloxane (PDMS) substrates that have been] [...] used for the evaluation of the baseline data-driven reconstruction technique [(see Section 6.1.2) is considered] to allow for a quantitative comparison of the reconstruction performance of both approaches. The following evaluation focuses on the VOF version of the developed PINNs, *VoF-PINNs*, which achieved superior performance among the three candidate models on the synthetic validation data.” from [DKK<sup>+</sup>24, Section 3.2]



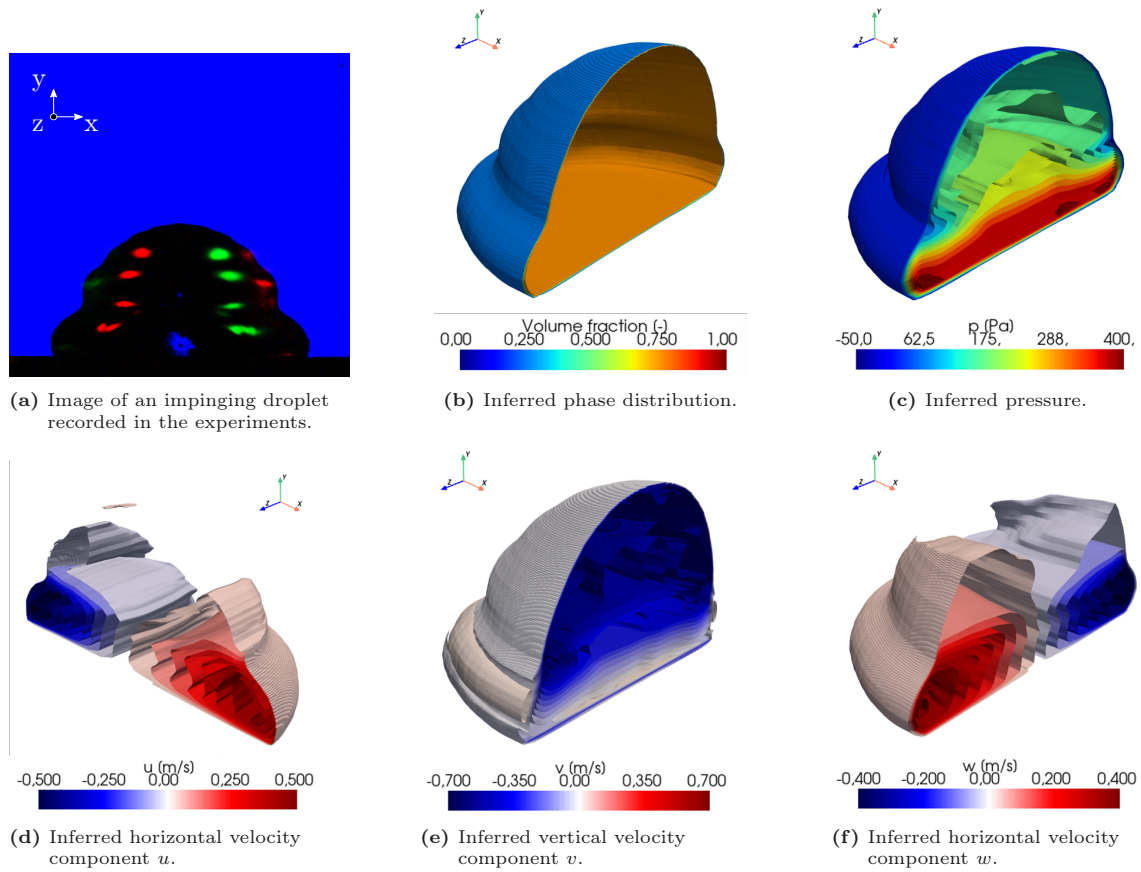
**Figure 7.9:** Out-of-plane reconstruction by the *DFS-GP* (left half) and *VoF-PINNs* (right half) for three different recordings from the experiments involving droplet impingement on the structured PDMS substrate at an observation angle of  $\varphi = 45^\circ$ . The droplet geometry reconstructed by *VoF-PINNs* is mirrored vertically to allow for a direct comparison of the reconstructed contours. Figure adopted from [DKK<sup>+</sup>24].

**Gas-liquid interface reconstruction** “Overall, [...] a substantial improvement of the reconstruction accuracy by the proposed PINNs [was found] in comparison to the data-driven baseline *DFS-GP* for the prediction of the gas-liquid interface on experimental data. The visual inspection of the reconstruction results for experiments with both substrates at different observation angles reveals that the gas-liquid interfaces reconstructed by the PINNs are smoother and more consistent over time in comparison to the reconstruction of *DFS-GP*. A comparison of the reconstruction results obtained by *DFS-GP* and *VoF-PINNs* is shown in Figure 7.9. The curvature of the interface reconstruction by *VoF-PINNs* appears to be more physical, as small-scale features with a high curvature that can be observed for the purely data-driven reconstruction are not present in the reconstruction by the PINNs. At the considered *We*-number, such features are unphysical, as the surface tension of the gas-liquid interface counteracts the formation of high curvatures. These results indicate that the consideration of surface tension in the momentum equation of the PINNs had a positive regularizing effect on the optimization of the neural network that led to a more physical reconstruction of the gas-liquid interface. Furthermore, the temporal evolution of the gas-liquid interface appears to be smoother for the reconstruction of the PINNs in comparison to the data-driven baseline, as the volume of the droplet remains more stable in time, unlike the fluctuation of the volume observed in the purely data-driven reconstruction.



**Figure 7.10:** Temporal evolution of the normalized integral volume of the reconstruction for droplet impingement on the structured PDMS substrate at  $\varphi = 90^\circ$  observation angle with *VoF-PINNs* and *DFS-GP*. Figure adopted from [DKK<sup>+</sup>24].

These results are illustrated in Figure 7.10, which shows the temporal evolution of the normalized integral droplet volume for the reconstruction of *VoF-PINNs* and *DFS-GP* for experiments that featured droplet impingement on the structured PDMS substrate at an observation angle of  $\varphi = 90^\circ$ . Throughout the entire period of time, the reconstruction of the PINNs lies significantly closer to the ground truth volume, indicated by the black line, and hence remains more conservative. Both reconstruction techniques exhibit larger errors immediately after the impingement of the droplet, which correlates with the strongest droplet deformation. However, the maximum error is significantly reduced by *VoF-PINNs* as compared to *DFS-GP*. Furthermore, the prediction of *DFS-GP* underestimates the volume of the droplet on average, while the average prediction of *VoF-PINNs* is close to the ground truth volume. Consequently, *VoF-PINNs* achieves a significantly lower uncertainty and bias error of the reconstructed volume of the droplet in comparison to *DFS-GP*. Specifically, *VoF-PINNs* reaches an average uncertainty of  $\sigma_V = 2.0\%$  and bias error of  $\delta_V = 2.2\%$  in comparison to  $\sigma_V = 6.2\%$  and bias error of  $\delta_V = 5.7\%$  by *DFS-GP*, effectively reducing the errors by a factor of three. It should be noted, that while a dedicated effort was made to ensure a fair comparison of the PINNs with the data-driven baseline model, both approaches require a different optimization, including tuning of hyperparameters and weighting of the loss terms. Consequently, the optimization of both networks and, in turn, their reconstruction performance could still be



**Figure 7.11:** Snapshot from the experiment involving droplet impingement on the structured PLA substrate (a) and iso-contours of the prediction by *VoF-PINNs*, phase distribution  $\alpha$  indicating the gas-liquid interface (b), pressure  $p$  (c), horizontal velocity component  $u$  (d), vertical velocity component  $v$  (e), and horizontal velocity component  $w$  (f). Figure adopted from [DKK<sup>+</sup>24].

further improved, thus changing the comparison of the performance. It should be noted, that a second version of the baseline data-driven model was trained with the same set of hyperparameters as the PINNs for a direct comparison, which yielded very similar results to *DFS-GP*. The substantial improvement of the reconstruction accuracy of the gas-liquid interface by the PINNs, however, suggests that the introduction of further prior knowledge through the physical constraints during training is an effective measure to enhance the trained model of the droplet dynamics. The lower errors by the PINNs in comparison to the data-driven baseline are consistently achieved for the reconstruction of all considered experiments, involving the impingement of droplets on structured PLA and PDMS substrates at different observation angles, as detailed in Appendix Table C.8. This includes varying degrees of gas-liquid interface deformation, as the impingement on the hydrophobic PDMS substrate leads to noticeably greater deformation of the droplet, demonstrating that the developed PINNs are adaptable towards different experimental conditions. Consequently, the developed PINNs, which were trained on two cases of numerical simulation, are able to generalize to different fluid mechanical conditions. Furthermore, the enhancement of the predictive accuracy for the gas-liquid interface that was observed for the reconstruction of the synthetic validation data carried over well to the prediction of experiments, which indicates that the approach of training on synthetic data is suitable for the proposed PINNs.” from [DKK<sup>+</sup>24, Section 3.2]

**Velocity and pressure reconstruction** “Figure 7.11 shows the inferred velocity and pressure distribution by *VoF-PINNs* for one snapshot of the experiments involving droplet impingement on

the structured PLA substrate. As can be seen, the PINNs are able to infer the continuous three-dimensional flow field inside of the droplet from a single input image. The comparison to snapshots at similar time instances and with comparable droplet geometries in the numerical simulation reveals that for some samples the velocity and pressure were inferred in a physically reasonable manner, while other frames yielded unphysical results. In particular, the velocity prediction failed at time instances that coincide with the inflection points of the droplet’s motion at which the acceleration of the gas-liquid interface approached zero and the velocity changed sign. Notably, the first of these inflection points, which occurred during the transition from the spreading phase to the receding phase (see [Section 2.2.2 and] Rioboo et al. [180, 181] for further details), was correctly predicted by the PINNs. However, the second inflection point at the end of the receding phase and the transition to the second spreading phase resulted in a high uncertainty of the velocity prediction, indicated by large fluctuations of the predicted velocity field. This behavior repeated at the following inflection points. Furthermore, the deformation of the droplet led to a self-occlusion of the gas-liquid interface in the input image, which resulted in increased errors in the velocity prediction. In contrast, the pressure prediction remained physically reasonable and similar to the validation data of the numerical simulation throughout the dynamic deformation of the droplet. In particular, the elevated dynamic pressure close to liquid-solid contact upon impact and the negative Laplace pressure at the negatively curved regions of the interface (*cp.* Figure 7.11c) were successfully reconstructed for different experimental conditions. While the inference of the velocity and pressure yielded accurate results for the synthetic validation data, these results could only partly be reproduced on the real data distribution, which involved different experimental conditions and, thereby, different droplet geometries in comparison to the training data. Consequently, the ability for the prediction of the velocity and pressure did not generalize as well as the prediction of the gas-liquid interface. This issue is likely related to the scarcity of information available for the velocity and pressure prediction. These results further substantiate the hypothesis that the predictive capability for the phase distribution generalizes well to novel droplet geometries due to the extracted image features, which provide additional information on the gas-liquid interface, but only indirectly for the velocity and pressure. The inference of the velocity and pressure thus relies more on the trained model of droplet dynamics. Consequently, the addition of further information in the images related to these quantities would likely improve the generalization capability toward different experimental conditions. Furthermore, the optimization of the PINNs was focused on the accurate prediction of the gas-liquid interface by a specific weighting of the loss terms. Therefore, improved results for the prediction of the velocity and pressure distributions can be expected for an increased relative weighting of the velocity and pressure data loss terms.” from [DKK<sup>+</sup>24, Section 3.2]

## 7.4 Concluding remarks regarding physics-informed learning of droplet dynamics

The presented results demonstrate that the introduction of additional prior knowledge to the reconstruction framework through physical constraints based on the governing equations of two-phase flows significantly enhances the accuracy of the gas-liquid interface reconstruction while extending the capabilities toward the prediction of three-dimensional field information for the velocity and pressure. The physics-informed losses, in particular the residuals of the interface evolution equation, couple the learning of the gas-liquid interface location and the velocity distribution at the interface. The accurate reconstruction of the interface location and the velocity field around the gas-liquid for the synthetic validation data in combination with low residuals of the interface evolution equation reveal that this relationship was learned well by the PINNs. Furthermore, the



consideration of surface tension in the momentum equation was found to result in a more physically correct prediction of the interface as indicated by a smoother surface with lower local curvature. This saliently demonstrates that the governing equations were successfully encoded in the neural network during optimization.

The improvements in the accuracy of the gas-liquid interface reconstruction that were observed for the synthetic validation data translated well to the reconstruction of real image data recorded in the experiments, leading to a significantly reduced uncertainty and enhanced temporal consistency of the reconstruction. Thereby, the developed PINNs exhibit significant generalization capabilities, demonstrated by the accurate reconstruction of experiments that considered different fluid mechanical conditions for the droplet impingement in comparison to the synthetic training data, which led to different states of droplet deformation. This generalization capability can be attributed to the prediction on the basis of image features extracted by the hourglass network that provide additional information on the spatial location of the gas-liquid interface to the PINNs. Moreover, these results confirm that the synthetic training data based on two cases of numerical simulation is suitable for training PINNs for a highly accurate gas-liquid interface reconstruction in the experiments.

“The developed PINNs offer prospects for the prediction of continuous three-dimensional velocity and pressure fields for both phases. The accurate quantitative prediction of velocity and pressure distributions on synthetic data, along with the promising reconstruction results for the flow topology on real measurement data, demonstrate that the inference of these hidden quantities from optical experiments using PINNs is feasible, though further development is needed to ensure reliable results. Since the proposed PINNs were optimized for the accurate prediction of the gas-liquid interface, adjusting the loss weighting could improve velocity and pressure learning. Additionally, incorporating more measurement data for the velocity or pressure could be beneficial for enhancing the accuracy of these predictions. The direct processing of images obtained by particle image velocimetry by the proposed convolutional-feature enhanced PINNs appears likewise feasible and promising.” from [DKK<sup>+</sup>24, Section 4]

“The simplicity of the VOF approach in combination with a superior accuracy of the predicted phase distribution and interface location in comparison to the phase-field approach renders it more suitable for the considered two-phase droplet flows. However, the explicit modeling of interface thickness in the phase-field approach enables the learning of thin interfaces, which may be valuable for certain applications. The control of interface thickness was identified as an important parameter for optimizing phase-field PINNs. Additionally, predicting the chemical potential separately in phase-field PINNs significantly improved neural network optimization and predictive accuracy, while reducing computational demands by avoiding the calculation of fourth-order derivatives in the Cahn-Hilliard equations, leading to faster training with fewer computational resources.” from [DKK<sup>+</sup>24, Section 4] The challenge of learning the sharp interfaces in two-phase flows by means of PINNs was successfully addressed by the local refinement of sampling points around the interface, as demonstrated by the particularly accurate reconstruction of the flow topology at the interface. “A further improvement of the reconstruction accuracy can potentially be achieved by feature expansion techniques, such as Fourier feature mapping [219], which has been shown to mitigate the spectral bias of neural networks towards lower frequencies and, thus, might facilitate the learning of a sharp gas-liquid interface. This approach has recently been applied successfully for the learning of 3D representations [219] and PINNs [243].” from [DKK<sup>+</sup>24, Section 4] Currently, the calculation of the derivatives in the physics-informed network does not consider the implicit dependence of the derivatives on the computation of the pixel-aligned features in the hourglass network. In order to obtain exact derivatives and facilitate a more accurate calculation of the PDE residuals, this implicit dependence should be considered in future work.

The proposed PINNs are a promising approach for the post-processing of two-phase flow experiments, as the accurate measurement of the interface position can be easily obtained by the shadowgraphy technique for various two-phase flows, whereas a direct measurement of field information for the velocity and pressure across the whole domain is often complicated or impractical. Therefore, the prediction of the three-dimensional flow field from simple optical experiments opens up new possibilities for the research of two-phase flows. The physically correct reconstruction of a smooth interface recommends the PINN approach for the application in the distortion correction of PIV measurements conducted within the curved interface of adhering droplets. Here, the accuracy of the distortion correction can be significantly enhanced, as the proposed PINNs enhance both the global shape of the reconstructed droplet and the local curvature of the interface, which are crucial for correctly mapping the PIV image – distorted by refraction at the interface – back to the object plane inside the droplet.





## 8 Conclusions and Outlook

This thesis investigates the feasibility of the spatio-temporal reconstruction of gas-liquid interfaces in droplet flows by means of deep learning techniques. Particularly, a data-driven framework is developed for the time-resolved three-dimensional reconstruction of the interface from monocular optical measurements. The investigation of fluid flows inside liquid droplets by particle image velocimetry (PIV) necessitates a correction of the distortion caused by the refraction of light at the curved gas-liquid interface to ensure accurate velocity measurements. The distortion correction for significantly deformed and temporally variable gas-liquid interfaces requires the instantaneous three-dimensional shape of the interface. Moreover, the constraints of limited optical access in the experiments demand the development of single-camera techniques. The review of existing two-phase flow reconstruction techniques based on optical measurements reveals that monocular reconstruction of three-dimensional gas-liquid interfaces is an ill-posed problem due to ambiguities in the projection of the 3D interface onto a 2D image and the scarcity of information in the images. This issue is conventionally addressed by imposing constraints derived from geometrical assumptions, such as symmetries or simple shapes (*e.g.*, ellipsoids or spherical caps), which limit the spatial accuracy of the reconstruction. Therefore, the accurate reconstruction of highly deformed interfaces currently requires multi-camera setups. However, recent advancements in deep learning techniques have led to data-driven approaches that leverage prior knowledge from training data to facilitate the monocular volumetric reconstruction of complicated geometries despite incomplete information in the images. This context gives rise to the central research question of this thesis: *Can data-driven techniques be used for three-dimensional and time-resolved reconstruction of gas-liquid interfaces in droplet flows on the basis of optical measurement techniques?*

To address this question, first, a data-driven reconstruction framework was conceptualized with the central scope of incorporating additional information on the gas-liquid interface and available domain knowledge to alleviate the ill-posedness of monocular spatio-temporal reconstruction by additional constraints. This framework combines a novel and purposefully developed optical measurement technique with a data-driven reconstruction technique based on neural networks. The reconstruction task was formulated as a supervised learning problem, in which a neural network is trained on a labeled dataset comprised of input glare-point shadowgraphy images and output three-dimensional gas-liquid interface shapes obtained by numerical simulation. This approach enables the employment of state-of-the-art supervised deep learning techniques for volumetric reconstruction and the incorporation of domain knowledge from the numerical simulation of two-phase droplet flows. Synthetic training data generation is employed to address the inherent matching problem of the experimental and numerical data, which stems from uncertainty in the experiments, as well as modeling and approximation errors in the simulation, and leads to a different shape of the droplets. An image generation pipeline that employs physically-based rendering to accurately reproduce the optical phenomena involved in image formation in the experiments was developed to provide suitable training data for the network. The synthetic image generation was validated by comparison to the experimental data and, lastly, the successful generalization of the neural networks trained on the synthetic data to the real-world application confirmed the effectiveness of the approach.

The issue of limited information on the three-dimensional shape of the gas-liquid interface in the experimental image data was addressed by the purposeful development of a novel optical

measurement technique. Among the plethora of measurement techniques employed for interface reconstruction in two-phase flows, the shadowgraphy technique is favored due to its high spatial accuracy and simplicity, which facilitate its application in a variety of experimental setups. To improve upon this technique, a novel optical method was introduced that extends the canonical shadowgraphy by glare points from lateral light sources to encode additional 3D information on the gas-liquid interface in the images. A comparative experimental study with the standard shadowgraphy demonstrated that the additional glare points significantly improve the accuracy of the interface reconstruction, in particular, the depth estimation. This demonstrates that the glare points are an effective measure for encoding additional information on the three-dimensional shape of the gas-liquid interface in the images, which is successfully exploited by the neural network. Feature visualization techniques elucidated that the convolutional neural network progressively refines the interface prediction by the information contained in the glare points during feature extraction. Furthermore, the high spatial accuracy and contrast of the gas-liquid interface in the images offer a strong basis for the in-plane reconstruction, which – combined with the specialized convolutional neural network architecture that extracts local high-resolution features – allowed for a high accuracy and generalizability of the in-plane reconstruction. These results demonstrate that the harmonization of optical measurement techniques with deep learning image processing methods – through the integration of complementary approaches and coordinated joint development – leads to significant synergies.

The introduction of available domain knowledge from numerical simulations and analytical approaches for constraining the ill-posed problem of monocular spatio-temporal reconstruction was developed in two stages. In the first stage, the results of numerical simulations were incorporated into the reconstruction framework through supervised learning of the neural network on numerical data of the three-dimensional gas-liquid interface. The reconstruction results show that the data-driven approach leverages the learned knowledge of the droplet dynamics to fill in missing information in the input images in a physically reasonable way, where experimental data is unavailable. This was most saliently demonstrated by the successful reconstruction of self-obscuring gas-liquid interface shapes during droplet impingement. In summary, the prior knowledge acquired by the neural network from the numerical training data implicitly guides the training of a physically accurate model of the gas-liquid interface dynamics. In the second stage, the governing equations of two-phase flows were directly integrated during the optimization of the neural network by means of physics-informed neural networks (PINNs). The existing neural network architecture was extended for the prediction of the three-dimensional velocity and pressure distribution. Moreover, additional physics-informed loss terms were introduced, derived from the residuals of the single-field Navier-Stokes equation, the continuity equation, and the evolution equation of the phase distribution, which are encoded in the neural network during the joint optimization of both physics-informed and data losses. This allows the developed PINNs to seamlessly integrate both sources of prior knowledge; the data from numerical simulation and the analytical description of the physics by the governing equations.

The results of the PINN-based approach revealed a significant improvement in the accuracy of the gas-liquid interface reconstruction. Furthermore, the temporal consistency of the reconstruction was greatly enhanced, leading to a better conservation of the droplet's volume in time. The incorporation of the surface tension through the encoded momentum equations contributed to a smoother and more physically accurate interface reconstruction, resulting in a visibly improved quality of the reconstructed gas-liquid interface. Furthermore, the developed PINNs offer the added capability to predict the continuous three-dimensional velocity and pressure fields for both phases. While the model successfully predicted these latent fields for the synthetic validation data, the application to experimental data requires further development to ensure a reliable prediction.

For the scope of this thesis, however, the inclusion of velocity data at the interface, combined with the encoded interface evolution equation, encourages the neural network to learn a more precise localization of the gas-liquid interface, which emphasizes the positive effect of introducing additional prior knowledge to the network optimization. Overall, the improved accuracy of the interface reconstruction demonstrates the effectiveness of introducing inductive bias to the data-driven reconstruction through the physics-informed losses – along with domain knowledge from the numerical data – to provide the necessary regularization of the ill-posed problem of monocular volumetric reconstruction.

The capabilities and limitations of the developed reconstruction framework were assessed through an extensive comparative study, which involved experiments with progressively more complex droplet phenomena. The results demonstrated that a high accuracy of the gas-liquid interface reconstruction can be achieved for various droplet flows, including the droplet impingement on both flat and structured substrates and adhering droplets subjected to external shear flows. These findings elucidate the ability of the reconstruction framework to train versatile models of the droplet-gas-liquid interface dynamics. The reliable reconstruction of adhering droplets at significantly higher external flow velocities in comparison to the training data demonstrates a substantial capability for the extrapolation beyond the fluid mechanical conditions covered during training. Despite the more severe deformation and unknown shapes of the gas-liquid interface under these conditions, the model largely maintained its accuracy. These findings are further substantiated by the successful reconstruction of the interface dynamics during droplet impingement on substrates with unknown wettability. In particular, the in-plane reconstruction generalizes well to unfamiliar interface shapes, because the prediction of the neural network is based on spatially aligned local features extracted from the images by the convolutional neural network. Thereby, the shadowgraphy contour served as a robust and accurate basis for the in-plane reconstruction. Conversely, the out-of-plane reconstruction relies more heavily on the trained droplet model of the gas-liquid interface dynamics due to the sparse and localized information provided by the glare points. The numerical training data induces a bias of the trained model, particularly obvious in the out-of-plane reconstruction, which ensures robust depth estimation for unknown interface geometries. However, this bias leads to the introduction of increasing errors the further the experimental data deviate from the training data distribution. To optimize reconstruction quality, the neural network should be trained on a dataset that closely matches the expected experimental conditions. The preferential direction of the reconstruction should be taken into account during the design of the experimental setup by aligning the direction of the expected maximum deformation of the interface with the image plane to exploit the high accuracy of the in-plane reconstruction and assist the out-of-plane reconstruction.

The effectiveness of this approach was demonstrated by the reconstruction of adhering droplets, for which the experiment and reconstruction technique were set up complementary to leverage the symmetry of the problem and, thereby, simplify the out-of-plane reconstruction. The chosen deep learning technique based on convolutional neural networks for image processing and the implicit representation of the gas-liquid interface as a level-set function learned by a multi-layer perceptron proves well-suited for the accurate reconstruction of the dynamic interfaces in droplet flows. The implicit representation enables the reconstruction of topological changes, such as the formation of secondary droplets during impingement. Moreover, the reconstruction can be performed at arbitrary resolutions, including higher resolutions compared to the training data, due to the learned continuous approximation of the droplet dynamics. Remarkably, even a modest amount of training data – for the case of adhering droplets, just a single numerical simulation – is sufficient to train the network. Thereby, smooth meshes of the instantaneous three-dimensional gas-liquid interface

can be obtained, which are suitable for a distortion correction of PIV measurements conducted through the curved interfaces in droplet flows.

In conclusion, this thesis demonstrates the potential of deep learning techniques for the accurate and reliable reconstruction of the spatio-temporal interface dynamics in droplet flows on the basis of optical measurement techniques. In future work, the suitability of the reconstructed gas-liquid interfaces for distortion correction in PIV measurements should be investigated experimentally. The ray tracing approach proposed by Kang et al. [107] and Minor et al. [150] appears promising for the reconstruction of the original particle distribution by tracing the approximately reciprocal light path from the distorted images, through the reconstructed gas-liquid interface, to the position of the light sheet. Preliminary testing within the developed rendering setup has confirmed the feasibility of this correction approach. To assess its accuracy, corrected PIV measurements should be compared with validation data, data from numerical simulations, or a different velocity measurement technique.

While the reconstruction framework was presented alongside the results for adhering droplets subjected to external shear flow and droplet impingement on flat and structured surfaces, the reconstruction of other two-phase flow phenomena such as impacts at an angle, splashing droplets, sprays, and gaseous bubbles in liquids, appears likewise feasible. The general setup of the synthetic data generation based on the numerical simulation of the flow, followed by physically-based rendering can be adapted to provide training data for various two-phase flows, while the neural network can be trained with as little as one simulation case, which facilitates a straightforward application. Improvements in reconstruction accuracy may be achieved by enhancing the visual quality of synthetic data to more closely resemble real-world data. In addition to the optimization of the rendering setup, domain adaptation techniques based on deep learning [40] appear promising.

Although the focus of this thesis was on the reconstruction of the gas-liquid interface, initial promising results for the reconstruction of the velocity and pressure distributions through PINNs demonstrate the feasibility of such an approach and warrant further investigation. Particularly, the reliable prediction of these latent quantities has to be ensured through further optimization of the training dynamics and network architecture. The incorporation of additional measurement data for the velocity or pressure into the reconstruction framework may be beneficial. For the further development of the proposed PINNs, integrating well-established methods from numerical simulations of two-phase flows could prove valuable, as both fields share common principles, which facilitates a transfer of techniques and knowledge. The developed reconstruction framework saliently demonstrates the potential for deep learning techniques for the integration of diverse domain knowledge in experiments, numerical simulation, and analytical models. This approach effectively regularizes previously ill-posed problems, such as monocular volumetric reconstruction, by incorporating prior knowledge to provide the necessary inductive biases for a data-driven solution to the problem. Thereby, this thesis is intended to lay the foundation for future research that leverages data-driven techniques to integrate diverse domain knowledge to advance toward more unified models of fluid flows.

## 9 References

### Bibliography

- [1] S. J. Anagnostopoulos, J. D. Toscano, N. Stergiopoulos, and G. E. Karniadakis. Residual-based attention in physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 421:116805, 2024.
- [2] M. Andersson, S. Beale, M. Espinoza, Z. Wu, and W. Lehnert. A review of cell-scale multi-phase flow modeling, including water management, in polymer electrolyte fuel cells. *Applied Energy*, 180:757–778, Oct. 2016.
- [3] M. Andersson, V. Vukčević, S. Zhang, Y. Qi, H. Jasak, S. Beale, and W. Lehnert. Modeling of droplet detachment using dynamic contact angles in polymer electrolyte fuel cell gas channels. *International Journal of Hydrogen Energy*, 44(21):11088–11096, Apr. 2019.
- [4] R. Andrade, O. Skurtys, and F. Osorio. Drop impact behavior on food using spray coating: Fundamentals and applications. *Food Research International*, 54(1):397–405, 2013.
- [5] N. Ashgriz. *Handbook of Atomization and Sprays: Theory and Applications*. Springer, Heidelberg, 2011.
- [6] A. R. Barron. Approximation and estimation bounds for artificial neural networks. *Machine Learning*, 14(1):115–133, Jan. 1994.
- [7] B. Barwari, S. Burgmann, and U. Janoske. Deformation and movement of adhering droplets in shear flow. In *5th International Conference on Experimental Fluid Mechanics, Munich, July 2-4*, 2018.
- [8] B. Barwari, S. Burgmann, A. Bechtold, M. Rohde, and U. Janoske. Experimental study of the onset of downstream motion of adhering droplets in turbulent shear flows. *Experimental Thermal and Fluid Science*, 109:109843, 2019.
- [9] G. K. Batchelor. *An Introduction to Fluid Dynamics*. Cambridge University Press, 1 edition, Feb. 2000.
- [10] A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind. Automatic differentiation in machine learning: a survey. *J. Mach. Learn. Res.*, 18(1):5595–5637, Jan. 2017.
- [11] J. S. Bendat and A. G. Piersol. *Random data : analysis and measurement procedures*. Wiley series in probability and statistics. Wiley, 4th ed (online-ausg.) edition, 2010.
- [12] J. Berg. *Wettability: John C. Berg (ed.)*. Surfactant Science Series. Marcel Dekker, 1 edition, Apr. 1993.
- [13] C. Bilsing, U. Janoske, J. Czarske, L. Büttner, and S. Burgmann. 3D-3C measurements of flow reversal in small sessile drops in shear flow. *International Journal of Multiphase Flow*, 182:105017, Jan. 2025.

- [14] J.-D. Boissonnat and S. Oudot. Provably good sampling and meshing of surfaces. *Graphical Models*, 67(5):405–451, Sept. 2005.
- [15] M. Born, E. Wolf, A. B. Bhatia, P. C. Clemmow, D. Gabor, A. R. Stokes, A. M. Taylor, P. A. Wayman, and W. L. Wilcock. *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*. Cambridge University Press, 7 edition, 1999.
- [16] J. Brackbill, D. Kothe, and C. Zemach. A continuum method for modeling surface tension. *Journal of Computational Physics*, 100(2):335–354, June 1992.
- [17] M. Brunel, B. Delestre, M. Talbi, and M. Fromager. Interferometric imaging for the tomography of rough particles in a flow: A case study. *Optics Communications*, 479:126412, 2021.
- [18] S. L. Brunton, B. R. Noack, and P. Koumoutsakos. Machine Learning for Fluid Mechanics. *Annual Review of Fluid Mechanics*, 52(1):477–508, Jan. 2020.
- [19] D. Bucciarelli, S. Wendsche, M. Klemm, P. Sandbacka, C. Ehoumanand, A. Castagnini, G. Quade, and O. Ghotbi. *LuxCoreRender Project*. LuxCoreRender Project, 2018.
- [20] E. Buckingham. On Physically Similar Systems; Illustrations of the Use of Dimensional Equations. *Physical Review*, 4(4):345–376, Oct. 1914.
- [21] A. B. Buhendwa, S. Adami, and N. A. Adams. Inferring incompressible two-phase flow fields from the interface motion using physics-informed neural networks. *Machine Learning with Applications*, 4:100029, 2021.
- [22] S. Burgmann, M. Blank, U. Panchenko, and J. Wartmann.  $\mu$ PIV measurements of two-phase flows of an operated direct methanol fuel cell. *Experiments in Fluids*, 54, 2013.
- [23] S. Burgmann, B. Barwari, and U. Janoske. Oscillation of adhering droplets in shear flow. In *5th International Conference on Experimental Fluid Mechanics, Munich, July 2-4*, 07 2018.
- [24] S. Burgmann, B. Barwari, and U. Janoske. Inner flow structure of an adhering oscillating droplet in shear flow. In *19th Int. Symp. on the Application of Laser and Imaging Techniques to Fluid Mechanics, Lissabon, July 16-19*, 07 2018.
- [25] S. Burgmann, M. Dues, B. Barwari, J. Steinbock, L. Büttner, J. Czarske, and U. Janoske. Flow measurements in the wake of an adhering and oscillating droplet using laser-doppler velocity profile sensor. *Experiments in Fluids*, 62(3):47, 2021.
- [26] S. Burgmann, V. Krämer, M. Rohde, M. Dues, and U. Janoske. Inner and outer flow of an adhering droplet in shear flow. *International Journal of Multiphase Flow*, 153:104140, Aug. 2022.
- [27] J. W. Cahn and J. E. Hilliard. Free Energy of a Nonuniform System. III. Nucleation in a Two-Component Incompressible Fluid. *The Journal of Chemical Physics*, 31(3):688–699, Sept. 1959.
- [28] S. Cai, Z. Mao, Z. Wang, M. Yin, and G. E. Karniadakis. Physics-informed neural networks (PINNs) for fluid mechanics: a review. *Acta Mechanica Sinica*, 37(12):1727–1738, 2021.
- [29] S. Cai, Z. Wang, F. Fuest, Y. J. Jeon, C. Gray, and G. E. Karniadakis. Flow over an espresso cup: inferring 3-D velocity and pressure fields from tomographic background oriented Schlieren via physics-informed neural networks. *Journal of Fluid Mechanics*, 915:A102, May 2021.



- 
- [30] K. Calvin, D. Dasgupta, G. Krinner, A. Mukherji, P. W. Thorne, C. Trisos, J. Romero, and P. Aldunce. Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change [Core Writing Team, H. Lee and J. Romero (eds.)]. Technical report, Intergovernmental Panel on Climate Change (IPCC), Geneva, Switzerland., July 2023.
  - [31] A. B. D. Cassie and S. Baxter. Wettability of porous surfaces. *Transactions of the Faraday Society*, 40(0):546–551, 1944.
  - [32] H. Chen, D. Y. Goswami, and E. K. Stefanakos. A review of thermodynamic cycles and working fluids for the conversion of low-grade heat. *Renewable and Sustainable Energy Reviews*, 14(9):3059–3067, Dec. 2010.
  - [33] N. Chen, R. Ma, A. Chen, and C. Cui. *PC-PINNs: Physics-informed neural networks for solving the phase-field model of pitting corrosion*, pages 244–252. CRC Press, 07 2024.
  - [34] Z. Chen and H. Zhang. Learning implicit fields for generative shape modeling. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5932–5941. IEEE, 2019.
  - [35] C. Cheng and G.-T. Zhang. Deep Learning Method Based on Physics Informed Neural Network with Resnet Block for Solving Fluid Flow Problems. *Water*, 13(4):423, Feb. 2021.
  - [36] W.-L. Cheng, W.-W. Zhang, H. Chen, and L. Hu. Spray cooling and flash evaporation cooling: The current development and application. *Renewable and Sustainable Energy Reviews*, 55: 614–628, Mar. 2016.
  - [37] C. B. Choy, D. Xu, J. Gwak, K. Chen, and S. Savarese. 3D-R2N2: A unified approach for single and multi-view 3D object reconstruction. In B. Leibe, J. Matas, N. Sebe, and M. Welling, editors, *Computer Vision – ECCV 2016*, pages 628–644, Cham, 2016. Springer International Publishing.
  - [38] J. Collier. *Convective Boiling and Condensation*. The Advanced Book Program. McGraw-Hill, London, 1981.
  - [39] B. O. Community. *Blender - a 3D modelling and rendering package*. Blender Foundation, Stichting Blender Foundation, 2018.
  - [40] G. Csurka. *Domain Adaptation in Computer Vision Applications*. Springer International Publishing, Cham, 2017.
  - [41] S. Cuomo, V. S. Di Cola, F. Giampaolo, G. Rozza, M. Raissi, and F. Piccialli. Scientific machine learning through physics-informed neural networks: Where we are and what’s next. *Journal of Scientific Computing*, 92(3), 2022.
  - [42] G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals, and Systems*, 2(4):303–314, Dec. 1989.
  - [43] A. Dalili, S. Chandra, J. Mostaghimi, H. C. Fan, and J. C. Simmer. Bubble entrapment and escape from sprayed paint films. *Progress in Organic Coatings*, 97:153–165, 2016.
  - [44] P.-G. De Gennes, F. Brochard-Wyart, and D. Quéré. *Capillarity and Wetting Phenomena*. Springer New York, New York, NY, 2004.
  - [45] P. Debye. Das elektromagnetische Feld um einen Zylinder und die Theorie des Regenbogens. *Phys. Z.*, 9(1464):775–778, 1908.

- [46] S. Dehaeck, Y. Tsoumpas, and P. Colinet. Extended glare point velocimetry and sizing for bubbly flows. *Experiments in Fluids*, 39(2):407–419, 2005.
- [47] S. Dehaeck, Y. Tsoumpas, and P. Colinet. Analyzing droplets through digital holography and a 1D wavelet transform technique. In *Digital Holography and Three-Dimensional Imaging*, page DW3A.5, Washington, D.C., 2013. OSA.
- [48] S. Dehaeck, Y. Tsoumpas, and P. Colinet. Analyzing closed-fringe images using two-dimensional fan wavelets. *Applied optics*, 54(10):2939–2952, 2015.
- [49] M. Delmas, M. Monthieux, and T. Ondarçuhu. Contact angle hysteresis at the nanometer scale. *Phys. Rev. Lett.*, 106:136102, 2011.
- [50] F. J. Dent. *Fluid-based micropatterning for the fabrication of biomimetic functional surfaces*. PhD thesis, University of Leeds, April 2024. Accessed 04.11.2024.
- [51] M. Dreisbach. Particle Detection by means of Neural Networks in Defocusing Particle Tracking Velocimetry on Synthetic and Experimental Data. Master’s thesis, Karlsruhe Institute of Technology, Karlsruhe, Germany, 2020.
- [52] V. Dumoulin and F. Visin. A guide to convolution arithmetic for deep learning. *arXiv:1603.07285*, 2018.
- [53] E. B. Dussan V. On the ability of drops to stick to surfaces of solids. part 3. the influences of the motion of the surrounding fluid on dislodging drops. *Journal of Fluid Mechanics*, 174: 381–397, 1987.
- [54] C. A. Dutra Fraga Filho. *Smoothed Particle Hydrodynamics: Fundamentals and Basic Applications in Continuum Mechanics*. Springer International Publishing, Cham, 2019.
- [55] M. Duxenneuner, O. Pust, P. Fischer, E. Windhab, and J. Cooper-White. Flow visualization of a forming droplet in a micro channel with FlexPIV. In *8th International Symposium on Particle Image Velocimetry, PIV’09 : August 25 - 28, 2009, Melbourne, Australia*, 08 2009.
- [56] M. R. Duxenneuner, P. Fischer, E. J. Windhab, and J. J. Cooper-White. Simultaneous visualization of the flow inside and around droplets generated in microchannels. *Microfluidics and Nanofluidics*, 16(4):743–755, 2014.
- [57] N. Dyn and D. Levin. Subdivision schemes in geometric modelling. *Acta Numerica*, 11: 73–144, Jan. 2002.
- [58] D. R. Easterling, G. A. Meehl, C. Parmesan, S. A. Changnon, T. R. Karl, and L. O. Mearns. Climate Extremes: Observations, Modeling, and Impacts. *Science*, 289(5487):2068–2074, Sept. 2000.
- [59] H. Eivazi, M. Tahani, P. Schlatter, and R. Vinuesa. Physics-informed neural networks for solving reynolds-averaged Navier–Stokes equations. *Physics of Fluids*, 2022.
- [60] A. Esposito, A. D. Montello, Y. G. Guezennec, and C. Pianese. Experimental investigation of water droplet–air flow interaction in a non-reacting pem fuel cell channel. *Journal of Power Sources*, 195(9):2691–2699, 2010.
- [61] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (VOC) challenge. *International Journal of Computer Vision*, 88(2): 303–338, Jun 2010.

- 
- [62] H. Fan, H. Su, and L. Guibas. A point set generation network for 3D object reconstruction from a single image. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2463–2471. IEEE, 2017.
  - [63] J. Fan, M. Wilson, and N. Kapur. Displacement of liquid droplets on a surface by a shearing air flow. *Journal of Colloid and Interface Science*, 356(1):286–292, 2011.
  - [64] V. Fink, X. Cai, A. Stroh, R. Bernard, J. Kriegseis, B. Frohnapfel, H. Marschall, and M. Wörner. Drop bouncing by micro-grooves. *International Journal of Heat and Fluid Flow*, 70:271–278, 2018.
  - [65] S. C. Fu, W. T. Leung, and C. Y. H. Chao. Detachment of droplets in a fully developed turbulent channel flow. *Aerosol Science and Technology*, 48(9):916–923, 2014.
  - [66] Y. Fu and Y. Liu. 3D bubble reconstruction using multiple cameras and space carving method. *Measurement Science and Technology*, 29(7):075206, 2018.
  - [67] A. Fujiwara, Y. Danmoto, K. Hishida, and M. Maeda. Bubble deformation and flow structure measured by double shadow images and PIV/LIF. *Experiments in Fluids*, 36(1):157–165, 2004.
  - [68] V. H. Gada and A. Sharma. On Derivation and Physical Interpretation of Level Set Method–Based Equations for Two-Phase Flow Simulations. *Numerical Heat Transfer, Part B: Fundamentals*, 56(4):307–322, Nov. 2009.
  - [69] P. Ge, S. Wang, J. Zhang, and B. Yang. Micro-/nanostructures meet anisotropic wetting: from preparation methods to applications. *Materials Horizons*, 7(10):2566–2595, 2020.
  - [70] Gemini Team. Gemini: A Family of Highly Capable Multimodal Models, 2023. arXiv:2312.11805.
  - [71] R. Girdhar, D. F. Fouhey, M. Rodriguez, and A. Gupta. Learning a predictable and generative vector representation for objects. In B. Leibe, J. Matas, N. Sebe, and M. Welling, editors, *Computer Vision – ECCV 2016*, pages 484–499, Cham, 2016. Springer International Publishing.
  - [72] A. R. Glover, S. M. Skippon, and R. D. Boyle. Interferometric laser imaging for droplet sizing: a method for droplet-size measurement in sparse spray systems. *Applied optics*, 34(36):8409–8421, 1995.
  - [73] C. Gong, Y. Song, G. Huang, W. Chen, J. Yin, and D. Wang. BubDepth: A neural network approach to three-dimensional reconstruction of bubble geometry from single-view images. *International Journal of Multiphase Flow*, 152:104100, 2022.
  - [74] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
  - [75] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.
  - [76] A. K. Gupta and S. Basu. Deformation of an oil droplet on a solid substrate in simple shear flow. *Chemical Engineering Science*, 63(22):5496–5502, 2008.
  - [77] G. M. Hale and M. R. Querry. Optical constants of water in the 200-nm to 200-microm wavelength region. *Applied optics*, 12(3):555–563, 1973.

- [78] A. Hannun, C. Case, J. Casper, B. Catanzaro, G. Diamos, E. Elsen, R. Prenger, S. Satheesh, S. Sengupta, A. Coates, and A. Y. Ng. Deep speech: Scaling up end-to-end speech recognition. *arXiv:1412.5567*, 2014.
- [79] F. H. Harlow and J. E. Welch. Numerical Calculation of Time-Dependent Viscous Incompressible Flow of Fluid with Free Surface. *The Physics of Fluids*, 8(12):2182–2189, Dec. 1965.
- [80] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *arXiv:1512.03385*, 2015.
- [81] T. He, J. Collomosse, H. Jin, and S. Soatto. Geo-pifu: Geometry and pixel aligned implicit functions for single-view human reconstruction. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- [82] G. Healey and T. O. Binford. Local shape from specularity. *Computer Vision, Graphics, and Image Processing*, 42(1):62–86, 1988.
- [83] E. Hecht. *Optics*. Pearson, Harlow, 2017.
- [84] A. A. Heydari, C. A. Thompson, and A. Mehmood. SoftAdapt: Techniques for Adaptive Loss Weighting of Neural Networks with Multi-Part Loss Functions. *arXiv:1912.12355*, Dec. 2019.
- [85] M. Higashine, K. Katoh, T. Wakimoto, and T. Azuma. Profiles of Liquid Droplets on Solid Plates in Gravitational and Centrifugal Fields. *Journal of the Japanese Society for Experimental Mechanics*, 8:49–54, 2008.
- [86] C. Hirt and B. Nichols. Volume of fluid (VOF) method for the dynamics of free boundaries. *Journal of Computational Physics*, 39(1):201–225, Jan. 1981.
- [87] C. Hirt, A. Amsden, and J. Cook. An arbitrary Lagrangian-Eulerian computing method for all flow speeds. *Journal of Computational Physics*, 14(3):227–253, Mar. 1974.
- [88] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [89] R. L. Hoffman. A study of the advancing interface. I. Interface shape in liquid—gas systems. *Journal of Colloid and Interface Science*, 50(2):228–241, Feb. 1975.
- [90] M. Honkanen. Reconstruction of a three-dimensional bubble surface from high-speed orthogonal imaging of dilute bubbly flow. In A. A. Mammoli and C. A. Brebbia, editors, *Computational Methods in Multiphase Flow V*, WIT Transactions on Engineering Sciences, pages 469–480. WIT Press Southampton, UK, 2009.
- [91] J. W. Horbach and T. Dang. 3D reconstruction of specular surfaces using a calibrated projector-camera setup. *Machine Vision and Applications*, 21(3):331–340, 2010.
- [92] B. Horn. Shape from shading: A method for obtaining the shape of a smooth opaque object from one view. Technical report, Massachusetts Institute of Technology, 1970.
- [93] K. Hornik, M. Stinchcombe, and H. White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, Jan. 1989.
- [94] E. A. Hovenac. Calculation of far-field scattering from nonspherical particles using a geometrical optics approach. *Applied optics*, 30(33):4739–4746, 1991.

- 
- [95] H. Hu, B. Wang, K. Zhang, W. Lohry, and S. Zhang. Quantification of transient behavior of wind-driven surface droplet/rivulet flows using a digital fringe projection technique. *Journal of Visualization*, 18(4):705–718, 2015.
  - [96] P. J. Huber. Robust Estimation of a Location Parameter. *The Annals of Mathematical Statistics*, 35(1):73–101, Mar. 1964.
  - [97] I. Ihrke, B. Goidluecke, and M. Magnor. Reconstructing the geometry of flowing water. In *Tenth IEEE International Conference on Computer Vision (ICCV’05) Volume 1*, pages 1055–1060 Vol. 2. IEEE, 2005.
  - [98] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv:1502.03167*, 2015.
  - [99] J. Donahue, L. A. Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, T. Darrell, and K. Saenko. Long-term recurrent convolutional networks for visual recognition and description. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2625–2634, 2015.
  - [100] D. Jacqmin. Calculation of Two-Phase Navier–Stokes Flows Using Phase-Field Modeling. *Journal of Computational Physics*, 155(1):96–127, Oct. 1999.
  - [101] A. D. Jagtap, K. Kawaguchi, and G. Em Karniadakis. Locally adaptive activation functions with slope recovery for deep and physics-informed neural networks. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 476(2239):20200334, July 2020.
  - [102] H. A. Jakobsen. *Chemical Reactor Modeling: Multiphase Reactive Flows*. Springer International Publishing, Cham, 2014.
  - [103] X. Jin, S. Cai, H. Li, and G. E. Karniadakis. NSFnets (Navier-Stokes flow nets): Physics-informed neural networks for the incompressible Navier-Stokes equations. *Journal of Computational Physics*, 426:109951, 2021.
  - [104] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Žídek, A. Potapenko, A. Bridgland, C. Meyer, S. A. A. Kohl, A. J. Ballard, A. Cowie, B. Romera-Paredes, S. Nikolov, R. Jain, J. Adler, T. Back, S. Petersen, D. Reiman, E. Clancy, M. Zielinski, M. Steinegger, M. Pacholska, T. Berghammer, S. Bodenstein, D. Silver, O. Vinyals, A. W. Senior, K. Kavukcuoglu, P. Kohli, and D. Hassabis. Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873):583–589, Aug. 2021.
  - [105] Y. C. Jung and B. Bhushan. Dynamic Effects of Bouncing Water Droplets on Superhydrophobic Surfaces. *Langmuir*, 24(12):6262–6269, June 2008.
  - [106] S. G. Kandlikar and M. E. Steinke. Contact angles and interface behavior during rapid evaporation of liquid on a heated surface. *International Journal of Heat and Mass Transfer*, 45(18):3771–3780, 2002.
  - [107] K. H. Kang, S. J. Lee, C. M. Lee, and I. S. Kang. Quantitative visualization of flow inside an evaporating droplet using the ray tracing method. *Measurement Science and Technology*, 15(6):1104, 2004.
  - [108] N. Kantarci, F. Borak, and K. O. Ulgen. Bubble column reactors. *Process Biochemistry*, 40(7):2263–2283, June 2005.

- [109] L. Karlsson, H. Lycksam, A.-L. Ljung, P. Gren, and T. S. Lundström. Experimental study of the internal flow in freezing water droplets on a cold surface. *Experiments in Fluids*, 60(12):182, 2019.
- [110] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, May 2021.
- [111] M. Kerker. *The Scattering of Light and Other Electromagnetic Radiation*. Elsevier, 1969.
- [112] D. Khojasteh, M. Kazerooni, S. Salarian, and R. Kamali. Droplet impact on superhydrophobic surfaces: A review of recent developments. *Journal of Industrial and Engineering Chemistry*, 42:1–14, Oct. 2016.
- [113] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, Conference Track Proceedings*, 2015.
- [114] D. P. Kingma and M. Welling. Auto-encoding variational bayes, 2013. arXiv:1312.6114.
- [115] J. L. Kinsey. Laser-induced fluorescence. *Annual Review of Physical Chemistry*, 28(1):349–372, 1977.
- [116] E. Kiyani, M. Kooshkbaghi, K. Shukla, R. B. Koneru, Z. Li, L. Bravo, A. Ghoshal, G. E. Karniadakis, and M. Karttunen. Characterization of partial wetting by cmas droplets using multiphase many-body dissipative particle dynamics and data-driven discovery based on PINNs. *Journal of Fluid Mechanics*, page A7, 2024.
- [117] G. Klinker, S. A. Shafer, and T. Kanade. The measurement of highlights in color images. *International Journal of Computer Vision*, 2:7–32, 1988.
- [118] L. P. Kobbelt, M. Botsch, U. Schwanke, and H.-P. Seidel. Feature sensitive surface extraction from volume data. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 57–66. ACM, Aug. 2001.
- [119] G. König, K. Anders, and A. Frohn. A new light-scattering technique to measure the diameter of periodically generated moving droplets. *Journal of Aerosol Science*, 17(2):157–167, 1986.
- [120] C. Korte and A. M. Jacobi. Condensate retention effects on the performance of plain-fin-and-tube heat exchangers: Retention data and modeling. *Journal of Heat Transfer*, 123(5):926–936, 2001.
- [121] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 1097–1105. Curran Associates, Inc., 2012.
- [122] V. Krämer, B. Barwari, S. Burgmann, M. Rohde, S. Rentschler, C. Holzknecht, C. Gmelin, and U. Janoske. Numerical analysis of an adhering droplet applying an adapted feedback deceleration technique. *International Journal of Multiphase Flow*, 145:103808, Dec. 2021.
- [123] E. C. Kumbur, K. V. Sharp, and M. M. Mench. Liquid droplet behavior and instability in a polymer electrolyte fuel cell flow channel. *Journal of Power Sources*, 161(1):333–345, 2006.
- [124] A. Laurentini. The visual hull concept for silhouette-based image understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(2):150–162, 1994.



- 
- [125] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. In *Proceedings of the IEEE*, pages 2278–2324, 1998.
  - [126] Y. LeCun, L. Bottou, G. B. Orr, and K. R. Müller. *Efficient BackProp*, pages 9–50. Springer Berlin Heidelberg, Berlin, Heidelberg, 1998.
  - [127] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521:436–444, 2015.
  - [128] H.-C. Lee, E. Breneman, and C. Schulte. Modeling light reflection for computer color vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(4):402–409, 1990.
  - [129] J. E. Lennard-Jones. Cohesion. *Proceedings of the Physical Society*, 43(5):461–482, Sept. 1931.
  - [130] T. Lewiner, H. Lopes, A. W. Vieira, and G. Tavares. Efficient Implementation of Marching Cubes’ Cases with Topological Guarantees. *Journal of Graphics Tools*, 8(2):1–15, Jan. 2003.
  - [131] C. Li, D. Pickup, T. Saunders, D. Cosker, D. Marshall, P. Hall, and P. Willis. Water surface modeling from a single viewpoint video. *IEEE transactions on visualization and computer graphics*, 19(7):1242–1251, 2013.
  - [132] C.-H. Lin, C. Kong, and S. Lucey. Learning efficient point cloud generation for dense 3D object reconstruction. In *AAAI Conference on Artificial Intelligence (AAAI)*, volume 32, 2018.
  - [133] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie. Feature pyramid networks for object detection. *arXiv:1612.03144*, 2017.
  - [134] Z. Liu, K. Yuan, Y. Liu, Y. Qiu, and G. Lei. Fluid sloshing thermo-mechanical characteristic in a cryogenic fuel storage tank under different gravity acceleration levels. *International Journal of Hydrogen Energy*, 47(59):25007–25021, July 2022.
  - [135] D. Lohse. Fundamental fluid dynamics challenges in inkjet printing. *Annu. Rev. Fluid Mech.*, 54(1):349–382, 2022.
  - [136] W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. *ACM SIGGRAPH Computer Graphics*, 21(4):163–169, 1987.
  - [137] J. Lu and G. Tryggvason. Effect of bubble deformability in turbulent bubbly upflow in a vertical channel. *Physics of Fluids*, 20(4):040701, Apr. 2008.
  - [138] L. Lu, X. Meng, Z. Mao, and G. E. Karniadakis. DeepXDE: A deep learning library for solving differential equations. *SIAM Review*, 63(1):208–228, 2021.
  - [139] A. L. Maas, A. Y. Hannun, A. Y. Ng, et al. Rectifier nonlinearities improve neural network acoustic models. In *Proceedings of the 30th International Conference on Machine Learning*, page 3. Atlanta, GA, 2013.
  - [140] S. Madani and A. Amirfazli. Oil drop shedding from solid substrates by a shearing liquid. *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, 441:796–806, 2014.
  - [141] Z. Mao, A. D. Jagtap, and G. E. Karniadakis. Physics-informed neural networks for high-speed flows. *Computer Methods in Applied Mechanics and Engineering*, 360:112789, Mar. 2020.
  - [142] A. Marchand, J. H. Weijs, J. H. Snoeijer, and B. Andreotti. Why is surface tension a force parallel to the interface? *American Journal of Physics*, 79(10):999–1008, Oct. 2011.



- [143] A. U. M. Masuk, A. Salibindla, and R. Ni. A robust virtual-camera 3D shape reconstruction of deforming bubbles/droplets with additional physical constraints. *International Journal of Multiphase Flow*, 120:103088, 2019.
- [144] MathWorks. Introduction to deep learning: what are convolutional neural networks?, 2017. Accessed: June 16th, 2020.
- [145] T. Maurer, A. Mebus, and U. Janoske. Water Droplet Motion on an Inclining Surface. In *3rd International Conference on Fluid Flow, Heat and Mass Transfer, Ottawa, Canada, May 2-3, May 2016*.
- [146] D. Meagher. Geometric modeling using octree encoding. *Computer Graphics and Image Processing*, 19(2):129–147, 1982.
- [147] L. Mescheder, M. Oechsle, M. Niemeyer, S. Nowozin, and A. Geiger. Occupancy networks: Learning 3D reconstruction in function space. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4455–4465. IEEE, 2019.
- [148] G. Mie. Beiträge zur optik trüber medien, speziell kolloidaler metallösungen. *Annalen der Physik*, 330(3):377–445, 1908.
- [149] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *Computer Vision – ECCV 2020*, pages 405–421, Cham, 2020. Springer International Publishing.
- [150] G. Minor, P. Oshkai, and N. Djilali. Optical distortion correction for liquid droplet visualization using the ray tracing method: Further considerations. *Measurement Science and Technology*, 18:L23, 2007.
- [151] G. Minor, N. Djilali, D. Sinton, and P. Oshkai. Flow within a water droplet subjected to an air stream in a hydrophobic microchannel. *Fluid Dynamics Research*, 41(4):045506, 2009.
- [152] S. Mirjalili, S. S. Jain, and M. Dodd. Interface-capturing methods for two-phase flows: An overview and recent developments. *Center for Turbulence Research Annual Research Briefs*, 2017(117-135):13, 2017.
- [153] U. Mock, T. Michel, C. Tropea, I. Roisman, and J. Rühle. Drop impact on chemically structured arrays. *Journal of Physics: Condensed Matter*, 17(9):595–605, 2005.
- [154] A. Moreira and M. Panão. Spray-wall impact. In *Handbook of Atomization and Sprays: Theory and Applications*, pages 441–455. Springer, Heidelberg, 2011.
- [155] A. Moreira, A. S. Moita, and M. R. Panão. Advances and challenges in explaining fuel spray impingement: How much of single droplet impact research is useful? *Progress in Energy and Combustion Science*, 36(5):554–580, 2010.
- [156] N. J. W. Morris and K. N. Kutulakos. Dynamic refraction stereo. *IEEE transactions on pattern analysis and machine intelligence*, 33(8):1518–1531, 2011.
- [157] R. F. Mudde. GRAVITY-DRIVEN BUBBLY FLOWS. *Annual Review of Fluid Mechanics*, 37(1):393–423, Jan. 2005.
- [158] D. Murakami, H. Jinnai, and A. Takahara. Wetting Transition from the Cassie–Baxter State to the Wenzel State on Textured Polymer Surfaces. *Langmuir*, 30(8):2061–2067, Mar. 2014.

- 
- [159] V. Nair and G. E. Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, ICML'10, page 807–814, Madison, WI, USA, 2010. Omnipress.
  - [160] A. Newell, K. Yang, and J. Deng. Stacked hourglass networks for human pose estimation. In *Computer Vision – ECCV 2016*, volume 9912 of *Lecture Notes in Computer Science*, pages 483–499. Springer International Publishing, Cham, 2016.
  - [161] M. Niemeyer, L. Mescheder, M. Oechsle, and A. Geiger. Differentiable volumetric rendering: Learning implicit 3D representations without 3D supervision. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3501–3512. IEEE, 2020.
  - [162] W. Nitsche and A. Brunn. *Strömungsmesstechnik*. Springer, Heidelberg, 2006.
  - [163] P. Niyogi and F. Girosi. Generalization bounds for function approximation from scattered noisy data. *Advances in Computational Mathematics*, 10(1):51–80, 1999.
  - [164] W. F. Noh and P. Woodward. Slic (simple line interface calculation). In A. I. van de Vooren and P. J. Zandbergen, editors, *Proceedings of the Fifth International Conference on Numerical Methods in Fluid Dynamics June 28 – July 2, 1976 Twente University, Enschede*, pages 330–340, 1976.
  - [165] C. L. Novak, S. A. Shafer, and R. G. Willson. Obtaining accurate color images for machine-vision research. In M. H. Brill, editor, *Perceiving, Measuring, and Using Color*, volume 1250, pages 54 – 68. International Society for Optics and Photonics, SPIE, 1990.
  - [166] H. M. Nussenzveig. The theory of the rainbow. *Scientific American*, 236(4):116–127, 1977.
  - [167] T. Ondarçuhu and A. Piednoir. Pinning of a contact line on nanometric steps during the dewetting of a terraced substrate. *Nano Letters*, 5(9):1744–1750, 2005.
  - [168] OpenAI. GPT-4 Technical Report, Mar. 2024. arXiv:2303.08774.
  - [169] J. J. Park, P. Florence, J. Straub, R. Newcombe, and S. Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 165–174. IEEE, 2019.
  - [170] M. Pharr, W. Jakob, and G. Humphreys. *Physically based rendering: from theory to implementation*. The MIT Press, Cambridge, Massachusetts, fourth edition edition, 2023.
  - [171] N. Qian. On the momentum term in gradient descent learning algorithms. *Neural Networks*, 12(1):145–151, 1999.
  - [172] Y. Qian, M. Gong, and Y.-H. Yang. Stereo-based 3D reconstruction of dynamic fluid surfaces by global optimization. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6650–6659. IEEE, 2017.
  - [173] R. Qiu, R. Huang, Y. Xiao, J. Wang, Z. Zhang, J. Yue, Z. Zeng, and Y. Wang. Physics-informed neural networks for phase-field method in two-phase flow. *Physics of Fluids*, 34(5), 2022.
  - [174] D. Quéré. Wetting and Roughness. *Annual Review of Materials Research*, 38(1):71–99, Aug. 2008.
  - [175] M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics informed deep learning (part i): Data-driven solutions of nonlinear partial differential equations. *arXiv:1711.10561*, 2017.

- [176] M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- [177] M. Raissi, A. Yazdani, and G. E. Karniadakis. Hidden fluid mechanics: Learning velocity and pressure fields from flow visualizations. *Science (New York, N.Y.)*, 367(6481):1026–1030, 2020.
- [178] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen, and I. Sutskever. Zero-Shot Text-to-Image Generation, Feb. 2021. arXiv:2102.12092.
- [179] G. Riegler, A. O. Ulusoy, and A. Geiger. OctNet: Learning deep 3D representations at high resolutions. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6620–6629. IEEE, 2017.
- [180] R. Rioboo, C. Tropea, and M. Marengo. Outcomes from a drop impact on solid surfaces. *Atomization and Sprays*, 11(2):155–166, 2001.
- [181] R. Rioboo, M. Marengo, and C. Tropea. Time evolution of liquid drop impact onto solid, dry surfaces. *Experiments in Fluids*, 33:112–124, 2002.
- [182] I. Ríos-López, P. Karamaounas, X. Zabulis, M. Kostoglou, and T. D. Karapantsios. Image analysis of axisymmetric droplets in wetting experiments: A new tool for the study of 3D droplet geometry and droplet shape reconstruction. *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, 553:660–671, 2018.
- [183] P. Roach, N. J. Shirtcliffe, and M. I. Newton. Progress in superhydrophobic surface development. *Soft Matter*, 4(2):224–240, 2008.
- [184] H. Robbins and S. Monro. A Stochastic Approximation Method. *The Annals of Mathematical Statistics*, 22(3):400–407, Sept. 1951.
- [185] I. V. Roisman, R. Rioboo, and C. Tropea. Normal impact of a liquid drop on a dry surface: model for spreading and receding. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 458(2022):1411–1430, 2002.
- [186] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-Resolution Image Synthesis with Latent Diffusion Models. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10674–10685, New Orleans, LA, USA, June 2022. IEEE.
- [187] O. Ronneberger, P. Fischer, and T. Brox. U-Net: Convolutional networks for biomedical image segmentation. In N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241, Cham, 2015. Springer International Publishing.
- [188] M. J. Rosen. *Surfactants and Interfacial Phenomena*. Wiley, 1 edition, July 2004.
- [189] A. Roth, E. Kristensson, and E. Berrocal. Snapshot 3D reconstruction of liquid surfaces. *Optics express*, 28(12):17906–17922, 2020.
- [190] G. Ryskin and L. G. Leal. Numerical solution of free-boundary problems in fluid mechanics. Part 1. The finite-difference technique. *Journal of Fluid Mechanics*, 148:1–17, Nov. 1984.
- [191] S. Saito, Z. Huang, R. Natsume, S. Morishima, A. Kanazawa, and H. Li. PIFu: pixel-aligned implicit function for high-resolution clothed human digitization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019.

- 
- [192] S. Saito, T. Simon, J. Saragih, and H. Joo. PIFuHD: Multi-Level Pixel-Aligned Implicit Function for High-Resolution 3D Human Digitization. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 81–90, Seattle, WA, USA, June 2020. IEEE.
  - [193] O. E. Sala, F. Stuart Chapin, Iii, J. J. Armesto, E. Berlow, J. Bloomfield, R. Dirzo, E. Huber-Sanwald, L. F. Huenneke, R. B. Jackson, A. Kinzig, R. Leemans, D. M. Lodge, H. A. Mooney, M. Oesterheld, N. L. Poff, M. T. Sykes, B. H. Walker, M. Walker, and D. H. Wall. Global Biodiversity Scenarios for the Year 2100. *Science*, 287(5459):1770–1774, Mar. 2000.
  - [194] N. Samkhaniani, H. Marschall, A. Stroh, B. Frohnepfel, and M. Wörner. Numerical simulation of drop impingement and bouncing on a heated hydrophobic surface. *Journal of Physics: Conference Series*, 2116(1):012073, Nov. 2021.
  - [195] D. Sarkar. *Thermal Power Plant*. Elsevier, 2015.
  - [196] L. H. Saw, Y. J. King, M. C. Yew, T. Ching Ng, W. T. Chong, and N. A. Pambudi. Feasibility study of mist cooling for lithium-ion battery. *Energy Procedia*, 142:2592–2597, Dec. 2017.
  - [197] D. Scherer, A. Müller, and S. Behnke. Evaluation of pooling operations in convolutional architectures for object recognition. In K. Diamantaras, W. Duch, and L. S. Iliadis, editors, *Artificial Neural Networks – ICANN 2010*, pages 92–101, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg. ISBN 978-3-642-15825-4.
  - [198] A. D. Schleizer and R. T. Bonnecaze. Displacement of a two-dimensional immiscible droplet adhering to a wall in shear and pressure-driven flows. *Journal of Fluid Mechanics*, 383:29–54, 1999.
  - [199] H. Schlichting and K. Gersten. *Boundary-Layer theory*. Springer, Berlin Heidelberg, ninth edition, 2017.
  - [200] G. K. Seevaratnam, H. Ding, O. Michel, J. Heng, and O. K. Matar. Laminar flow deformation of a droplet adhering to a wall in a channel. *Chemical Engineering Science*, 65(16):4523–4534, 2010.
  - [201] G. S. Settles. *Schlieren and Shadowgraph Techniques*. Springer Berlin Heidelberg, 2001.
  - [202] S. A. Shafer. Describing light mixtures through linear algebra. *Journal of the Optical Society of America*, 72(2):299–300, 1982.
  - [203] S. A. Shafer. Optical phenomena in computer vision. Technical Report TR135, Carnegie-Mellon University, Computer Science Department, Pittsburgh, Pa., 1984.
  - [204] P. Sharma, W. T. Chung, B. Akoush, and M. Ihme. A review of physics-informed machine learning in fluid mechanics. *Energies*, 16(5):2343, 2023.
  - [205] Y. Shin, J. Darbon, and G. E. Karniadakis. On the convergence of physics informed neural networks for linear second-order elliptic and parabolic type PDEs. *Communications in Computational Physics*, 28(5):2042–2074, June 2020.
  - [206] C. Shorten and T. M. Khoshgoftaar. A survey on Image Data Augmentation for Deep Learning. *Journal of Big Data*, 6(1):60, Dec. 2019.
  - [207] A. Shrivastava, T. Pfister, O. Tuzel, J. Susskind, W. Wang, and R. Webb. Learning from simulated and unsupervised images through adversarial training. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2242–2251, 07 2017.

- [208] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556*, 2014.
- [209] S. Solomon, G.-K. Plattner, R. Knutti, and P. Friedlingstein. Irreversible climate change due to carbon dioxide emissions. *Proceedings of the National Academy of Sciences*, 106(6): 1704–1709, Feb. 2009.
- [210] J. H. Spurk and N. Aksel. *Fluid Mechanics*. Springer, Cham, 3rd ed edition, 2020.
- [211] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(56):1929–1958, 2014.
- [212] S. Succi. *The Lattice Boltzmann Equation for Fluid Dynamics and Beyond*. Oxford University PressOxford, June 2001.
- [213] C. Sun, A. Shrivastava, S. Singh, and A. Gupta. Revisiting Unreasonable Effectiveness of Data in Deep Learning Era. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 843–852, Venice, Oct. 2017. IEEE.
- [214] M. Sussman, P. Smereka, and S. Osher. A Level Set Approach for Computing Solutions to Incompressible Two-Phase Flow. *Journal of Computational Physics*, 114(1):146–159, Sept. 1994.
- [215] I. Sutskever, O. Vinyals, and Q. V. Le. Sequence to sequence learning with neural networks. *arXiv:1409.3215*, 2014.
- [216] D. F. Swinehart. The beer-lambert law. *Journal of Chemical Education*, 39(7):333, 1962.
- [217] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. *arXiv:1409.4842*, 2014.
- [218] M. Takeda, H. Ina, and S. Kobayashi. Fourier-transform method of fringe-pattern analysis for computer-based topography and interferometry. *J. Opt. Soc. Am.*, 72(1):156, 1982.
- [219] M. Tancik, P. P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. T. Barron, and R. Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *NeurIPS*, 2020.
- [220] G. Taubin. Estimation of planar curves, surfaces, and nonplanar space curves defined by implicit equations with applications to edge and range image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(11):1115–1138, 1991.
- [221] G. Taubin. Curve and surface smoothing without shrinkage. In *Proceedings of IEEE International Conference on Computer Vision*, pages 852–857, Cambridge, MA, USA, 1995. IEEE Comput. Soc. Press.
- [222] A. Theodorakakos, T. Ous, M. Gavaises, J. M. Nouri, N. Nikolopoulos, and H. Yanagihara. Dynamics of water droplets detached from porous surfaces of relevance to pem fuel cells. *Journal of Colloid and Interface Science*, 300(2):673–687, 2006.
- [223] L. Thompson. The role of oil detachment mechanisms in determining optimum detergency conditions. *Journal of Colloid and Interface Science*, 163(1):61–73, 1994.
- [224] V. Thoreau, B. Malki, G. Berthome, L. Boulange-Petermann, and J. C. Joud. Physico-chemical and dynamic study of oil-drop removal from bare and coated stainless-steel surfaces. *Journal of Adhesion Science and Technology*, 20(16):1819–1831, 2006.

- 
- [225] T. Tieleman and G. Hinton. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural networks for machine learning*, 4(2):26–31, 2012.
  - [226] A. Tomiyama, G. P. Celata, S. Hosokawa, and S. Yoshida. Terminal velocity of single bubbles in surface tension force dominant regime. *International Journal of Multiphase Flow*, 28(9):1497–1519, 2002.
  - [227] L. S. Tong and Y. S. Tang. *Boiling Heat Transfer and Two-Phase Flow*. Routledge, 2 edition, May 2018.
  - [228] M. Toprak, N. Samkhaniani, and A. Stroh. Drop rebounding on heated micro-textured surfaces. *International Journal of Heat and Mass Transfer*, 227:125498, Aug. 2024.
  - [229] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, A. Rodriguez, A. Joulin, E. Grave, and G. Lample. LLaMA: Open and Efficient Foundation Language Models, Feb. 2023. arXiv:2302.13971.
  - [230] G. Tryggvason, B. Bunner, A. Esmaeeli, D. Juric, N. Al-Rawahi, W. Tauber, J. Han, S. Nas, and Y.-J. Jan. A Front-Tracking Method for the Computations of Multiphase Flow. *Journal of Computational Physics*, 169(2):708–759, May 2001.
  - [231] G. Tryggvason, R. Scardovelli, and S. Zaleski. *Direct Numerical Simulations of Gas-Liquid Multiphase Flows*. Cambridge University Press, 1 edition, Jan. 2001.
  - [232] H. C. van de Hulst. *Light scattering by small particles*. Dover books on physics. LSC Communications and Dover Publications, Inc, New York, unabridged and corrected republication of the work originally published in 1957 by John Wiley & Sons, Inc., n.y edition, 1981.
  - [233] H. C. van de Hulst and R. T. Wang. Glare points. *Applied Optics*, 30(33):4755–4763, 1991.
  - [234] J. D. Van Der Waals. The thermodynamic theory of capillarity under the hypothesis of a continuous variation of density. *Journal of Statistical Physics*, 20(2):200–244, Feb. 1979. Translation by J. S. Rowlingson.
  - [235] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention Is All You Need. *arXiv:1706.03762*, 2017.
  - [236] R. Vinuesa and S. L. Brunton. Enhancing computational fluid dynamics with machine learning. *Nature Computational Science*, 2(6):358–366, June 2022.
  - [237] W. A. Wakeham, M. J. Assael, A. Marmur, J. D. Coninck, T. D. Blake, S. A. Theron, and E. Zussman. Material Properties: Measurement and Data; Contact Angle. In C. Tropea, A. L. Yarin, and J. F. Foss, editors, *Springer Handbook of Experimental Fluid Mechanics*, pages 106–118. Springer, Heidelberg, 2007.
  - [238] J. D. Walker. Multiple rainbows from single drops of water and other liquids. *American Journal of Physics*, 44(5):421–433, 1976.
  - [239] F. Wang and A. Theuwissen. Linearity analysis of a cmos image sensor. *Electronic Imaging*, 29(11):84–90, 2017.
  - [240] H. Wang, Z. Wang, M. Du, F. Yang, Z. Zhang, S. Ding, P. Mardziel, and X. Hu. Score-CAM: Score-Weighted Visual Explanations for Convolutional Neural Networks. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 111–119, Seattle, WA, USA, 2020. IEEE.



- [241] N. Wang, Y. Zhang, Z. Li, Y. Fu, W. Liu, and Y.-G. Jiang. Pixel2Mesh: Generating 3D mesh models from single rgb images. In *Computer Vision – ECCV 2018*, volume 11215 of *Lecture Notes in Computer Science*, pages 55–71. Springer International Publishing, Cham, 2018.
- [242] S. Wang and P. Perdikaris. Deep learning of free boundary and Stefan problems. *Journal of Computational Physics*, 428:109914, Mar. 2021.
- [243] S. Wang, H. Wang, and P. Perdikaris. On the eigenvector bias of Fourier feature networks: From regression to solving multi-scale PDEs with physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 384:113938, Oct. 2021.
- [244] S. Wang, X. Yu, and P. Perdikaris. When and why PINNs fail to train: A neural tangent kernel perspective. *Journal of Computational Physics*, 449:110768, Jan. 2022.
- [245] S. Wang, B. Li, Y. Chen, and P. Perdikaris. PirateNets: Physics-informed Deep Learning with Residual Adaptive Networks. *arXiv:2402.00326*, Feb. 2024.
- [246] X. Wang, B. Xu, Z. Chen, D. Del Col, D. Li, L. Zhang, X. Mou, Q. Liu, Y. Yang, and Q. Cao. Review of droplet dynamics and dropwise condensation enhancement: Theory, experiments and applications. *Advances in colloid and interface science*, 305:102684, 2022.
- [247] R. N. Wenzel. Resistance of solid surfaces to wetting by water. *Industrial & Engineering Chemistry*, 28(8):988–994, 1936.
- [248] R. Westermann, L. Kobbelt, and T. Ertl. Real-time exploration of regular volume data by adaptive reconstruction of isosurfaces. *The Visual Computer*, 15(2):100–111, Apr. 1999.
- [249] F. A. Williams. *Combustion Theory: The Fundamental Theory of Chemically Reacting Flow Systems*. CRC Press, 2 edition, Mar. 2018.
- [250] M. Wörner, N. Samkhaniani, X. Cai, Y. Wu, A. Majumdar, H. Marschall, B. Frohnapfel, and O. Deutschmann. Spreading and rebound dynamics of sub-millimetre urea-water-solution droplets impinging on substrates of varying wettability. *Applied Mathematical Modelling*, 95: 53–73, 2021.
- [251] J. Wu, C. Zhang, T. Xue, W. T. Freeman, and J. B. Tenenbaum. Learning a probabilistic latent space of object shapes via 3D generative-adversarial modeling. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, NIPS’16, pages 82–90, Red Hook, NY, USA, 2016. Curran Associates Inc.
- [252] M. Wörner. Numerical modeling of multiphase flows in microfluidics and micro process engineering: a review of methods and applications. *Microfluidics and Nanofluidics*, 12(6): 841–886, May 2012.
- [253] D. Xia, L. M. Johnson, and G. P. López. Anisotropic Wetting Surfaces with One-Dimensional and Directional Structures: Fabrication Approaches, Wetting Properties and Potential Applications. *Advanced Materials*, 24(10):1287–1302, Mar. 2012.
- [254] H. Xu, W. Zhang, and Y. Wang. Explore missing flow dynamics by physics-informed deep learning: The parameterized governing systems. *Physics of Fluids*, 33(9):095116, Sept. 2021.
- [255] J. Xu, Z. Li, B. Du, M. Zhang, and J. Liu. Reluplex made more practical: Leaky relu. In *2020 IEEE Symposium on Computers and communications (ISCC)*, pages 1–7. IEEE, 2020.



- [256] A. Yarin. Drop impact dynamics: Splashing, spreading, receding, bouncing. . . . *Annual Review of Fluid Mechanics*, 38:159–192, 2005.
- [257] A. L. Yarin, I. V. Roisman, and C. Tropea. *Collision Phenomena in Liquids and Solids*. Cambridge University Press, 2017.
- [258] T. Young. III. An essay on the cohesion of fluids. *Philosophical Transactions of the Royal Society of London*, 95:65–87, Dec. 1805.
- [259] D. Youngs. Time-dependent multi-material flow with large fluid distortion. In M. KW and B. MJ, editors, *Numerical Methods for Fluid Dynamics*, volume 24, pages 273–285. Academic Press, New York, 1982.
- [260] M. Yu and H. Quan. Fluid surface reconstruction based on specular reflection model. *Computer Animation and Virtual Worlds*, 24(5):497–510, 2013.
- [261] P. Yue, J. J. Feng, C. Liu, and J. Shen. A diffuse-interface method for simulating two-phase flows of complex fluids. *Journal of Fluid Mechanics*, 515:293–317, Sept. 2004.
- [262] Y. Zeng, J. Fu, and H. Chao. Learning joint spatial-temporal transformations for video inpainting. In *Computer Vision – ECCV 2020*, volume 12361 of *Lecture Notes in Computer Science*, pages 528–543. Springer International Publishing, Cham, 2020.
- [263] K. Zhang, T. Wei, and H. Hu. An experimental investigation on the surface water transport process over an airfoil by using a digital image projection technique. *Experiments in Fluids*, 56(9), 2015.
- [264] X. Zhang, B. A. Tuna, S. Yarusevych, and S. D. Peterson. Flow development over isolated droplet-inspired shapes. *International Journal of Heat and Fluid Flow*, 88:108756, Apr. 2021.
- [265] L. Zhi-yong, P. Xiao-feng, and X.-D. Wang. Oscillation characteristics of droplets on solid surfaces with air flow. *Heat Transfer—Asian Research*, 35:13–19, 2006.

## List of Conference Contributions

- [DBM<sup>+</sup>22] Maximilian Dreisbach, Sebastian Blessing, Frank Michaux, André Brunn, Alexander Stroh, and Jochen Kriegseis. Three-dimensional detection of a gas-liquid interface by means of oblique colored light sources. In *Proceedings of the 20th International Symposium on Application of Laser and Imaging Techniques to Fluid Mechanics, Lisbon, Portugal, July 11–14, 2022*.
- [DBSK23] Maximilian Dreisbach, Sebastian Blessing, Alexander Stroh, and Jochen Kriegseis. Volumetric reconstruction of drop impact dynamics by means of color-coded glare points and deep learning. In *11th International Conference on Multiphase Flow (ICMF 2023), Kobe, Japan, April 2–7, 2023*.
- [DKK<sup>+</sup>24a] Maximilian Dreisbach, Elham Kiyani, Jochen Kriegseis, George Karniadakis, and Alexander Stroh. Convolutional feature-enhanced physics-informed neural networks for the spatio-temporal reconstruction of two-phase flow dynamics. In *77th DFD Annual Meeting of the Division of Fluid Dynamics (APS DFD Meeting 2024), Salt Lake City, United States of America, November 24–26, 2024*.

- [DKK<sup>+</sup>24b] Maximilian Dreisbach, Elham Kiyani, Jochen Kriegseis, George Karniadakis, and Alexander Stroh. Physics-informed neural networks for the prediction of hidden fluid mechanics in droplet impingement. In *1st European Fluid Dynamics Conference (EFDC1 2024)*, Aachen, Germany, September 16–20, 2024.
- [DLP<sup>+</sup>21] Maximilian Dreisbach, Robin Leister, Matthias Probst, Pascal Friederich, Alexander Stroh, and Jochen Kriegseis. Particle detection by means of machine learning in defocusing PTV. In *14th International Symposium on Particle Image Velocimetry, Illinois Institute of Technology (Illinois Tech)*, August 1–8, 2021.
- [SDK23] Christian Sax, Maximilian Dreisbach, and Jochen Kriegseis. Blasengrößenbestimmung durch Interferometric Particle Imaging (IPI) in Vorwärts-, Seitwärts- und Rückstreuung. In *30. Experimentelle Strömungsmechanik (GALA)*, München, Germany, September 5–7, 2023.

## List of Journal Publications

- [DBB<sup>+</sup>23] Maximilian Dreisbach, Sebastian Blessing, André Brunn, Frank Michaux, Alexander Stroh, and Jochen Kriegseis. Three-dimensional encoding of a gas–liquid interface by means of color-coded glare points. *Experiments in Fluids*, 64(3):53, 2023.
- [DKS24] Maximilian Dreisbach, Jochen Kriegseis, and Alexander Stroh. Spatio-temporal reconstruction of droplet impingement dynamics by means of color-coded glare points and deep learning. *Measurement Science and Technology*, 36(1):015304, 2024.
- [DLP<sup>+</sup>22] Maximilian Dreisbach, Robin Leister, Matthias Probst, Pascal Friederich, Alexander Stroh, and Jochen Kriegseis. Particle detection by means of neural networks and synthetic training data refinement in defocusing particle tracking velocimetry. *Measurement Science and Technology*, 33(12):124001, 2022.
- [SDLK23] Christian Sax, Maximilian Dreisbach, Robin Leister, and Jochen Kriegseis. Deep learning and hybrid approach for particle detection in defocusing particle tracking velocimetry. *Measurement Science and Technology*, 34(9):095909, 2023.
- [SSL<sup>+</sup>23a] Christian Sax, Alexander Stroh, Robin Leister, Christian Denda, Patrick Bürk, Maximilian Dreisbach, and Jochen Kriegseis. Fluid-mechanical evaluation of different clutch geometries based on experimental and numerical investigations. *Forschung im Ingenieurwesen*, 87(4):1297–1306, 2023.
- [SSL<sup>+</sup>23b] Christian Sax, Alexander Stroh, Robin Leister, Maximilian Dreisbach, Jochen Kriegseis, Christian Denda, and Patrick Bürk. Fluid-mechanical evaluation of different clutch geometries based on experimental and numerical investigations. In VDI Wissensforum GmbH, editor, *Kupplungs- und Bremssysteme für mobile und stationäre Anwendungen*, pages 109–124. VDI Verlag, 2023.

## List of Preprints

- [DHK<sup>+</sup>25] Maximilian Dreisbach, Itzel Hinojos, Jochen Kriegseis, Alexander Stroh, and Sebastian Burgmann. Interface reconstruction of adhering droplets for distortion correction using glare points and deep learning, 2025. arXiv:2501.03453.

- [DKK<sup>+</sup>24] Maximilian Dreisbach, Elham Kiyani, Jochen Kriegseis, George Karniadakis, and Alexander Stroh. PINNs4Drops: Convolutional feature-enhanced physics-informed neural networks for reconstructing two-phase flows, 2024. arXiv:2411.15949.

## Code and Data Repositories

- [280] Maximilian Dreisbach. Droplet-PIFu: Interface reconstruction of adhering droplets for distortion correction using glare points and deep learning (code), 2025. GitHub repository. <https://github.com/MaxDreisbach/Droplet-PIFu>.
- [281] Maximilian Dreisbach. GPS-Processing: Image processing for glare-point shadowgraphy experiments (code), 2025. GitHub repository. <https://github.com/MaxDreisbach/GPS-Processing>.
- [282] Maximilian Dreisbach. RenderGPS: Synthetic glare-point shadowgraphy data generation in blender with the luxcore package (code), 2025. GitHub repository. <https://github.com/MaxDreisbach/RenderGPS>.
- [283] Maximilian Dreisbach. Spatio-temporal interface reconstruction by means of glare points and deep learning (research data), 2025. KITopen, Karlsruhe Institute of Technology. <https://dx.doi.org/10.35097/mmnxkbqqeye8p5tx>.
- [284] Maximilian Dreisbach, Itzel Hinojos, Jochen Kriegseis, Alexander Stroh, and Sebastian Burgmann. Interface reconstruction of adhering droplets for distortion correction using glare points and deep learning (research data), 2025. KITopen, Karlsruhe Institute of Technology. <https://dx.doi.org/10.35097/egqrfznmr9yp2s7f>.
- [285] Maximilian Dreisbach and Elham Kiyani. PINNs4Drops: Convolutional feature-enhanced physics-informed neural networks for reconstructing two-phase flows (code), 2025. GitHub repository. <https://github.com/MaxDreisbach/PINNs4Drops>.
- [286] Maximilian Dreisbach, Jochen Kriegseis, and Alexander Stroh. Spatio-temporal reconstruction of droplet impingement dynamics by means of color-coded glare points and deep learning (research data), 2024. KITopen, Karlsruhe Institute of Technology. <https://dx.doi.org/10.35097/AcElpeTrdkOvxYWf>.

## Co-Supervised Theses

- [Fü23] Noah Fürst. Experimental investigation of bubble entrapment during drop impact on smooth substrates. Bachelors thesis, ISTM, KIT, 2023.
- [Gad22] Ubada Gad. 3D detection of a gas-liquid interface by means of colored light sources at varying elevation and azimuth angles. Bachelors thesis, ISTM, KIT, 2022.
- [Har23] Kilian Harlos. Experimental investigation of regime boundaries for a drop impingement on flat and structured surfaces. Bachelors thesis, ISTM, KIT, 2023.
- [HS25] Itzel Alejandra Hinojos Salayandia. Development of a digital twin for optical experiments based on shadowgraphy and glare points. Masters thesis, ISTM, KIT, 2025.
- [Luy23] Antonius Luy. Extension of a drop impact test-rig by lateral glare point for multi-view RGB-shadowgraphy. Bachelors thesis, ISTM, KIT, 2023.

- [Pra23] Daniel Prado. Image processing techniques for two-phase flows in microchannels. Bachelors thesis, ISTM, KIT, 2023.
- [Sax22] Christian Sax. Particle detection by means of deep learning and conventional algorithms for defocusing particle tracking velocimetry data. Masters thesis, ISTM, KIT, 2022.





# Acronyms and symbols

## Acronyms

<b>1D</b>	One-Dimensional
<b>2D</b>	Two-Dimensional
<b>2D2C</b>	Two-Dimensional Two-Component
<b>3D</b>	Three-Dimensional
<b>AD</b>	Automatic Differentiation
<b>Adam</b>	Adaptive Moment estimation
<b>ASC</b>	Analytical Spectral Correction
<b>CFD</b>	Computational Fluid Dynamics
<b>CMOS</b>	Complementary Metal Oxide Semiconductor
<b>CNN</b>	Convolutional Neural Network
<b>CSF</b>	Continuum Surface Force
<b>DF</b>	Droplet impingement on Flat substrates
<b>DFS</b>	Droplet impingement on Flat and Structured substrates
<b>DFSSh</b>	Droplet impingement on Flat and Structured substrates and droplets in external Shear flows
<b>DL</b>	Deep Learning
<b>DNS</b>	Direct Numerical Simulation
<b>DS</b>	Droplet impingement on Structured substrates
<b>DSh</b>	Droplets in external Shear flows
<b>FC</b>	Fully Connected
<b>FDM</b>	Fused Deposition Modeling
<b>fps</b>	Frames Per Second
<b>GAN</b>	Generative Adversarial Network
<b>GO</b>	Geometrical Optics
<b>GP</b>	Glare Points
<b>GT</b>	Ground Truth
<b>GPU</b>	Graphics Processing Units



<b>IOU</b>	Intersection Over Union
<b>IPI</b>	Interferometric Particle Imaging
<b>ISC</b>	In-Situ color Correction
<b>LED</b>	Light-Emitting Diode
<b>LIF</b>	Laser Induced Fluorescence
<b>LMT</b>	Lorenz-Mie Theory
<b>LSTM</b>	Long Short-Term Memory network
<b>ML</b>	Machine Learning
<b>MLP</b>	Multi-Layer Perceptron
<b>MSE</b>	Mean Squared Error
<b>NIR</b>	Neutral Interface Reflection
<b>NSE</b>	Navier-Stokes Equations
<b>PBR</b>	Physically Based Rendering
<b>PDE</b>	Partial Differential Equation
<b>PDMS</b>	Polydimethylsiloxane
<b>PEMFC</b>	Proton-Exchange Membrane Fuel Cell
<b>PIFu</b>	Pixel-aligned Implicit Function
<b>PINNs</b>	Physics-Informed Neural Networks
<b>PIV</b>	Particle Image Velocimetry
<b>PLA</b>	Poly lactide
<b>PLIC</b>	Piecewise Linear Interface Calculation
<b>PMMA</b>	Polymethylmethacrylate
<b>PTV</b>	Particle Tracking Velocimetry
<b>px</b>	pixels
<b>ReLU</b>	Rectified Linear unit
<b>R</b>	Reconstruction
<b>RGB</b>	Red Green Blue
<b>RMSPProp</b>	Root Mean Square Propagation
<b>Score-CAM</b>	Score-based Class Activation Map
<b>SiO<sub>x</sub></b>	Silicon Oxide
<b>SDF</b>	Signed Distance Function
<b>SGD</b>	Stochastic Gradient Descent
<b>SGDM</b>	Stochastic Gradient Descent with momentum

<b>SMG</b>	Spatially Modulated Gradient
<b>SPD</b>	Spectral Power Distribution
<b>VAE</b>	Variational Autoencoder
<b>VoF</b>	Volume-of-Fluid method

### Latin letter -Upper case

$A$	area
$A_r$	aspect ratio
$A_{\parallel\perp}$	glare point amplitude of parallelly or perpendicularly polarized light
$B$	bias parameter of neuron or neural network
$Bo$	Bond number
$C$	order parameter
$\mathcal{C}$	correction matrix
$Ca$	Capillary number
$D$	diameter
$\mathcal{D}$	dataset
$E$	energy
$F$	force
$F_{\text{mix}}$	mixing energy
$\mathcal{F}$	convolutional features
$Fr$	Froude number
$H$	parameters of convolutional layer
$I$	intensity
$I_A$	intensity of activation map
$L$	characteristic length
$L(\lambda)$	spectral power distribution of light
$\mathcal{L}$	loss function
$L_1$	L1 distance
$L_2$	L2 distance
$M$	mobility parameter
$N_A$	numerical aperture
$R$	area ratio
$Re$	Reynolds number

$R_p$	reproduction scale
$\mathcal{R}$	residual
$S(Z)$	sigmoid function
$S(\lambda)$	spectral sensitivity of camera
$S_{\parallel\perp}$	electrical field amplitude of parallelly or perpendicularly polarized light
$T$	time
$T(\lambda)$	spectral transmittance
$\mathcal{T}$	transfer matrix
$U$	characteristic velocity
$V$	volume
$W$	weight parameters of neuron or neural network
$W^*$	updated weight parameter
$X$	input of neuron or neural network
$Y$	ground truth output
$\hat{Y}$	output of neuron or neural network
$Z$	weighted inputs of neuron
$We$	Weber number

### Latin letter - Lower case

$a$	droplet radius
$a_{\{1,2,3\}}$	ellipsoidal droplet half-axis
$a_s$	adaptive activation coefficient
$a(Z)$	activation function
$b$	lens radius
$d$	diameter
$d_0$	initial droplet diameter
$d_g$	direction of green light ray
$d_l(x)$	direct light transport coefficient
$d_r$	direction of red light ray
$f$	force
$f_r$	frame rate
$f_{\text{mix}}$	free energy of the interface
$g$	gravitational acceleration

$g_0$	initial condition
$g_\Gamma$	boundary condition
$h$	hidden layer
$k_j$	volume force
$l$	length
$l_c$	capillary length
$l_{HG}$	level of hourglass network
$l(\lambda, x)$	light transport coefficient
$n$	number of quantity
$n_e$	unit vector
$n_i$	interface unit vector
$n_r$	index of refraction
$n_s$	adaptive activation scale factor
$p$	pressure
$p_a$	position of aperture
$p_{gp}$	order of glare point
$p_i$	position on image plane
$p_s$	position of surface points
$q$	predicted quantity
$r$	radius
$r_d(\lambda, x)$	interface reflectance coefficient
$r_f(x, \lambda)$	refraction coefficient
$r_l(\lambda, x)$	spectral reflectance coefficient
$r_s(\lambda, x)$	body reflectance coefficient
$s$	distance
$s(\lambda)$	spectral responsivity of camera
$t$	time
$t^*$	dimensionless time
$t_k$	transfer coefficient
$t_l(\lambda, x)$	transmission coefficient
$u$	velocity
$u_0$	impact velocity
$u_B$	bulk velocity

$u_{\text{CL}}$	contact line velocity
$\hat{u}$	prediction of physics-informed neural network
$v$	velocity
$w$	velocity
$w$	dimensionless droplet coordinate
$x$	spatial coordinate
$x_s$	scattering size parameter
$y$	spatial coordinate
$z$	spatial coordinate

### Greek letters - Upper case

$\Gamma$	parameter of leaky ReLU
$\Delta_{\text{B}}$	relative position of blue glare point
$\Delta_{\text{G}}$	relative position of green glare point
$\Delta_{\text{R}}$	relative position of red glare point
$\Delta_{\text{s}}$	displacement of droplet
$\Delta\Theta$	contact angle hysteresis
$\Theta$	contact angle
$\Theta_{\text{adv}}$	advancing contact angle
$\Theta_{\text{CB}}$	contact angle in Cassie-Baxter state
$\Theta_{\text{eq}}$	equilibrium contact angle
$\Theta_{\text{rec}}$	receding contact angle
$\Theta_{\text{W}}$	contact angle in Wenzel state
$\Phi$	elevation angle
$\Omega$	domain
$\Omega_{\text{g}}$	liquid domain
$\Omega_{\text{l}}$	gaseous domain

### Greek letters - Lower case

$\alpha$	volume fraction
$\alpha_i$	angle
$\alpha_{\text{lr}}$	learning rate
$\beta$	angle

$\gamma$	surface tension coefficient
$\delta$	bias error
$\epsilon$	capillary width
$\varepsilon_{ij}$	strain rate tensor
$\zeta$	unknown parameters in partial differential equation
$\theta$	scattering angle
$\theta_c$	rainbow angle
$\kappa$	interface curvature
$\lambda$	wave length of light
$\lambda_{\text{mix}}$	magnitude of mixing energy
$\mu$	dynamic viscosity
$\mu_g$	dynamic viscosity of gas
$\mu_M$	dynamic viscosity of fluid mixture
$\mu_l$	dynamic viscosity of liquid
$\xi$	material property of fluid
$\rho$	density
$\rho_g$	density of gas
$\rho_M$	density of fluid mixture
$\rho_l$	density of liquid
$\sigma$	uncertainty
$\sigma_{ij}$	stress tensor
$\tau$	angle of incident light ray
$\tau'$	angle of refracted light ray
$\tau_{ij}$	viscous stress tensor
$\varphi$	angle
$\phi$	spatial phase distribution
$\chi$	indicator function
$\psi$	chemical potential
$\omega$	weighting coefficient

### Mathematical operators

$ \dots $	absolute value
$\ \dots\ $	Euclidean norm

$\partial(\dots)$	partial derivative
$D(\dots)$	material derivative
$\delta_{ij}$	Kronecker delta
$\Delta$	difference
$\nabla$	Laplace operator
<b>a</b>	vector/tensor (bold print)
$(\dots)^\mathsf{T}$	vector/matrix transpose
i	imaginary unit



# List of Figures

2.1	Static wetting of an ideal solid surface by a liquid droplet. . . . .	9
2.2	Regimes of droplet impingement. . . . .	11
2.3	Visualization of an adhering droplet subjected to an external shear flow. . . . .	12
2.4	Sketch of the parallel light focused shadowgraphy. . . . .	18
2.5	Scattering of light at a spherical water droplet in air in the framework of geometrical optics. . . . .	19
2.6	Schematic of a fully-connected feedforward neural network. . . . .	22
2.7	Schematic representation of the process inside a neuron. . . . .	22
2.8	Comparison of the hyperbolic tangent, sigmoid, ReLU, and leaky ReLU activation functions. . . . .	23
2.9	Spatial convolution of a 3x3 filter over a 5x5 input. . . . .	26
2.10	Schematic of a simple convolutional neural network. . . . .	27
2.11	Schematic of a residual block. . . . .	27
2.12	Schematic of the <i>U-net</i> architecture. . . . .	28
2.13	Schematics of physics-informed neural networks. . . . .	29
3.1	Overview of the reconstruction framework. . . . .	39
3.2	Sketch of the PIFu neural network architecture. . . . .	40
4.1	Sketch and photograph of the experimental setup of the applied glare-point shadowgraphy method. . . . .	43
4.2	Spectral power distribution of the LED light sources, spectral sensitivity of the camera, and transmittance of the lenses. . . . .	45
4.3	Raw color-imaging of an impinging water droplet on a substrate with SiOx-coating. . . . .	48
4.4	Processed color-imaging of an impinging water droplet on a substrate with SiOx-coating. . . . .	48
4.5	Mean image intensities of the unprocessed color channels in comparison to the ASC-corrected color images. . . . .	49
4.6	Processed color-imaging with overexposure. . . . .	49
4.7	Red image channel of the frame with a spherical water droplet before impact for different azimuth angles. . . . .	50
4.8	Glare-point intensity measured from the ISC corrected red channel from a recording of a spherical droplet in the experiments. . . . .	50
4.9	Mie glare-point images calculated by means of Equation (2.36). . . . .	50
4.10	Comparison of the glare-point images at particular scattering angles. . . . .	50
4.11	Mean image intensities of the color-corrected red and green image channels. . . . .	52
4.12	Frame with the lowest glare point intensity at different elevation angles. . . . .	52
4.13	Renderings of an ellipsoidal cap. . . . .	54
4.14	Geometrical description of the aspherical droplet and definition of the relative glare point positions. . . . .	55
4.15	Positions of the red, green, and blue glare points relative to the shadowgraph contour of the droplet. . . . .	56
5.1	Optical setup modeled in the rendering environment <i>Blender</i> . . . . .	62
5.2	Comparison of synthetic images and glare-point shadowgraphy recordings of spherical droplets at various scattering angles. . . . .	64

5.3	Comparison of the glare-point intensity between the synthetically rendered images and images recorded in the experiments. . . . .	65
5.4	Theoretical intensity of the glare points according to Lorenz-Mie theory. . . . .	65
5.5	Comparison of synthetic images and recordings from the experiment at similar time steps. . . . .	66
6.1	Symmetries of the droplets gas-liquid interface in different experiments. . . . .	69
6.2	Experimental setup for the investigation of adhering droplets subjected to external shear flows and resulting raw recordings. . . . .	71
6.3	Sketch of the measurement setup for experiments involving droplet impingement on patterned substrates with gradients in wettability. . . . .	72
6.4	Side view of the gas-liquid interface at $t = 1.1$ ms, 8.0 ms, and 19.1 ms after droplet impingement on the <i>SMG2</i> -substrate. . . . .	73
6.5	Bottom view of the gas-liquid interface at $t = 1.1$ ms, 8.0 ms, and 19.1 ms after droplet impingement on the <i>SMG2</i> -substrate. . . . .	73
6.6	Image of a droplet that is deposited on a structured substrate, binary mask, and masked image. . . . .	74
6.7	Image of a spheroidal droplet before impact; intersection and union areas of the reconstruction and ground truth. . . . .	75
6.8	Average 3D-IOU during training on the $\mathcal{D}_{DFS}$ and $\mathcal{D}_{DSH}$ synthetic datasets. . . . .	77
6.9	Cross-sections of the volume fraction $\alpha$ ; prediction by the neural network, ground truth, and error. . . . .	78
6.10	Score-CAM feature importance for <i>DFS-GP</i> and <i>DFS</i> on synthetic and real data. . .	80
6.11	Relative feature importance for <i>DFS-GP</i> in the intermittent levels and the final level of the hourglass network. . . . .	80
6.12	Reconstruction results for droplet impingement on the flat SiOx substrate. . . . .	82
6.13	Reconstruction of a frame with self-occlusion of the gas-liquid interface and respective input image. . . . .	83
6.14	Contours of the shadowgraph images for deposited droplets on the structured PLA and PDMS substrates. . . . .	83
6.15	Temporal evolution of the normalized integral volume of the reconstruction for droplet impingement on the structured PLA substrate. . . . .	84
6.16	In-plane and out-of-plane contours of the reconstructed droplet shapes. Figures adapted from [DKS24]. . . . .	85
6.17	Reconstructed three-dimensional gas-liquid interfaces and corresponding experimental input images for droplet impingement on the structured hydrophobic PDMS substrate. . . . .	86
6.18	Temporal evolution of the reconstructed volume for droplet impingement on the structured PDMS substrate. . . . .	86
6.19	Comparison of the depth estimation for a time step with a high volumetric error during the reconstruction of droplet impingement on the structured PDMS substrate by <i>DFS-GP</i> and <i>DFS</i> . . . . .	87
6.20	Reconstruction of the droplet during the wetting of the structured surface by <i>DFS-GP</i> . . . . .	88
6.21	Reconstruction by <i>DFS-GP</i> and the respective input images of droplet impingement with partial rebound and jet. . . . .	88
6.22	Envelopes of the droplet motion in the numerical simulation and in the experiments at different bulk velocities of external flow. . . . .	90
6.23	Temporal evolution of the reconstructed droplet volume for adhering droplets at different bulk velocities of external flow. . . . .	90
6.24	Input images and out-of-plane projection of the respective three-dimensional interfaces for the reconstruction of adhering droplets. . . . .	91
6.25	Temporal evolution of the reconstructed in-plane contour for different bulk velocities of external flow. . . . .	92

6.26	Temporal evolution of the reconstructed out-of-plane contour for different bulk velocities of external flow. . . . .	92
6.27	Temporal evolution of the reconstructed wetted area for different bulk velocities of external flow. . . . .	93
6.28	Average 3D-IOU of the reconstruction for the models trained on data with only green glare points <i>DFS-g</i> , <i>DSh-g</i> , and <i>DFSSh-g</i> on the respective validation datasets during training of the network in comparison to models trained on the same datasets with red and green glare points <i>DFS-GP</i> , <i>DSh-GP</i> , and <i>DFSSh-GP</i> . . .	95
6.29	Temporal evolution of the reconstructed volume for droplet impingement on the patterned PDMS substrate. . . . .	96
6.30	Temporal evolution of bottom view projection of the interface reconstruction by <i>DFS-g</i> for droplet impingement on the <i>SMG2</i> substrate. . . . .	96
6.31	Volumetric reconstruction at different resolution for the same input image from the experiment at $u_B = 8.32 \text{ m/s}$ . . . . .	98
6.32	Smoothed volumetric reconstruction at $512^3$ grid nodes resolution. . . . .	98
6.33	Input images with overexposure and higher-order glare points and the respective reconstruction. . . . .	100
7.1	Schematic of the Convolutional Feature-Enhanced PINN Framework architecture. . .	107
7.2	Evolution of the loss terms during training of the <i>VoF-PINNs</i> . . . . .	114
7.3	Average 3D-IOU of the reconstruction for the validation datasets during training of <i>VoF-PINNs</i> , <i>PF-PINNs-v1</i> , and <i>PF-PINNs-v2</i> in comparison to the baseline <i>DFS-GP2</i> . . .	115
7.4	Phase distribution predicted by <i>VoF-PINNs</i> and <i>PF-PINNs-v2</i> in comparison to the ground truth simulation data, and the absolute error. . . . .	116
7.5	Snapshot from the validation data set and respective in-plane reconstruction by <i>VoF-PINNs</i> . . . . .	117
7.6	Prediction of <i>VoF-PINNs</i> for the horizontal components of the velocity in comparison to the ground truth simulation data, and the absolute error. . . . .	118
7.7	Prediction of <i>VoF-PINNs</i> for the vertical components of the velocity in comparison to the ground truth simulation data, and the absolute error. . . . .	118
7.8	Prediction of <i>VoF-PINNs</i> for the pressure distribution in comparison to the ground truth simulation data, and the absolute error. . . . .	118
7.9	Out-of-plane reconstruction by the <i>DFS-GP</i> and <i>VoF-PINNs</i> for three different recordings from the experiments. . . . .	120
7.10	Temporal evolution of the reconstructed volume for droplet impingement on the structured PDMS substrate with <i>VoF-PINNs</i> and <i>DFS-GP</i> . . . . .	121
7.11	Snapshot from the experiment involving droplet impingement on the structured PLA substrate and iso-contours of the prediction by <i>VoF-PINNs</i> . . . . .	122
A.1	Ellipsoidal droplet projected in its rotation plane . . . . .	168
A.2	Red glare-point positions on ellipsoidal droplets. . . . .	169
A.3	Time series of an impinging water droplet with raw channels. . . . .	170
A.4	Time series of an impinging water droplet with color-corrected channels. . . . .	171
B.1	Modified rendering setup for adhering droplets in external shear flows. . . . .	173
B.2	Adhering droplet recorded in the experiments and synthetically rendered image. . . .	174
C.1	Temporal evolution of the reconstructed volume for adhering droplets. . . . .	175
C.2	Cross-sections of the predicted volume fraction $\alpha$ for adhering droplets. . . . .	176
C.3	Photographs of the substrates with structured surface, 3D-printed polylactide and Polydimethylsiloxane. . . . .	177
C.4	Confocal microscopy images of the patterned PDMS substrates with gradients in wettability. . . . .	178

C.5	Snapshot from the validation data set and respective out-of-plane reconstruction by <i>VoF-PINNs</i> . . . . .	180
C.6	Relative $L_1$ and $L_2$ errors of the predicted velocity and pressure fields. . . . .	181
C.7	Absolute $L_1$ and $L_2$ errors of the predicted velocity and pressure fields. . . . .	181

## List of Tables

5.1	Overview of the different training datasets used throughout this work. . . . .	68
6.1	Uncertainty and bias error of the reconstructed integral volume for <i>DFS-GP</i> and <i>DFS</i> . . . . .	82
6.2	Uncertainty of the reconstructed integral volume for different external flow velocities. . . . .	91
6.3	Average time required for the reconstruction of one time step at different output resolutions. . . . .	98
C.1	Contact angles of the PLA and PDMS-substrates in parallel and transversal direction. . . . .	177
C.2	Displacement $\Delta_s$ of the droplet during impingement on the gradiented patterned substrates. . . . .	177
C.3	Uncertainty $\sigma_V$ and bias error $\delta_V$ of the reconstructed integral volume for <i>DFS-GP</i> and <i>DFS</i> . . . . .	178
C.4	Uncertainty $\sigma_V$ of the reconstructed integral volume for <i>DFS2024D</i> , <i>DSFSH2024D</i> , and <i>DSH2024D</i> . . . . .	179
C.5	Bias error $\delta_V$ of the reconstructed integral volume for <i>DFS2024D</i> , <i>DSFSH2024D</i> , and <i>DSH2024D</i> . . . . .	179
C.6	Absolute and relative $L_1$ and $L_2$ errors of the predicted flow quantities $u, v, w$ and $p$ by <i>VoF-PINNs</i> . Table adopted from [DKK <sup>+</sup> 24]. . . . .	180
C.7	Absolute and relative $L_1$ and $L_2$ errors of the predicted flow quantities $u, v, w$ and $p$ by <i>PF-PINNs-v2</i> . Table adopted from [DKK <sup>+</sup> 24]. . . . .	180
C.8	Uncertainty $\sigma_V$ and bias error $\delta_V$ of the reconstructed integral volume for <i>DFS-GP</i> and <i>VoF-PINNs</i> . . . . .	182



## A Appendix: Glare-point shadowgraphy

### A.1 Analytical glare-point derivation

“In this section, the analytical derivation of the position of a zeroth order glare point on an ellipsoidal droplet that is arbitrarily rotated around its third primary half-axis is detailed.

Points on the ellipsoid can be parameterized as

$$\begin{aligned}\mathbf{p}_s(\alpha, \beta) &= \mathbf{a}_0 + \mathbf{a}_1 \cos \alpha \cos \beta \\ &+ \mathbf{a}_2 \cos \alpha \sin \beta + \mathbf{a}_3 \sin \alpha,\end{aligned}\tag{A.1}$$

where  $\mathbf{a}_0$  denotes the ellipsoid’s center point and  $\mathbf{a}_{\{1,2,3\}}$  denote the ellipsoid’s (orthogonal) primary half-axes, while  $\alpha \in [-\frac{\pi}{2}; \frac{\pi}{2}]$  and  $\beta \in [0; 2\pi[$  are the angular coordinates of the point. Furthermore, the normal vector at each point of the ellipsoid surface is known by

$$\begin{aligned}\mathbf{n}_{e,s}(\alpha, \beta) &= \mathbf{a}_2 \times \mathbf{a}_3 \cos \alpha \cos \beta \\ &+ \mathbf{a}_3 \times \mathbf{a}_1 \cos \alpha \sin \beta \\ &+ \mathbf{a}_1 \times \mathbf{a}_2 \sin \alpha.\end{aligned}\tag{A.2}$$

The droplet’s aspect ratio is consequentially defined as

$$A_r = \frac{\|\mathbf{a}_1\|}{\|\mathbf{a}_2\|}.\tag{A.3}$$

The camera recording the scene is modeled as a pinhole aperture and an image plane behind it. The droplet’s center point  $\mathbf{a}_0$  and the known position of the pinhole aperture  $\mathbf{p}_a$  from the optical axis. The image plane has a distance of  $s_i$  from  $\mathbf{p}_a$  in direction of the optical axis and has the normal vector

$$\mathbf{n}_{e,i} = \frac{\mathbf{p}_a - \mathbf{a}_0}{\|\mathbf{p}_a - \mathbf{a}_0\|}.\tag{A.4}$$


The parallel light of color  $c$  producing the glare point is emitted in known direction  $\mathbf{d}_{c,1}$ . The direction of the light ray reflected from the droplet’s surface to the pinhole aperture is denoted

$$\mathbf{d}_{c,0} = \frac{\mathbf{p}_a - \mathbf{p}_{s,c}}{\|\mathbf{p}_a - \mathbf{p}_{s,c}\|},\tag{A.5}$$

where  $\mathbf{p}_{s,c}$  is the position of the glare point of color  $c$  on the droplet’s surface.

A projection of the droplet in its rotational plane including all needed parameters is depicted in Figure A.1, where the third direction as well as the camera model are omitted for easier visualization.





**Figure A.1:** Ellipsoidal droplet projected in its rotation plane with red and green light beams, primary axes  $\mathbf{a}_1$  and  $\mathbf{a}_2$ , and rotation angle  $\varphi$ . Figure adopted from [DKS24].

The reflection condition – i.e. the angle of incidence equals the angle of reflection – can be written as

$$\begin{aligned} \mathbf{d}_{c,0} \cdot \mathbf{n}_{e,s} &= -\mathbf{d}_{c,1} \cdot \mathbf{n}_{e,s} \\ \Leftrightarrow \mathbf{n}_{e,s} \cdot (\mathbf{d}_{c,1} + \mathbf{d}_{c,0}) &= 0, \end{aligned} \quad (\text{A.6})$$

since  $\mathbf{d}_{c,i}$  and the normal vector on the ellipsoid  $\mathbf{n}_{e,s}$  (which corresponds to the point  $\mathbf{p}_{s,c}$ ) all have unit length.

Furthermore, the light direction vectors  $\mathbf{d}_{c,i}$  and the normal vector of the reflection plane must all lie in one plane, and therefore the conditions

$$\mathbf{d}_{c,0} \cdot (\mathbf{d}_{c,1} \times \mathbf{n}_{e,s}) = 0 \quad (\text{A.7})$$

must hold.

Solving the nonlinear system of equations formed by Equation (A.6) and Equation (A.7) yields the angular coordinates  $\alpha$  and  $\beta$  of the point  $\mathbf{p}_{s,c}$  on the ellipsoid surface that reflects the light with the given initial direction into the aperture without crossing the ellipsoid. Since the ellipsoid is a convex body, the solution is unique.

Subsequently, the intersection between the linear equation

$$\mathbf{x} = \mathbf{p}_a + k\mathbf{d}_{c,0} \quad (\text{A.8})$$


and the image plane is given by

$$\mathbf{n}_{e,i} \cdot (\mathbf{p}_i - \mathbf{x}) = 0, \quad (\text{A.9})$$

where  $\mathbf{p}_i$  is an arbitrary point on the image plane, yields the position  $\mathbf{p}_{i,c}$  of the glare-point imprint on the camera sensor.

For a given aspect ratio  $A_r$  and droplet rotation  $\varphi$ , this set of equations can be numerically solved to reproduce the results from the renderings shown in Figure 4.15. The analytical results using the full set of parameters, as also used in the rendering setup, are shown in Figure A.2 for the red glare point.

When the pixel positions of the glare points on the sensor  $\Delta_R$  and  $\Delta_G$  are known, the system can be solved numerically for the aspect ratio  $A_r$  and the rotation angle  $\varphi$ , yielding information about the droplet shape for a known optical setup. Since the position of a single glare point is




**Figure A.2:** Red glare-point positions on ellipsoidal droplets calculated using the set of equations derived above. Figure adopted from [DKS24].


enough to close the nonlinear system of equations, the addition of a second glare point allows for an estimation of measurement errors and consequentially presents a measure for the validity of the chosen assumptions.” from [DBB<sup>+</sup>23, Appendix A]

## A.2 Time series of RGB images

“Comparison of the full series of RGB-frames shortly before and after the drop impact between the raw images (Figure A.3) and the color-corrected images that result from the processing of the raw images with the ISC function [see 4.2](Figure A.4).” from [DBB<sup>+</sup>23, Appendix C]



**Figure A.3:** Time series of an impinging water droplet onto a substrate with SiOx-coating; (a) RGB snap shots, captured with a *Photron Fastcam Nova R2* at 7,500 fps; separated (b) blue, (c) green and (d) red raw channels of (a). Figure adopted from [DBB<sup>+</sup>23].




**Figure A.4:** Time series of an impinging water droplet onto a substrate with SiOx-coating; (a) RGB snap shots, captured with a *Photron Fastcam Nova R2* at 7,500 fps; separated (b) blue, (c) green and (d) red color-corrected channels of (a). Figure adopted from [DBB<sup>+</sup>23].



## B Appendix: Synthetic image rendering

### B.1 Capability of the rendering setup as a digital twin of the experiment


The rendering setup developed for synthetic data generation in this work accurately reproduces the system behavior of the optical setup in the experiments as demonstrated in Section 5.3 and, thereby, serves as a digital twin for the real experiments. The predictive capabilities of this digital twin allow for the testing of different experimental setups in a virtual environment before the construction of a physical test rig, thereby saving time and resources. In this capacity, the rendering setup was employed during the development of the experimental setup for the investigations involving adhering droplets subjected to external shear flows. The wind tunnel necessary to generate the channel flow, in which the droplet was placed, was modeled in the rendering setup to accurately reproduce the dimensions and optical properties of its real counterpart in the experiments. The wind tunnel consisted of acrylic glass (PMMA), which was modeled in the rendering setup as a transparent medium with a constant index of refraction  $n_r = 1.459$ . The modified rendering setup with the included channel is illustrated in Figure B.1.



**Figure B.1:** Modified rendering setup including the channel employed in the experiments involving adhering droplets in external shear flows.

During the testing in the rendering setup, it was found that the lateral light sources needed to be raised in order to illuminate the droplet at a steeper elevation angle of  $\Phi = 60^\circ$  to ensure a consistent glare-point signal during the dynamic deformation of the gas-liquid interface. With these changes implemented, the testing of the modified experimental setup in the rendering environment revealed, that the glare-point shadowgraphy technique was applicable to the experiments involving adhering droplets. It was found that the transmission of light from the backlight and the lateral light sources through the transparent walls of the wind tunnel influenced the intensity of the glare points and the background in the shadowgraph, but not their respective locations. The two-fold refraction of the light emitted by the lateral light sources during the first transmission through walls of the channel at an angle of incidence (see Figure B.1), introduces astigmatism. Since the reflective glare points are not dependent on the translation of the incident light distribution, as previously elaborated in Section 5.2, only the shape of the glare points was influenced by this astigmatism, but

not their location. The second transmission through the walls occurs at an approximately normal incident angle, as the side walls of the wind tunnel are aligned parallel to the object plane, and thus no additional astigmatism arises here. Therefore, the alignment of the wind tunnel and objective lens in the experiments was determined as crucial for imaging with minimal optical aberration. The most relevant effect of the channel was the reduction of the light intensity, which occurred due to the reflection of the incident light at the walls of the channel, and to some minor degree due to scratches and surface roughness. As the calculation of the light interaction with the channel resulted in additional computational effort, instead the intensity of the light sources was modified to reproduce the illumination observed in the experiments without explicitly representing the wind tunnel in the rendering environment for the synthetic data generation. Thereby, the rendering time was significantly reduced. Figure B.2 shows a synthetic image of an adhering droplet subjected to external shear flow rendered on the basis of a gas-liquid interface geometry extracted from the simulation of [DHK<sup>+</sup>25] in comparison to an image recorded in the experiments.




**Figure B.2:** Adhering droplet recorded in the experiments (left) and synthetically rendered image (right).



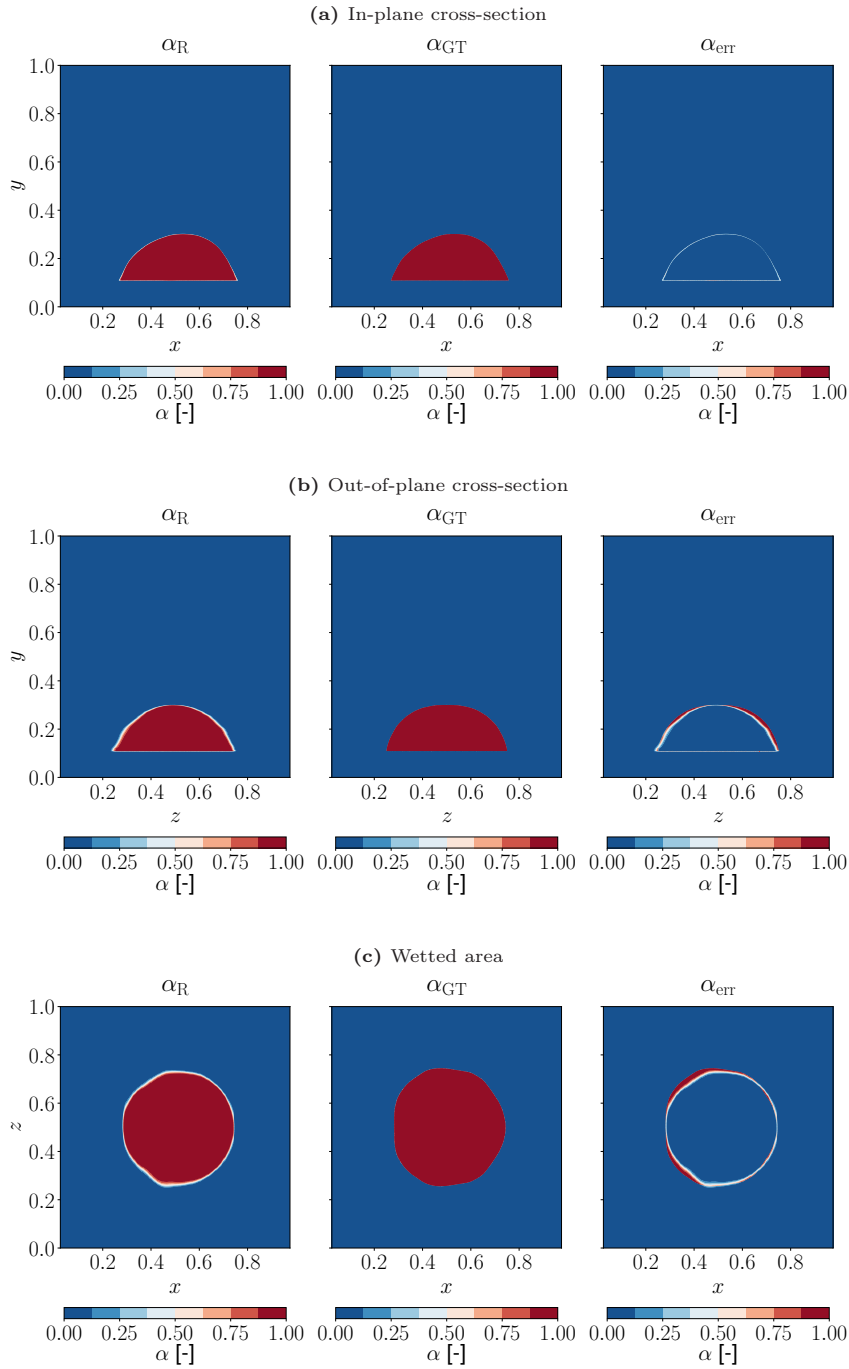
## C Appendix: Spatio-temporal reconstruction

### C.1 Adhering droplets subjected to shear flows




**Figure C.1:** Temporal evolution of the normalized integral volume of the reconstruction for adhering droplets in external shear flows at different bulk velocities  $u_B = 5.85$  m/s,  $u_B = 7.58$  m/s, and  $u_B = 8.32$  m/s. Figure adopted from [DHK<sup>+</sup>25].

“The temporal evolution of the normalized integral volume of the reconstructed droplet is presented in Figure 6.23 for one representative experiment at each evaluated velocity  $u_B$  of the external flow. These results are further detailed in Figure C.1, showing the evolution of the reconstructed droplet volume for all cases at the respective velocities  $u_B = 5.85$  m/s,  $u_B = 7.58$  m/s, and  $u_B = 8.32$  m/s. Figure C.2 shows the reconstructed phase distribution represented by the volume fraction  $\alpha$  in comparison to the ground truth data in three orthogonal cross-sections for one example of the validation data with high errors. As can be seen, the out-of-plane reconstruction exhibits significantly larger errors in comparison to the in-plane reconstruction, which falls in line with the observations of the error distribution for a low error sample in Figure 6.9.” from [DHK<sup>+</sup>25, Appendix]



**Figure C.2:** Cross-sections of the volume fraction  $\alpha$ ; prediction by the neural network (left), ground truth (middle), and deviation of the prediction from the ground truth (right). In subfigures (a) and (c) the main flow direction is from left to right and in subfigure (b) the main flow direction is aligned to the image plane. Figure adopted from [DHK<sup>+</sup>25].

## C.2 Droplet impingement on structured substrates



**Figure C.3:** Photographs of the substrates with structured surface, 3D-printed polylactide (PLA) on the left and Polydimethylsiloxane (PDMS) on the right. The detail (below) shows the side view on the substrate at 0° observation angle. Figure adopted from [DKS24].

**Table C.1:** Contact angles of the PLA and PDMS-substrates in parallel and transversal direction. Table adopted from [DKS24].


case	$\Theta_{adv}$	$\Theta_{rec}$	$\Theta_{eq}$	$\Delta\Theta$
PDMS 0°	115	88	107	27
PDMS 90°	107	74	97	33
PLA 0°	113	59	76	54
PLA 90°	101	52	63	49

The properties of the substrates used for the droplet impingement experiments in Section 6.1.2 are further detailed in the following. Figure C.3 shows photographs of the structured PLA and PDMS substrate in the view from the top and the side, while Table C.1 details the measured contact angles for these substrates.

**Table C.2:** Displacement  $\Delta_s$  of the droplet during impingement on the gradiented patterned substrates and standard deviation of the displacement  $\sigma_s$ . The droplet position was calculated as the mean of the left and right contact line locations detected in the side view shadowgraph.

substrate	$\Delta_s$ [mm]	$\sigma_s$ [mm]	samples
<i>SMG0</i>	1.165	0.069	6
<i>SMG1</i>	0.588	0.166	7
<i>SMG2</i>	1.845	0.169	7
<i>SMG3</i>	1.183	0.142	7
<i>SMG4</i>	-0.028	0.043	7

Figure C.4 shows examples of the patterned substrates with gradients in wettability imaged by confocal microscopy. The main feature of these substrates is the obvious gradient in the surface patterning, which consists of densely packed hemispheres with around  $d_p = 10\mu\text{m}$  diameter transitioning into a flat surface. Further images of the patterned substrates with different lengths of gradients and details on the manufacturing of these substrates can be found in [50]. Table C.2



**Figure C.4:** Confocal microscopy images of the patterned PDMS substrates with gradients in wettability (*SMG*).

presents the measured horizontal displacement  $\Delta_s$  of impinging droplets on these substrates. The displacement of the droplets from the point of impact to the resting position at equilibrium was measured from the detected contours in the side view shadowgraph.

**Table C.3:** Uncertainty  $\sigma_V$  and bias error  $\delta_V$  of the reconstructed integral volume for *DFS-GP* and *DFS* in percent of the ground truth volume for droplet impingement on different substrates. Table adopted from [DKS24].

case $\sigma_V, \delta_V$ (in %)	<i>DFS-GP</i>		<i>DFS</i>	
	$\sigma_V$	$\delta_V$	$\sigma_V$	$\delta_V$
flat	3.5	4.7	5.1	6.2
PLA 0°	7.0	9.9	7.9	11.7
PLA 45°	2.9	0.3	3.2	6.3
PLA 90°	7.2	4.5	6.8	6.1
PDMS 0°	9.3	5.1	9.5	6.9
PDMS 45°	3.9	7.8	2.9	9.3
PDMS 90°	3.7	7.2	4.3	10.0

In extension to the results presented in Sections 6.3.2 and 6.3.3, the accuracy of the reconstructed droplet interfaces during impingement on the polylactide (PLA) and Polydimethylsiloxane (PDMS) substrates on the basis of images recorded at different observation angles is detailed in Table C.3.

**Table C.4:** Uncertainty  $\sigma_V$  of the reconstructed integral volume for *DFS2024D*, *DSFSH2024D*, and *DSH2024D* in percent of the ground truth volume for droplet impingement on different substrates.

case	<i>DFS2024D</i>	<i>DSFSH2024D</i>	<i>DSH2024D</i>
<i>SMG0</i>	6.6	7.5	25.2
<i>SMG1</i>	5.3	5.9	24.8
<i>SMG2</i>	10.5	11.5	43.3
<i>SMG3</i>	8.6	10.2	26.5
<i>SMG4</i>	9.7	12.3	23.9
avg.	8.1	9.5	28.7

**Table C.5:** Bias error  $\delta_V$  of the reconstructed integral volume for *DFS2024D*, *DSFSH2024D*, and *DSH2024D* in percent of the ground truth volume for droplet impingement on different substrates.

case	<i>DFS2024D</i>	<i>DSFSH2024D</i>	<i>DSH2024D</i>
<i>SMG0</i>	2.5	1.0	0.1
<i>SMG1</i>	19.0	22.0	37.0
<i>SMG2</i>	17.6	6.3	19.6
<i>SMG3</i>	4.6	3.4	0.4
<i>SMG4</i>	17.6	25.5	26.5
avg.	12.3	11.6	16.7

The accuracy of the reconstructed droplet interfaces during impingement on the patterned substrates with gradients in wettability is outlined in Tables C.4 and C.5, which detail the uncertainty and bias error, respectively, supplementing the presented results in Section 6.3.6.

### C.3 Physics-informed neural networks

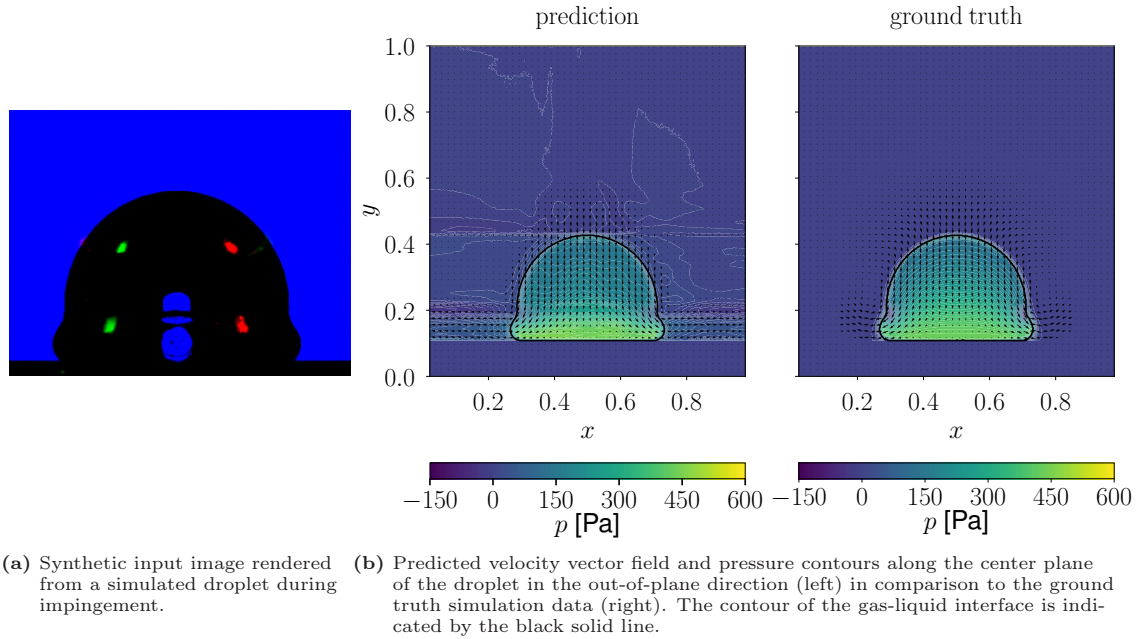
**Table C.6:** Absolute and relative  $L_1$  and  $L_2$  errors of the predicted flow quantities  $u, v, w$  and  $p$  by *VoF-PINNs*.  
Table adopted from [DKK<sup>+</sup>24].

quantity	abs. $L_1$	abs. $L_2$	rel. $L_1$	rel. $L_2$
$u$ [m/s]/[%]	0.0187	0.0500	3.44	7.43
$v$ [m/s]/[%]	0.0229	0.0403	3.26	5.43
$w$ [m/s]/[%]	0.0240	0.0539	4.64	8.71
$p$ [Pa]/[%]	12.81	19.179	4.58	6.77

**Table C.7:** Absolute and relative  $L_1$  and  $L_2$  errors of the predicted flow quantities  $u, v, w$  and  $p$  by *PF-PINNs-v2*.  
Table adopted from [DKK<sup>+</sup>24].


quantity	abs. $L_1$	abs. $L_2$	rel. $L_1$	rel. $L_2$
$u$ [m/s]/[%]	0.0242	0.0548	5.0	9.19
$v$ [m/s]/[%]	0.0344	0.0511	5.01	7.15
$w$ [m/s]/[%]	0.0376	0.0711	7.53	12.39
$p$ [Pa]/[%]	20.095	27.700	7.48	10.12

“The results [for the extension of the reconstruction framework by physics-informed neural networks] presented in Section 7.3 are further elaborated in the following. Figure C.5b shows the out-of-plane pressure and velocity fields in the center plane of the droplet predicted by *VoF-PINNs* (left) in comparison to the ground truth (right) for one snapshot of the synthetic validation dataset displayed in Figure C.5a. Tables C.6 and C.7 detail the relative and absolute  $L_1$  and  $L_2$  errors for the prediction of the velocity and pressure distributions, averaged over the validation dataset for *VoF-PINNs* and *PF-PINNs-v2*, respectively. The evolution of the relative and absolute errors for




**Figure C.5:** Snapshot from the validation data set (a) and respective out-of-plane reconstruction (b) by *VoF-PINNs*. Figure adapted from [DKK<sup>+</sup>24].





**Figure C.6:** Relative  $L_1$  (dashed lines) and  $L_2$  (solid lines) errors of the predicted velocity and pressure fields. Figure adopted from [DKK<sup>+</sup>24].



**Figure C.7:** Absolute  $L_1$  (dashed lines) and  $L_2$  (solid lines) errors of the predicted velocity and pressure fields. Figure adopted from [DKK<sup>+</sup>24].

*VoF-PINNs* on the subset of the validation data that features droplet impingement on a structured surface are plotted in Figures C.6 and C.7, respectively.

The results regarding the accuracy of the reconstruction for droplet impingement experiments with *VoF-PINNs* are detailed in Table C.8. The images recorded in experiments involving the impingement of droplets on the structured PLA and PDMS substrates at different observation angles with respect to the orientation of the surface structure are reconstructed by *VoF-PINNs* and the uncertainty, as well as bias errors of the integral reconstructed volume are compared to the results by [...] [*DFS-GP*].” from [DKK<sup>+</sup>24, Appendix]



**Table C.8:** Uncertainty  $\sigma_V$  and bias error  $\delta_V$  of the reconstructed integral volume for *DFS-GP* and *VoF-PINNs* in percent of ground truth volume for droplet impingement on different substrates. Table adopted from [DKK<sup>+</sup>24].

case $\sigma_V, \delta_V$ (in %)	<i>DFS-GP</i>		<i>VoF-PINNs</i>	
	$\sigma_V$	$\delta_V$	$\sigma_V$	$\delta_V$
PLA 0°	7.0	9.9	2.3	5.1
PLA 45°	2.9	0.3	1.4	1.1
PLA 90°	7.2	4.5	2.2	0.2
PDMS 0°	9.3	5.1	3.3	3.0
PDMS 45°	3.9	7.8	1.3	3.7
PDMS 90°	3.7	7.2	1.8	0.01

## D Appendix: Code and data repositories

All data and code that support the findings of this thesis have been made publicly available. This includes the raw and processed image files from the experiments and the respective code for image processing, code for the generation of synthetic training data, code related to the training, as well as the deployment of the neural networks for interface reconstruction, and the weights of the trained neural networks. In the following chapter, the data and code repositories related to the individual components of the developed reconstruction framework in this thesis are specified.

The raw and processed image data from the glare-point shadowgraphy experiments involving the impingement of liquid droplets on structured polylactide (PLA) and Polydimethylsiloxane (PDMS) substrates used in Section 6.3 are published on KITopen [286]. The image data from droplet impingement experiments on patterned substrates with gradients in wettability are published on KITopen [283], and the image data recorded in the experiments involving adhering droplets subjected to external shear flows are published on KITopen [284]. The processed image files of the aforementioned experiments are provided in a format that can directly be inputted to the trained neural networks to obtain the frame-wise reconstruction of the gas-liquid interface. The code for processing raw recordings from glare-point shadowgraphy experiments, including the color-correction techniques developed in Chapter 4, is published on GitHub [281].

The code for the generation of synthetic training data replicating the recordings from glare-point shadowgraphy experiments by means of physics-based rendering, utilized in Chapter 5, is published on GitHub [282]. The related Blender file, containing the modeled rendering setup, is published on KITopen [281].

The code for the training, deployment, and evaluation of the PIFu neural network, used for the reconstruction of the gas-liquid interface from glare-point shadowgraphy images in Chapter 6, is published on GitHub [280]. The code for the training and evaluation of the physics-informed neural networks (PINNs), developed in Chapter 7, is published on GitHub [285]. Both code repositories include documentation detailing the requirements and the setup of runtime environments, as well as the execution of the codes. The weights of the neural networks and PINNs trained on the different synthetic datasets are published on KITopen [283].